Conservative Holdings, Aggressive Trades: Ambiguity, Learning, and Equilibrium Flows*

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

January 7, 2023

*We thank Ella Dias Saraiva-Patelli, Bernard Dumas, Philipp Illeditsch, Alessandro Melone, Martin Schneider, 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, Italy. E-mail: Alex.Weissensteiner@unibz.it

Conservative Holdings, Aggressive Trades: Ambiguity, Learning, and Equilibrium Flows

Abstract

We propose an equilibrium asset pricing model in which agents learn about the parameters that drive economic fundamentals and differ in their aversion to ambiguity. 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: ambiguity-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. These predictions are consistent with observed portfolio flows of retail and institutional investors around dividend surprises. Our model highlights that heterogeneity of preferences and learning about volatility of fundamentals 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. 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 interactions between heterogeneous agents who, upon the arrival of new information, learn about the fundamentals of the economy. 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 emphasizes that subjective risk premia inferred from investors' return expectations for a variety of asset classes are much less counter-cyclical than objective risk premia identified from in-sample predictive regressions (see, e.g., Nagel and Xu, 2022). Such a discrepancy can be reconciled through representative-agent asset pricing models in which subjective expectations are time varying as in settings with perpetual learning (Collin-Dufresne, Johannes, and Lochstoer, 2016; Nagel and Xu, 2021). While learning in representative-agent models can explain asset pricing puzzles, these models are not designed to study portfolio flows.

In this paper, we develop an equilibrium asset pricing model in which heterogeneous agents learn about both the moments of the endowment process and differ in their aversion to parameter uncertainty. We show that, when some agents are averse to parameter 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 observations. Specifically, we show that while in equilibrium ambiguity-averse agents are more conservative in their holdings, they are more aggressive in their trades. 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 imply that larger dividend realizations are associated with higher expected market risk premia. These results highlight the importance of learning about volatility of fundamentals when interpreting empirical evidence.

[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).

We first illustrate the main mechanism in a simple two-period heterogeneous agent model that we can fully 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 variance *and* volatility of the 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 increase the agents' volatility estimate. When some agents are ambiguity averse, these belief updates generate equilibrium 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 implies gains from trade in which more (less) ambiguity averse agents increase (decrease) their position in the risky asset in times of high uncertainty, i.e., after both positive and negative surprises. In contrast, when no agent is ambiguity averse the equilibrium risk premium is proportional only to the dividend variance, and a standard "no-trade" result emerges where surprises do not generate equilibrium flows.

We then extend the simple model to an infinite-horizon model where overlapping generations of agents learn about the mean and the variance of the endowment process and differ in their attitude towards ambiguity. When both the mean and the variance are not known, the predictive distribution of dividends, which is normal if only the mean is unknown, becomes a fat-tailed Student t distribution. Unfortunately, with t-distributed dividends expected utility might not be well defined, see, e.g., Geweke (2001). We overcome this difficulty by imposing an a-priori restriction to the dividend variance, referring to recent developments in Bayesian learning techniques with truncated distributions (see, e.g., Weitzman, 2007; Bakshi and Skoulakis, 2010). Specifically, we assume that the unknown variance can take values on an arbitrarily large, but finite interval. This assumption implies that the predictive distribution of dividends is a "dampened Student-t"—i.e., a Student-t with thinner tails—and allows us to fully characterize the equilibrium with learning about both mean and variance.

Because in our model the true mean and variance are assumed to be constant parameters, agents might eventually learn these parameters perfectly. When this happens, the role of parameter uncertainty and ambiguity vanishes in the long run. In reality, parameter uncertainty is unlikely to disappear even after observing a long history of data. To capture this realistic feature, we assume the existence of "information leakages" that occurs as generations overlap: new generation partially forget the accumulated knowledge handed over to them by the older generation. In a representative agent economy, this assumption coincides with the idea of "fading memory" as in, e.g., Nagel and Xu (2021), or "age-related experiential learning", as in Malmendier and Nagel (2016), Collin-Dufresne, Johannes, and Lochstoer (2016), and Malmendier, Pouzo, and Vanasco (2020).

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 show that, in line 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, our model shows that (i) heterogeneous attitude towards ambiguity; (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 the aversion to ambiguity 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 that studies asset prices under parameter uncertainty and learning.[3] 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. By emphasizing the importance of modelling the process of learning about volatility, our paper echoes Weitzman (2007, p.1111) who claims that "for asset pricing implications [...] the most critical issue involved in Bayesian learning [...] is the unknown variance". When

[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] 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.

dealing with parameter uncertainty and learning, the vast majority of the asset pricing literature assumes that the mean is unknown, but the variance is known. This assumption is typically motivated by greater analytical tractability and the impression that with a large sample and continuous observations, it is easy to learn the variance. In reality, however, information reaches market participants in a lumpy fashion, such as during FOMC communication events or corporate earning announcements, and agents cannot avoid the effort to learn about volatility.[4]

Second, we contribute to the literature on asset pricing with heterogeneous agents.[5] 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 ambiguity aversion, and we generalize their results to the case of learning about the mean and the variance of the endowment process in an overlapping-generation economy.[6] Buss, Uppal, and Vilkov (2021) study the dynamics of asset demand in a multi-period 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 ambiguity and learn about both the mean and the variance of the dividend process. Because of agents' ambiguity aversion, learning about variance has a first-order effect on both equilibrium flows and asset prices. These effects are instead negligible in a model where agents are ambiguity-neutral or do not differ in their degree of ambiguity aversion.

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

[5]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.

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

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.[7] 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 invoking the "attention-grabbing" hypothesis, which assumes that individual investors have limited attention and rarely sell short.[8] We provide an alternative explanation to the attention-grabbing hypothesis by highlighting the role of agents' ambiguity attitude and learning.

The rest of the paper proceeds as follows. In Section 2 we provide intuition in a simple equilibrium model which is analytically tractable. Section 3 presents an overlapping-generations model with perpetual learning about the mean and variance of the dividend process. Section 4 contains our empirical analysis of the equilibrium flow dynamics. Section 5 concludes. Appendix A contains proofs, Appendix B provides technical details of Bayesian learning with unknown variance, and Appendix C illustrates our numerical procedure to determine the equilibrium.

2 A two-period model

In this section, we develop a simple equilibrium model to illustrate the effect of volatility on portfolio weights and risk premia when agents differ in their ambiguity aversion.

Assets. There are two dates and a single "tree" producing a perishable dividend $\tilde{d}$ at time 1. Agents live for two periods. In the first period, they can trade in claims over the

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

[8] 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).

dividend tree (the risky asset) at a price p and a riskless asset available in infinite supply. In the second period, they consume the dividend from their portfolio. Since consumption occurs only at the terminal date, the riskless rate in the economy is undetermined and assumed to be a constant r.

We assume that the dividend $\tilde{d}$ is normally distributed with unknown mean μ and known variance σ^2, $\tilde{d} \sim \mathcal{N}(\mu, \sigma^2)$.[9] At the initial date, agents have observed a history of t dividend realizations and calculate the time series average m and its standard error

$$s = \frac{\sigma}{\sqrt{t}}.$$

Preferences. The economy is populated by two types of agents, $i = S, A$, both having CARA utility

$$u(W) = -\frac{1}{\gamma} e^{-\gamma W},$$

with identical absolute risk aversion $\gamma > 0$. Agents differ in their attitude towards uncertainty about the mean estimate. Type-S agents are standard subjective expected utility investors. Type-A agents are averse against ambiguity in the estimated mean m.

Agents S use m as their subjective dividend expectation and account for its estimation error by inflating the variance σ^2 by the variance of the mean s^2. Therefore, the predictive distribution of the dividend $\tilde{d}$ for agent S is

$$\tilde{d} \sim^S \mathcal{N}\left(\mu^S, \sigma^2 \left(\frac{t+1}{t}\right)\right), \quad \text{where } \mu^S = m. \tag{1}$$

In contrast, A agents have "multiple priors" about the distribution of $\tilde{d}$. The set of priors is characterized by a "confidence interval" around the mean estimate, m, whose size depends on the standard error s and their aversion towards ambiguity. Specifically, the set of priors that A agents consider is

$$\tilde{d} \sim^A \mathcal{N}\left(\mu^A, \sigma^2 \left(\frac{t+1}{t}\right)\right), \tag{2}$$

where μ^A belongs to the confidence interval

$$\mathcal{P} \equiv [m - \kappa s, m + \kappa s], \tag{3}$$

[9] In Section 3 we generalize the model to a setting with unknown mean and variance.

with $\kappa > 0$ a preference parameter that captures the heterogeneity in attitude towards parameter uncertainty between the agents. When $\kappa = 0$, the set of priors $\mathcal{P}$ collapses to the singleton m, and A and S agents are identical. The parameter κ has also a classical statistical interpretation as a quantile of a distribution.[10]

Optimal Portfolios. Each agent $i = S, A$ is initially endowed with wealth W^i and chooses a portfolio of θ^i units of the risky assets. Therefore, the agents' terminal wealth is

$$\widetilde{W}^i = W^i (1 + r) + \theta^i \left(\tilde{d} - p(1 + r)\right), \quad i = S, A. \tag{4}$$

Agents S set θ^S to maximize their expected utility of terminal wealth, that is,

$$\max_{\theta^S} \mathbb{E}^S \left[-\frac{1}{\gamma} e^{-\gamma \widetilde{W}^S}\right], \tag{5}$$

subject to the budget constraint (4).

Agents A guard against parameter uncertainty by choosing portfolios that are robust to worst-case scenarios. This implies maximizing expected utility by using the "worst-prior" from the set $\mathcal{P}$ in equation (3). Formally, type-A agents buy portfolios θ^A that solve the following problem[11]

$$\max_{\theta^A} \min_{\mu^A \in \mathcal{P}} \mathbb{E}^A \left[-\frac{1}{\gamma_A} e^{-\gamma \widetilde{W}^A}\right], \tag{6}$$

subject to the budget constraint (4). The prior that minimizes A's expected utility in equation (6) is

$$\arg \min_{\mu^A \in \mathcal{P}} \mathbb{E}^A \left[u\left(W^A\right)\right] = \begin{cases} m - \kappa s, & \text{if } \theta^A > 0 \\ \mathcal{P} & \text{if } \theta^A = 0. \\ m + \kappa s, & \text{if } \theta^A < 0 \end{cases} \tag{7}$$

Therefore, the minimum expected utility for the ambiguity-averse agent A in equation (6) can be computed from the predictive distribution of $\tilde{d}$ in equation (2) where the belief μ^A

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

[11] 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 belief m, see, for example Anderson, Hansen, and Sargent (2000) and Hansen and Sargent (2008).

is replaced by either $m - \kappa s$, if $\theta^A > 0$ or $m_t + \kappa s$, if $\theta^A < 0$. This implies that agent A's problem in equation (6) is equivalent to the problem of agent S but with a "distorted" belief $\mu^A = m - \kappa s$ or $\mu^A = m + \kappa s$. The optimal portfolio of agent i is therefore

$$\theta^i = \frac{\mu^i - p(1+r)}{\gamma\sigma^2\left(\frac{t+1}{t}\right)}, \quad i = S, A. \tag{8}$$

Specifically, for agent A this implies

$$\theta^A = \begin{cases} \frac{m-\kappa s - p(1+r)}{\gamma\sigma^2\left(\frac{n+1}{n}\right)} > 0 & \text{if } p(1+r) < m - \kappa s, \\ 0 & \text{if } m - \kappa s < p(1+r) < m + \kappa s. \\ \frac{m+\kappa s - p(1+r)}{\gamma\sigma^2\left(\frac{n+1}{n}\right)} < 0 & \text{if } p(1+r) > m + \kappa s \end{cases} \tag{9}$$

Figure 1 shows the optimal demand θ^i as a function of the risky asset's price p. Agents A hold a more conservative portfolio than agents S, $|\theta^A| < |\theta^S|$. Moreover, as it is well-known, ambiguity aversion implies no-participation. That is, A agents hold a positive amount of the risky asset when p is sufficiently low, $p(1+r) < m - \kappa s$; short the risky asset when p is sufficiently high, $p(1+r) > m + \kappa s$; and do not participate otherwise.

Equilibrium. We determine the equilibrium price p by imposing market clearing, $\theta^A + \theta^S = 1$.

Proposition 1. *The equilibrium price p is given by*

$$p = \frac{1}{1+r} m - \lambda, \tag{10}$$

where λ is the subjective risk premium of the S agent,

$$\lambda = \begin{cases} \frac{\kappa}{2}\frac{\sigma}{\sqrt{t}} + \frac{\gamma}{2}\left(\frac{t+1}{t}\right)\sigma^2 & \text{if } \kappa \leq \kappa^*, \\ \gamma\left(\frac{t+1}{t}\right)\sigma^2 & \text{if } \kappa > \kappa^*. \end{cases} \quad \text{with} \quad \kappa^* \equiv \gamma\left(\frac{t+1}{\sqrt{t}}\right)\sigma. \tag{11}$$

The demand for the risky asset in equations (8) and (9) implies that either both agents hold long positions or only S agents participate. The expression for the equilibrium subjective risk premium in equation (11) shows that the participation of A-type agents is more likely when risk aversion, volatility, and the number of dividend observation is high. When both agents participate, i.e., $\kappa < \kappa^*$, the equilibrium risk premium is linear-quadratic in the

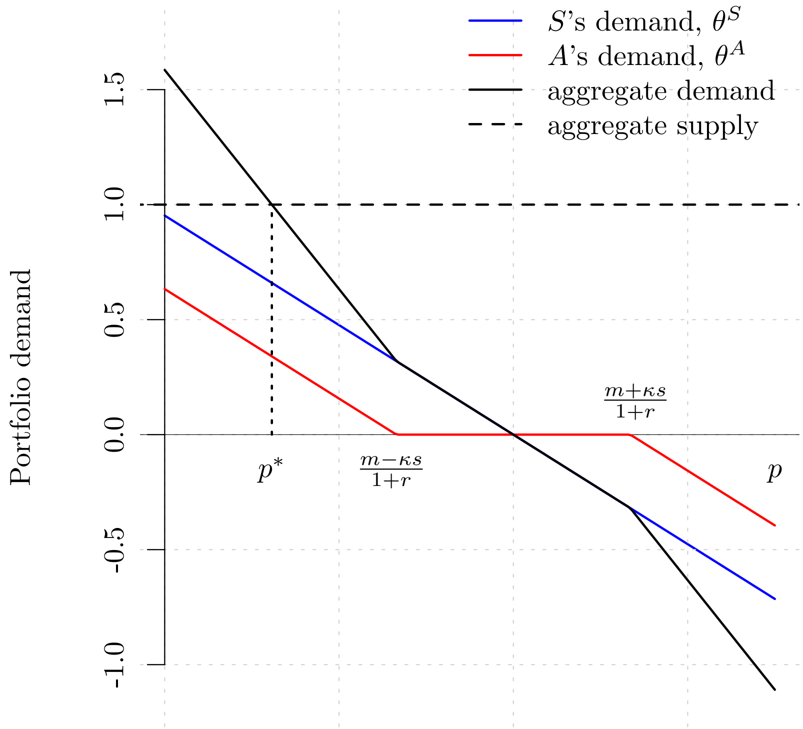

Figure 1: **Risky asset demand.** The figure shows the risky asset demand θ^S and θ^A from equations (8)–(9) as a function of the risky asset price p. The red line denotes type-A's demand; the blue line type-S's demand; the black line is aggregate demand; and the dashed line is the aggregate supply of the risky asset. Parameter values: $n = 20$, $\gamma = 1$, $\kappa = 0.15$, $\sigma = 0.1$.

dividend volatility σ. 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 S agents who are locally risk-neutral, A agents are locally risk-averse and demand a compensation for holding a vanishing amount of risk. Note that, from equation (11), the subjective risk premium λ depends only on the dividend volatility σ and the variance σ^2, and it will therefore be constant in an economy in which the dividend variance is a known constant.

Substituting the equilibrium price p from equation (10) in the agents' demand functions (8) and simplifying we obtain that, when both agents participate, the equilibrium weights are

$$\theta^A = \frac{1}{2} - \frac{\kappa}{2\gamma}\left(\frac{\sqrt{t}}{t+1}\right)\frac{1}{\sigma}, \tag{12}$$

$$\theta^S = \frac{1}{2} + \frac{\kappa}{2\gamma}\left(\frac{\sqrt{t}}{t+1}\right)\frac{1}{\sigma}. \tag{13}$$

Equations (12)–(13) show that the equilibrium portfolio holdings do not depend on the agents' beliefs about the mean but only on the dividend volatility, the risk and the ambiguity aversion parameters, and the number of observations. As the number of observations t increases, the portfolio weights of both agents converge to $1/2$, and the effect of ambiguity aversion, risk aversion, and dividend volatility vanishes.[12] This happens because the true value of the mean dividend is constant and will eventually be learned by both agents. In Section 3 we generalize the model to allow for perpetual learning.

Figure 2 illustrates the equilibrium portfolio weights θ^A and θ^S from equations (12)–(13) as a function of the dividend volatility σ. The figure shows that if $\kappa < \kappa^*$ or, equivalently, $\sigma > \sigma^* \equiv \frac{\sqrt{t}}{t+1}\frac{\kappa}{\gamma}$ (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 σ while S's demand is decreasing. As $\sigma \to \infty$, the portfolio holdings converge asymptotically to the constant weights $\theta^A = \theta^S = 1/2$.

Figure 3 provides an intuition for the structure of the equilibrium holdings in equations (12)–(13). The dotted curves in the figure represent "iso-portfolio" curves for both agents, that is, the combination of volatility σ and risk premium λ associated with the same risky asset demand from equations (8), red-dashed lines, and (9), blue-dashed lines. The solid black line traces the intersection of complementary iso-portfolio curves, that is the set of volatility and risk premia (σ, λ) for which the market clears, $\theta^S + \theta^A = 1$.

From equation (11) in Proposition 1, A agents participate in the market only if the risk premium exceeds the hurdle $\lambda^* = \frac{\kappa^2}{\gamma(t+1)} = \frac{\kappa}{\sqrt{t}}\sigma^*$. The $\theta^A = 0$ iso-portfolio line intersects the curve of the equilibrium risk premium line at (σ^*, λ^*). The red-shaded area indicates (σ, λ) combinations for which A agents do not participate. For values of $\sigma < \sigma^*$, only S agents hold the risky asset in equilibrium, and the equilibrium risk premium coincides with the $\theta^S = 100\%$ iso-curve, i.e., the highest blue-dashed line. For values of $\sigma > \sigma^*$, both agents

[12] The value $1/2$ corresponds to the perfect risk-sharing portfolio when agents have the same value of risk aversion γ and no aversion to ambiguity. Our analysis can easily be extended to the case of heterogeneous risk aversion. We refrain from it to highlight the role of heterogeneity in agents' attitude towards ambiguity.

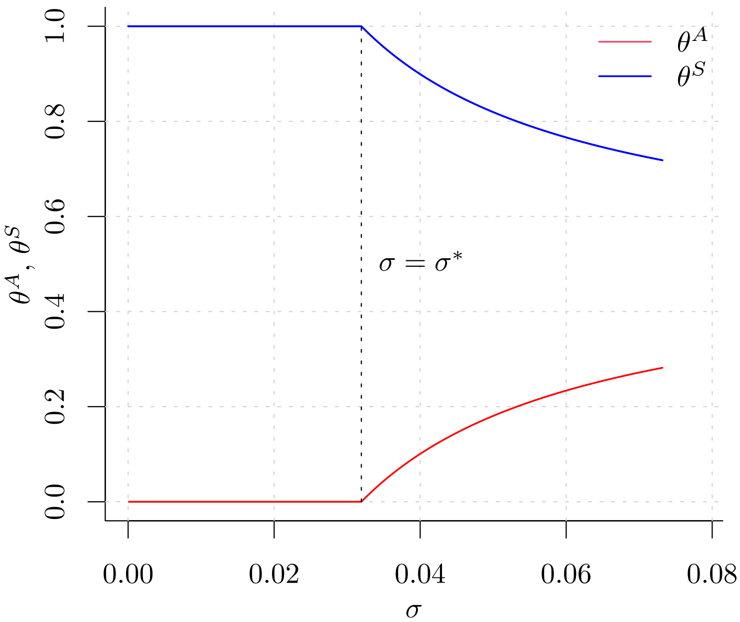

Figure 2: **Equilibrium portfolios.** The figure shows the equilibrium portfolios θ^A and θ^S from equations (12)–(13) as a function of the dividend volatility σ. The vertical dashed line indicates the participation threshold $\sigma^* \equiv \frac{\sqrt{t}}{t+1}\frac{\kappa}{\gamma}$. For values of $\sigma < \sigma^*$, ambiguity-averse A agents do not hold the risky asset, i.e., $\theta^A = 0$ and $\theta^S = 1$. Parameter values: $t = 20$, $\gamma = 1$, $\kappa = 0.15$.

participate in equilibrium. Lemma A.1 in Appendix A shows that in any equilibrium in which A agents participate, their iso-portfolio lines in Figure 3 are always flatter than those of S agents. Intuitively, because A agents hold fewer units of the risky asset than S agents, starting from an equilibrium in which both A and S participate, A agents require relatively less compensation than S agents to bear an additional unit of volatility while keeping the portfolio unchanged. This implies that the marginal rate of substitution between required risk premium and dividend volatility is strictly higher for S than for A. Because the risk premium in an equilibrium where both agents participate lies in between these two extremes, it will be perceived as "too high" (price too low) by A and "too low" (price too high) by S, and a gain from trade emerges in which S is willing to sell and A is willing to buy. To

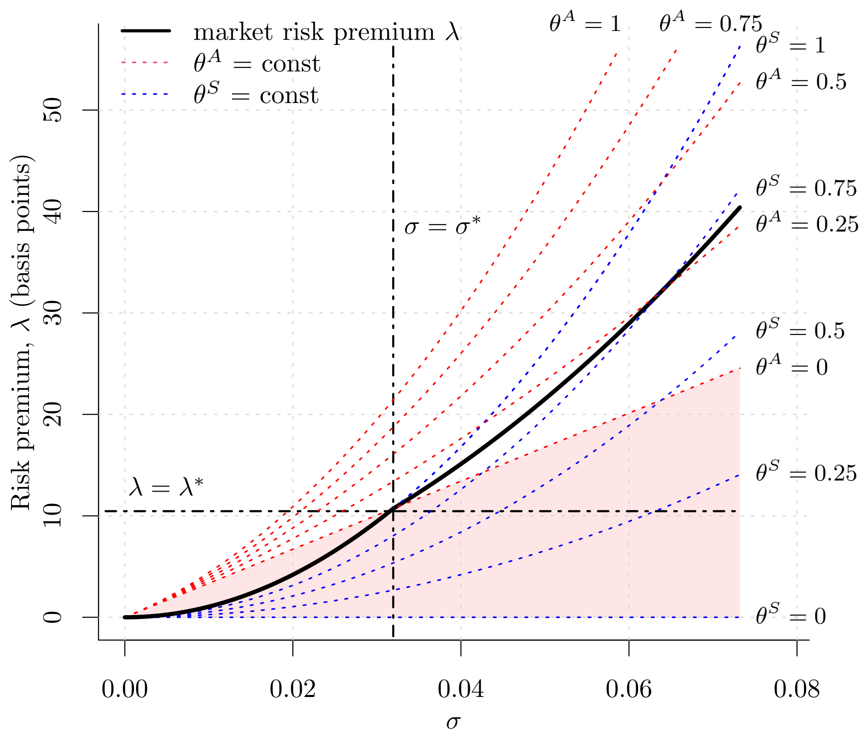

Figure 3: **Equilibrium portfolios and risk premia.** The figure shows iso-portfolios lines of type-A (red-dashed) and type S (blue-dashed) agents. These lines represent the set of volatility and risk premium values (σ, λ) that correspond to a given constant portfolio weight in equations (8) and (9). The solid black line represents the set of points (σ, λ) at which the market clears, $\theta^A + \theta^S = 1$, and λ is the equilibrium risk premium λ. The vertical dashed-dotted line indicates the participation threshold $\sigma^* \equiv \frac{\sqrt{t}}{t+1}\frac{\kappa}{\gamma}$, and the horizontal dashed-dotted line indicates the hurdle risk premium $\lambda^* = \frac{\kappa}{\sqrt{t}}\sigma^*$. Parameter values: $t = 20$, $\gamma = 1$, $\kappa = 0.15$.

summarize: In equilibrium, A agents hold a "conservative" portfolio but trade "aggressively" by increasing the holdings of the risky asset following an increase in volatility.

While the above analysis suggests a role for volatility in driving equilibrium flows between agents that differ in their ambiguity aversion, the stylized model of this section is not equipped to study how portfolio flows can emerge in equilibrium as a result of learning over time. In the next section, we develop a model to address this issue.

3 An overlapping-generation model

In this section, we extend the two-period model of the previous section to an infinite-horizon, overlapping-generation (OLG) setting. The main goal of this analysis is twofold. First, we confirm the main economic mechanism of the simple model of Section 2 in a dynamic context. Second, we provide a formal analysis of how portfolio flows can emerge as an equilibrium outcome when agents who differ in their ambiguity attitude learn about economic fundamentals.

3.1 Setup

We consider an infinite-horizon OLG model in which each generation consists of type-S and type-A agents in equal mass, as in Section 2, who live for two periods. The setup we consider is similar to De Long, Shleifer, Summers, and Waldmann (1990) and Lewellen and Shanken (2002), however, unlike De Long, Shleifer, Summers, and Waldmann (1990) there are no noise traders in our model, but agents differ in their attitude towards ambiguity. Unlike Lewellen and Shanken (2002), each generation consists of heterogeneous agents instead of a representative agent.

Assets. There is a riskless asset in perfectly elastic supply that pays the interest rate r in every period $t = 1, \ldots, \infty$ and a risky security in unit supply that pays the dividend d_t in each period t. Dividends are iid and are normally distributed,

$$d_t \sim \mathcal{N}(\mu, \sigma^2), \tag{14}$$

with constant mean μ and variance σ^2. Agents know that dividends are normally distributed, but they do not know the moments of the distribution.

Investors. Agents live for two periods with overlapping generations. There is no first-period consumption or labor supply. In the first period, agents only decide how to allocate their exogenous wealth between the risky and risk-free asset. In the second period, agents collect the dividend, liquidate their risky portfolio by selling it to the new incoming generation, and consume the proceeds. There is no bequest. As in Section 2, we assume that both agents have CARA preferences but differ in the way they determine expected end-of-period wealth from historical dividend data: S-agents are subjective expected utility investors with a *unique* prior over unknown parameters. A-agents are averse to ambiguity and entertain *multiple* priors over the unknown parameters.

Because investors are short-lived, their portfolio decisions do not contain an intertemporal hedging component. However, in equilibrium, to construct their portfolio, generation-t investors need to form beliefs about both future dividends d_{t+1} and asset prices p_{t+1}. To do so, they would need to know how generation-$(t+1)$ forms beliefs and so on, ad infinitum. As in the case of the two-period model of Section 2, we assume that each generation observes the past history of realized dividends and performs the same statistical analysis to estimate the dividend moments and their confidence intervals on which agents' beliefs are based. Because of this belief formation process, the set of priors of the ambiguity-averse agent A in each generation t can be characterized as confidence intervals.

Portfolio flows. The OLG structure of this economy allows us to formally define portfolio flows. Denoting by θ_t^i the t-generation risky asset demand of type-i agents, we define portfolio flows at time t as

$$\Delta\theta_t^i = \theta_t^i - \theta_{t-1}^i. \tag{15}$$

A positive flow $\Delta\theta_t^i > 0$ implies that the t-generation of type-i agents increases risky asset holdings relative to the $(t-1)$-generation. Such a positive flow represents an intra-generational trade in which type-i agents buy the risky asset from non-type-i agents. In Section 3.3, we show how such flows could result from agents' learning about the dividend variance. In preparation for this result, in Section 3.2 we derive the equilibrium for an OLG economy in which the dividend variance σ^2 is known. This case allows for a closed-form solution and generalizes the analysis of Section 2.

3.2 Equilibrium with known variance

Figure 4 illustrates the learning process underlying the construction of the equilibrium. The top part represents the "information processing" step. Both agents agree on the estimate of the unknown dividend mean and its standard error, that is, the sample mean serves as a state variable. The bottom part of the figure illustrates the "belief formation" step that affects agents' portfolio choice problem.

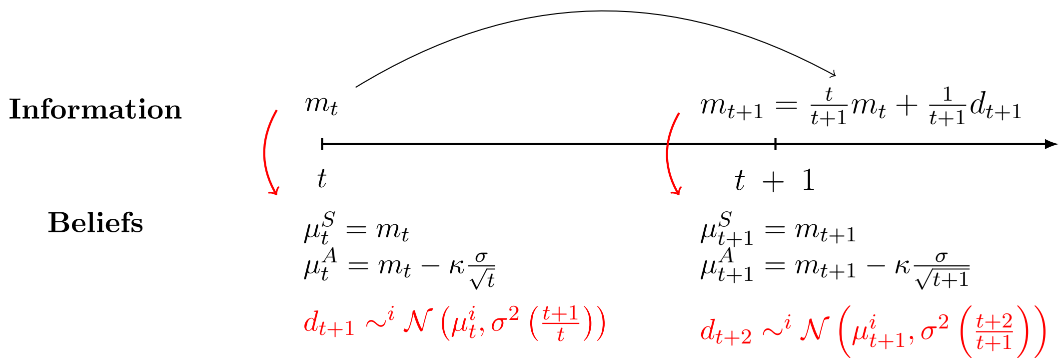

Figure 4: **Updating and belief formation: known variance.** The top part of the figure illustrates the "information processing" step through which the sample mean m_t is updated upon observation of a new dividend realization d_{t+1}. The bottom part of the figure illustrates the "belief formation" step, and shows how agents at time t use the information m_t from the information processing step to form predictive distributions about the future dividend d_{t+1}.

Information processing. At each time t, both types of agents observe a history of t realized dividends and agree on its sample mean m_t and the corresponding standard error s_t

$$m_t \equiv \frac{1}{t} \sum_{k=1}^{t} d_k, \quad s_t \equiv \frac{\sigma}{\sqrt{t}}. \tag{16}$$

The sample mean m_t is a state variable that is observable by all agents. The standard error s_t of m_t is a deterministic fraction of the dividend volatility. At each time t, agents understand that the next generation will compute the time $t+1$ sample mean m_{t+1} after observing the dividend d_{t+1} as

$$m_{t+1} = \frac{t}{t+1} m_t + \frac{1}{t+1} d_{t+1}. \tag{17}$$

Equation (17) represents the updating rule of the state variable m_t. The equation highlights that the state variable m_t is "handed" over by generation t to generation $t+1$, who updates following the observation of the d_{t+1} dividend. This updating rule is agreed upon by all agents. The state variable m_t and its law of motion (17) are the input for agents' belief formation process needed for the construction of an equilibrium.

Belief formation and portfolio choice. Agents in the economy base their beliefs about the unknown dividend mean μ on the observed state variable m_t. Type-S agents' prior about

μ is given by[13]

$$\mu \sim^S \mathcal{N}\left(\mu_t^S, \frac{\sigma^2}{t}\right), \quad \mu_t^S = m_t. \tag{18}$$

In contrast, ambiguity-averse agents entertain a *set of priors* $\mathcal{P}_t^\mu$ representing a confidence interval around the sample mean,

$$\mathcal{P}_t^\mu \equiv \left[m_t - \kappa\frac{\sigma}{\sqrt{t}},\ m_t + \kappa\frac{\sigma}{\sqrt{t}}\right]. \tag{19}$$

Specifically, for A agents the set of priors for the mean μ is

$$\mu \sim^A \mathcal{N}\left(\mu_t^A, \frac{\sigma^2}{t}\right), \quad \mu_t^A \in \mathcal{P}_t^\mu. \tag{20}$$

At each time t, agents are initially endowed with wealth W_t^i and choose a portfolio of θ_t^i units of the risky assets. Their wealth at time $t+1$ is

$$W_{t+1}^i = W_t^i(1+r) + \theta_t^i\left(p_{t+1} + d_{t+1} - p_t(1+r)\right), \quad i = S, A. \tag{21}$$

Hence, to determine their portfolios at time t, agents have to form expectations about future dividends d_{t+1} and prices p_{t+1}. From the beliefs (18) and (20), the subjective distribution of future dividends for i agents is[14]

$$d_{t+1} \sim^i \mathcal{N}\left(\mu_t^i, \sigma^2\left(\frac{t+1}{t}\right)\right), \quad \mu_t^S = m_t,\ \mu_t^A \in \mathcal{P}_t^\mu. \tag{22}$$

To determine their optimal portfolio, in addition to the future dividend, agents have to form expectations about the future price p_{t+1}. We conjecture, and later verify, that the equilibrium price p_t is affine in the state variable m_t. Therefore, t-generation agents need to form expectations about the future state variable m_{t+1} in order to determine their risky asset demand. Therefore, using the predictive distributions of the future dividend in equation (22), the agents' subjective distribution of the state variable m_{t+1} is

$$m_{t+1} \sim^i \mathcal{N}\left(m_t\left(\frac{t}{t+1}\right) + \frac{1}{t+1}\mu_t^i, \sigma^2\left(\frac{1}{t(t+1)}\right)\right), \quad i = S, A, \tag{23}$$

[13]This distribution can be thought of as the Bayesian update of a diffuse prior over μ. See, e.g., Section 2.5 of Gelman, Carlin, Stern, Dunson, Vehtari, and Rubin (2020) for a proof of this result.

[14]See, e.g., Section 2.5 of Gelman, Carlin, Stern, Dunson, Vehtari, and Rubin (2020) for a proof of this result.

with $\mu_t^S = m_t$ and $\mu_t^A \in \mathcal{P}_t^\mu$. Equation (23) shows that when forming beliefs about the future sample mean m_{t+1} agents' subjective distribution differs because of their different predictive distributions of the future dividend realization d_{t+1}. From equation (23) and the conjecture that the price p_{t+1} is affine in m_{t+1}, it follows that for both types of agents the predictive distribution of the price p_{t+1} is normal. Both types of t- and $t+1$-generation agents understand and agree on the value of the state variable m_t, which can be interpreted as the historical "knowledge" that is passed on from generation t to generation $t+1$. In Section 3.3, in order to model perpetual learning, we relax the assumption that the state variable is perfectly communicated across generations and allow for "information leakage".

Equilibrium. Given the assumption of CARA preferences for both agents, the optimal portfolios are

$$\theta_t^i = \frac{\mathbb{E}_t^i\left[p_{t+1} + d_{t+1}\right] - (1+r)p_t}{\gamma \mathrm{Var}_t\left[p_{t+1} + d_{t+1}\right]}, \quad i = S, A. \tag{24}$$

In equilibrium the price p_t is such that asset markets clear, that is $\theta_t^A + \theta_t^S = 1$. The following proposition characterizes the equilibrium price and the portfolio weights in an infinite-horizon overlapping-generation economy when variance is known.

Proposition 2. *The equilibrium price of the risky asset when both agents participate is*

$$p_t = \frac{1}{r} m_t - \Lambda_t, \tag{25}$$

where the risk premium Λ_t is given by

$$\Lambda_t = g_t \frac{\kappa}{2}\sigma + f_t \frac{\gamma}{2}\sigma^2, \tag{26}$$

with g_t and f_t deterministic functions of time defined in equations (A.23) and (A.24) of Appendix A. The equilibrium portfolio weights are

$$\theta_t^A = \frac{1}{2} - \frac{\kappa}{2\gamma}\left(\frac{r\sqrt{t}}{1+r(t+1)}\right)\frac{1}{\sigma}, \tag{27}$$

$$\theta_t^S = \frac{1}{2} + \frac{\kappa}{2\gamma}\left(\frac{r\sqrt{t}}{1+r(t+1)}\right)\frac{1}{\sigma}. \tag{28}$$

The equilibrium weights (27)–(28) are the infinite-horizon OLG equivalent of the equilibrium weights (12)–(13) in the two-period model of Section 2. The portfolio holdings highlight that only unexpected changes in volatility can generate trade among agents in equilibrium. As in the simple model of Section 2, A-agents hold conservative portfolios, $\theta_t^A < \theta_t^S$, but increase risk when σ increases, i.e., $\partial\theta_t^A/\partial\sigma > 0$. Unfortunately, a model with constant and known dividend variance is not rich enough to explain portfolio flows. In the next section, we extend the analysis to the case in which volatility is unknown and overlapping generations learn about both the mean and the variance of the dividend process. Such a model delivers a realistic description of flows in equilibrium and forms the foundation for our empirical analysis in Section 4.

3.3 Equilibrium with unknown variance

In this section, we develop an overlapping-generations model that addresses the two shortcomings of the analysis so far: First, the dividend variance is perfectly observable, resulting in the absence of flows between ambiguity-averse and ambiguity-neutral agents. Second, as the number of dividend observations increases, the effect of learning on asset prices vanishes as agents learn the true parameters in the limit. To address the first shortcoming, we assume that variance is not known and that agents learn about it. This implies that the subjective variance is time-varying and its dynamics gives rise to flows between agents of different type in equilibrium. An alternative way to introduce time-variation in volatility would be to just assume stochastic and observable volatility. In Section 3.6, we show that in a model with stochastic volatility ambiguity declines monotonically over time as any new dividend observation reduces the size of the mean confidence interval. To address the second shortcoming, we generalize the model so that uncertainty does not disappear after agents observe a long history of data. We achieve this goal by assuming that some of the information from past observations is gradually lost as generations overlap. In a representative agent economy, this setting coincides with the idea of "fading memory" as in, e.g., Nagel and Xu (2021), or "age-related experiential learning", as in Malmendier and Nagel (2016) and Collin-Dufresne, Johannes, and Lochstoer (2016).

Unknown variance. When estimating unknown mean and variance from normally distributed dividends, the predictive dividend distribution is Student-t, see Greene (2020). As a consequence of the heavy tails of this distribution, the expected CARA utility is not well-defined (see, e.g. Geweke, 2001; Weitzman, 2007). We overcome this difficulty by relying on the approach proposed by Bakshi and Skoulakis (2010) who provide a learning framework in which the dividend variance is constrained to a pre-specified, arbitrary finite interval. They replace the well-known normal-inverse Gamma framework for updating mean and variances with a normal-inverse *truncated* Gamma setup that preserves conjugacy.[15] The effect of restricting the variance is to dampen the fat tails of the Student-t distribution, which allows us to properly define the agents' portfolio choice problem.

Information leakage and perpetual learning Agents of each generation have a common knowledge about the dividend mean μ and the precision $\phi \equiv 1/\sigma^2$. This knowledge is summarized by prior distributions available to both types of agents at birth. When generation t "hands over" information to generation $(t+1)$ some information leakage occurs. As a result, the next generation perceives the inherited knowledge less precise when updating it with the new dividend observation, d_{t+1}, and hence, puts a weight on the new dividend observation that is higher than the one the previous generation would use. Formally, we model this information loss as a shock that affects the priors about μ and ϕ. When forming beliefs about the resale price p_{t+1} of the risky asset, agents of generation t, however, fully anticipate this information leakage and the information processing of the next generation. How agents S and A differ in their belief formation is explained below.

Figure 5 illustrates the learning process underlying the construction of the equilibrium for the case of unknown variance. The top part represents the "information processing" step, referring to the evolution of the observable state variables. The bottom part of the figure illustrates the "belief formation" step that affects agents portfolio choice problem.

Now we describe the "data processing" in a normal-inverse truncated Gamma setup in detail. The prior for the precision ϕ is a truncated Gamma with ν_t degrees of freedom and "shape" parameter b_t:

$$\phi \sim \mathrm{TG}\left[\frac{\nu_t}{2}, \frac{b_t}{2}; \underline{\phi}, \overline{\phi}\right], \quad 0 < \underline{\phi} < \overline{\phi} < \infty. \tag{29}$$

The shape parameter b_t is essentially the sum of historical squared errors. The parameters $\underline{\phi}$ and $\overline{\phi}$ are arbitrary truncation constants and all agents agree on these bounds for the precision. Truncating the precision is key to guarantee existence of agents' expected utility.

[15]See, e.g., Gelman, Carlin, Stern, Dunson, Vehtari, and Rubin (2020) for an introduction to Bayesian data analysis, including Bayesian updating of prior distributions to newly observed information and the concept of conjugacy.

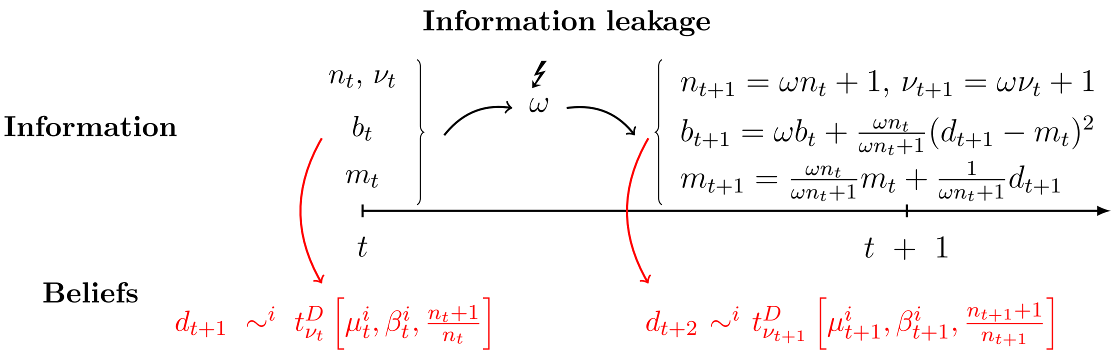

Figure 5: **Updating and belief formation: unknown variance.** The top part of the figure illustrates the "information processing" step with information leakage ($\omega < 1$) in which the state variables n_t, b_t and m_t are updated upon observation of a new dividend realization d_{t+1}. The bottom part of the figure illustrates the "belief formation" step, and shows how agents at time t use the information from the state variables in the information processing step to form predictive distributions about the future dividend.

Conditional on ϕ the prior of the mean μ is normally distributed:

$$\mu|\phi \sim \mathcal{N}\left(m_t, \frac{1}{n_t \phi}\right), \tag{30}$$

with n_t describing the precision of the prior about μ relative to the dividend variance.

The generation-t information set consists of the state variables m_t, b_t, n_t, and ν_t. Without information leakage, n_t corresponds to the number of observations used for the estimate of μ and $\nu_t = n_t - 1$. As we discuss below, information leakage leads to a down-weighting of historical data relative to more recent observations. With a new dividend observation, information is gained, but at the same time part of the historical information gets lost. So n_t and ν_t do not grow linearly with the number of observations, but their dynamics properly accounts for the net-informativeness of the priors of each generation. We therefore refer to n_t as "effective number of observations" in the sense it is used by Weitzman (2007), Bakshi and Skoulakis (2010), and Nagel and Xu (2021).

The parameter $\omega \in [0, 1]$ controls the amount of data handed over to the next generation, with $(1 - \omega)$ the extent of information loss. Appendix B.2 provides details about how we explicitly model the shocks to the priors about μ and ϕ and how the state variables are updated in a Bayesian way with new dividend information. Information gain and loss is described by the recurrence

$$n_{t+1} = \omega n_t + 1, \quad \nu_{t+1} = \omega \nu_t + 1. \tag{31}$$

For small t, $\nu_t \approx n_t - 1$. However, as t grows and information leakage becomes relevant, both n_t and ν_t approach a common asymptotic value $\overline{n} = \frac{1}{1-\omega}$. This implies that, for large t, only m_t and b_t are state variables of the economy. When this "steady state" is reached, the gain in estimation precision from observing new dividend information and the information leakage over time are exactly balanced.

The updated posteriors, which serve as the priors for generation $t+1$, are

$$\begin{aligned} \phi &\sim \mathrm{TG}\left[\frac{\nu_{t+1}}{2}, \frac{b_{t+1}}{2}; \underline{\phi}, \overline{\phi}\right], \\ \mu|\phi &\sim \mathcal{N}\left(m_{t+1}, \frac{1}{n_{t+1}\phi}\right), \end{aligned}$$

with the new estimate of the dividend mean m_{t+1} given by

$$m_{t+1} = \frac{\omega n_t}{\omega n_t + 1} m_t + \frac{1}{\omega n_t + 1} d_{t+1}, \tag{32}$$

and the updated sum of squared errors b_{t+1} given by

$$b_{t+1} = \omega b_t + \frac{\omega n_t}{\omega n_t + 1}(d_{t+1} - m_t)^2. \tag{33}$$

The parameter ω "blurs" information as it gets passed across generations. Intuitively, when $\omega < 1$, the $(t+1)$-generation partially forgets the information contained in m_t and "overweights" the more recent observation d_{t+1}. When $\omega = 1$ the $(t+1)$-generation does not forget the past and equation (32) corresponds to the standard updating of the sample mean used in equation (17). Likewise, equation (33) shows that the information b_t about historical variance is down-weighted to ωb_t before being updated with the newly observed squared error $(d_{t+1} - m_t)^2$.

Belief formation and portfolio choice. Although both types of agents observe the same state variables and process information in the same way, they differ in how they use such information to form beliefs about the distribution of their future wealth. Specifically, S agents time-t prior about the precision ϕ is truncated Gamma distributed with shape parameter $\beta_t^S = b_t$,

$$\phi \sim^S \text{TG}\left[\frac{\nu_t}{2}, \frac{\beta_t^S}{2}; \underline{\phi}, \overline{\phi}\right], \quad \beta_t^S = b_t, \quad 0 < \underline{\phi} < \overline{\phi} < \infty, \tag{34}$$

and, conditional on ϕ, their prior about the mean μ is normal with mean $\mu_t^S = m_t$,

$$\mu|\phi \sim^S \mathcal{N}\left(\mu_t^S, \frac{1}{n_t\phi}\right), \quad \mu_t^S = m_t. \tag{35}$$

In contrast, the ambiguity-averse A agents entertain the following set of truncated Gamma-normal priors

$$\begin{aligned} \phi &\sim^A \text{TG}\left[\frac{\nu_t}{2}, \frac{\beta_t^A}{2}; \underline{\phi}, \overline{\phi}\right], &(36)\\ \mu|\phi &\sim^A \mathcal{N}\left(\mu_t^A, \frac{1}{n_t\phi}\right), &(37)\end{aligned}$$

where (μ_t^A, β_t^A) belong to the set

$$\mathcal{P}_t^{(\mu,\sigma^2)} = \left\{(\mu_t^A, \beta_t^A) : \mu_t^A \in \mathcal{P}_t^m, \beta_t^A \in \mathcal{P}_t^b\right\}, \tag{38}$$

with

$$\begin{aligned} \mathcal{P}_t^m &= \left[m_t - \frac{\kappa}{\sqrt{n_t \mathbb{E}_t^A[\phi|\beta_t^A]}}, m_t + \frac{\kappa}{\sqrt{n_t \mathbb{E}_t^A[\phi|\beta_t^A]}}\right], &(39)\\ \mathcal{P}_t^b &= \left[\underline{\ell} b_t, \bar{\ell} b_t\right], \quad 0 < \underline{\ell} < 1 < \bar{\ell}. &(40)\end{aligned}$$

At each time t, agents are initially endowed with wealth W_t^i and choose a portfolio of θ_t^i units of the risky assets. Their wealth at time $t+1$ is

$$W_{t+1}^i = W_t^i(1+r) + \theta_t^i\left(p_{t+1} + d_{t+1} - p_t(1+r)\right), \quad i = S, A. \tag{41}$$

Hence, to determine their portfolios at time t, agents have to form expectations about future dividends d_{t+1} and prices p_{t+1}. Note that agents of generation t do not suffer from memory loss over the span of their life, thus their belief about the next period's dividend d_{t+1} is not subject to information leakage. However, they rationally anticipate that the $t+1$ generation uses a different information set to determine p_{t+1} from d_{t+1}. Lemma B.1

in Appendix B.1 shows that agents i's predictive distribution of future dividends is a non-standardized "dampened" Student-t[16]

$$d_{t+1} \sim^i t^D_{n_t}\left[\mu^i_t, \beta^i_t, \frac{n_t+1}{n_t}\right], \tag{42}$$

where $(\mu^S_t, \beta^S_t) = (m_t, b_t)$ and $(\mu^A_t, \beta^A_t) \in \mathcal{P}^{(\mu,\sigma^2)}_t$. The dampened Student-t distribution has thinner tails than the Student-t. This guarantees the existence of a moment generating function, which allows us to formally define the agent's expected utility and solve for optimal portfolios.

The following proposition characterizes the time-t generation expected utility for the case in which both the mean and variance are unobservable.

Proposition 3. *Suppose $d_t \sim \mathcal{N}(\mu, \sigma^2)$, with μ and σ unobservable and that agents $i = S, A$ form beliefs μ^i_t and β^i_t about μ and σ^2 as described in equations (35)–(34) and (37)–(36). Then, given a portfolio θ^i_t and state variables m_t, b_t, n_t, and ν_t, the time-t expected utility of future wealth W^i_{t+1} from equation (21) is*

$$\mathbb{E}^i_t[u(W^i_{t+1})] = \frac{1}{C(\beta^i_t, \nu_t; \underline{\phi}, \overline{\phi})}\int_{\underline{\phi}}^{\overline{\phi}} \mathbb{E}^i_t[u(W^i_{t+1})|\phi] f^i(\phi) d\phi, \tag{43}$$

with $f^i(\phi) = \phi^{\frac{\nu_t}{2}-1} e^{-\phi\frac{\beta^i_t}{2}}$ the density of the truncated Gamma distribution, $C(\beta^i_t, \nu_t; \underline{\phi}, \overline{\phi})$ an integration constant defined in equation (B.1) of Appendix B.2 and

$$\mathbb{E}^i_t[u(W^i_{t+1})|\phi] = -\frac{1}{\gamma} e^{-\gamma(1+r)(W_t - \theta p_t)} \sqrt{\frac{\phi}{2\left(\frac{n_t+1}{n_t}\right)\pi}} \int_{-\infty}^{\infty} e^{-\gamma\theta^i_t\left(\mu^i_t + e^i_{t+1} + p_{t+1}\right) - \frac{1}{2}\frac{\phi}{\left(\frac{n_t+1}{n_t}\right)}(e^i_{t+1})^2} de^i_{t+1}. \tag{44}$$

In equation (44), $e^i_{i+1} = d_{t+1} - \mu^i_t$ denotes agents i's dividend surprise, $p_t \equiv p_t(m_t, b_t, n_t, \nu_t)$, for all t, and the $t+1$ state variables are

$$n_{t+1} = \omega n_t + 1, \quad \nu_{t+1} = \omega \nu_t + 1, \quad \omega \in [0, 1], \tag{45}$$

$$m_{t+1} = \left(1 - \frac{1}{n_{t+1}}\right) m_t + \frac{1}{n_{t+1}} d_{t+1}, \tag{46}$$

$$b_{t+1} = \omega b_t + \omega \frac{n_t}{n_{t+1}} (d_{t+1} - m_t)^2. \tag{47}$$

[16] See Definitions B.1 and B.2 in Appendix B.1 for a formal definition of the density of a dampened t-distribution.

The above proposition characterizes the expected utility of future wealth for a given portfolio θ_t^i. The existence of the expected utility in equation (43) follows from the fact that the dividend variance is truncated over the bounded support $[1/\overline{\phi}, 1/\underline{\phi}]$, which results in dampened t-distributed dividends, see Bakshi and Skoulakis (2010). The boundedness of this variance implies boundedness of the risk premium in equilibrium. This guarantees that the equilibrium price p_{t+1} is finite, and hence the integral in equation (44) is well defined.

Optimal portfolios. In each generation, agents face the budget constraint (21) and choose their optimal portfolio by solving, respectively, the following maximization problems,

$$\max_{\theta_t^S} \mathbb{E}_t^S[u(W_{t+1}^S)], \tag{48}$$

and

$$\max_{\theta_t^A} \min_{(\mu_t^A, \beta_t^A) \in \mathcal{P}_t^{(\mu_t, \sigma^2)}} \mathbb{E}_t^A[u(W_{t+1}^A)], \tag{49}$$

where the expected utility $\mathbb{E}_t^i\left[u(W_{t+1}^i)\right]$ is derived in Proposition 3 and the set of prior $\mathcal{P}^{(\mu,\sigma^2)}$ is defined in equation (38). While the above setup allows to consider ambiguity aversion about both mean and variance, in the max-min setting of Gilboa and Schmeidler (1989) the ambiguity-averse agent always elects the highest possible return variance in the set $\mathcal{P}^{(\mu,\sigma^2)}$ when constructing optimal portfolios, see, e.g., Easley and O'Hara (2009). Therefore, the portfolio choice problem with ambiguity about both μ and σ^2 reduces to a problem with ambiguity only about μ, where β_t^A is fixed at the maximum over the set of its possible values, that is, $\overline{\ell} b_t$.

Equilibrium. The optimal demand of the risky asset is therefore determined by the first-order conditions

$$\frac{d\mathbb{E}_t^i[u(W_{t+1}^i)]}{d\theta_t^i} = 0, \quad i = S, A. \tag{50}$$

By Proposition 3, the optimal demand θ_t^i and hence the equilibrium price p_t depend on the subjective distribution of the future dividend d_{t+1} and price p_{t+1}.

To construct an equilibrium, we focus on large values of t, for which $n_t = \nu_t = \overline{n}$. This reduces the state variables to two only: m_t and b_t. Under this assumption, we first solve for the price p_t in a fictitious economy with τ periods, where we can use the terminal condition $p_{t+\tau} = 0$. We conjecture, and verify, that the equilibrium price in this fictitious economy takes the form

$$p(m_t, b_t, \tau) = h(t, \tau) m_t - \Lambda(b_t, \tau). \tag{51}$$

The price in equation (51) is the discounted expected future dividend income minus a risk premium $\Lambda(b_t, \tau)$ that depends on the current belief about the dividend volatility, characterized by b_t, and the time horizon τ. The time t equilibrium price in the overlapping-generation economy is given by the limit as $\tau \to \infty$ of $p(m_t, b_t, \tau)$

$$p(m_t, b_t) = \frac{1}{r} m_t - \Lambda(b_t). \tag{52}$$

Hence, to construct an equilibrium, we must determine the risk premium Λ as a univariate function of b_t. Unfortunately, a closed-form expression is not available for the case with learning about variance. In Appendix C we provide details of the numerical procedure we use to construct the equilibrium.

3.4 Equilibrium flows

Proposition 2 in Section 3.2 shows that, in an equilibrium when both types of agents participate, portfolio holdings depend on the dividend volatility, σ. This result suggests that, in an economy with both ambiguity-averse and ambiguity-neutral agents, changes in volatility play a key role in determining portfolio flows. When the volatility parameter is not known, as in Section 3.3, agents' learning about volatility naturally generates time-variation in the subjective beliefs about volatility, which in turn gives rise to inter-generational portfolio flows in equilibrium. In this section, we illustrate the implications of the equilibrium model of Section 3.3 for the dynamic of portfolio flows and risk premia.

Figure 6 shows that A-agents' (S-agents') equilibrium portfolio from the overlapping-generation economy of Section 3.3 is an *increasing* (*decreasing*) function of the standard deviation estimate $\hat{\sigma}_t = \sqrt{b_t/\nu_t}$. This pattern is consistent with the intuition we developed in the two-period model of Section 2 and with the OLG model with known variance in Section 3.2. In both cases, in fact, the portfolio weights of A-agents (S-agents) are increasing (decreasing) in the volatility σ. Figure 6 extends this intuition to a realistic model in which variance is unknown. In the figure we set an a-priori restriction of $\sigma \in [0.2, 0.6]$. When the estimated $\hat{\sigma}_t$ takes values outside this range, agents become very confident that the true σ is either at the upper or at the lower bound of the a-priori interval. Therefore, as the figure

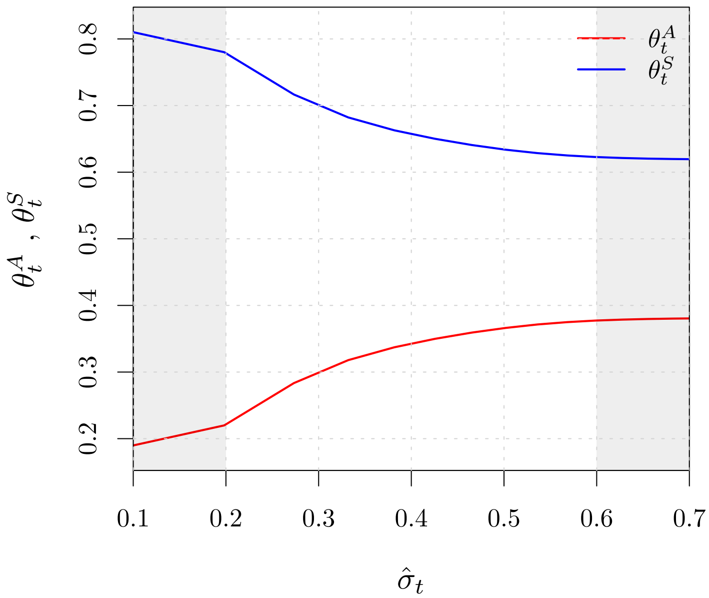

Figure 6: **Equilibrium portfolios and volatility.** The figure shows A and S agents' equilibrium portfolios as a function of the standard deviation estimate $\hat{\sigma}_t \equiv \sqrt{b_t/\nu_t}$, where b_t is the sum of squared error and ν_t the degrees of freedom. The dividend mean and variance are unknown to agents, as in Section 3.3. Parameters values: $\kappa = 1$, $\underline{\phi}$ and $\overline{\phi}$ such that the true volatility $\sigma \in [0.2, 0.6]$, $\overline{\ell} = 1.05$, $r = 0.1$, and $\nu_t = \overline{n} = 20$.

shows, the portfolio weights are less sensitive to changes in $\hat{\sigma}_t$ for values outside the range of $[0.2, 0.6]$.

The dependence of the portfolio weights θ_t^A and θ_t^S on $\hat{\sigma}_t$ illustrated in Figure 6 has a direct counterpart in terms of dividend "surprises", i.e., deviations of the realized dividend from the historical mean m_t. Figure 7 shows the equilibrium risky asset holdings of A-agents (left panel) and S-agents (right panel) as a function of this surprise. Consistent with the structure of the portfolios described in Figure 6, Figure 7 shows that large dividend surprises are associated with large subjective values of volatility. Different lines corresponds to different values of the ambiguity aversion parameter κ. Larger values of κ imply stronger ambiguity aversion and more conservative (aggressive) portfolios for A-agents (S-agents). The U-shape nature of the equilibrium portfolios in the left panel of Figure 7 implies that A-agents are more aggressive than S-agents in their trades. After large positive and negative surprises A-agents *increase* their risky asset holdings. Figure 6 confirms this finding by

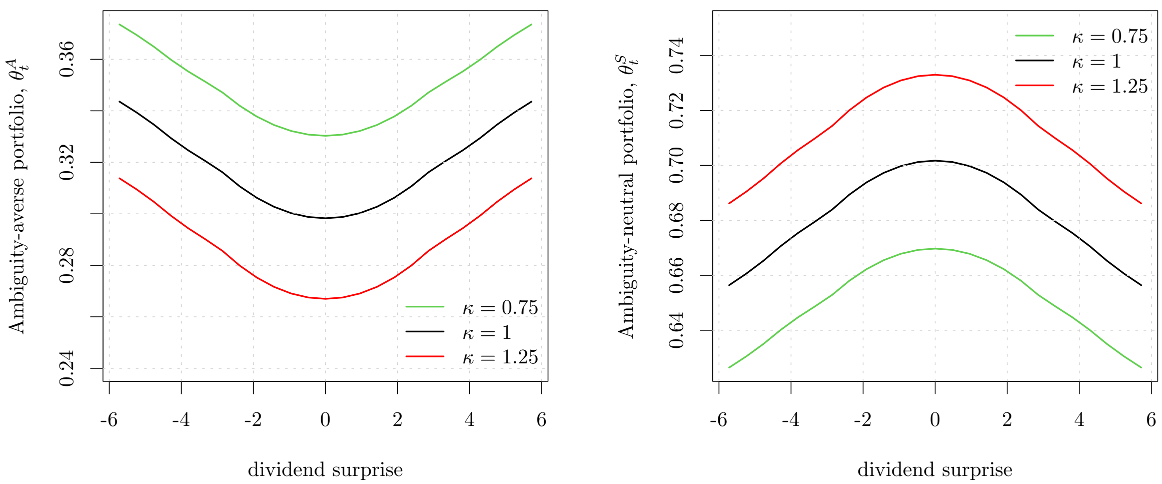

Figure 7: **Equilibrium portfolios and dividend surprise.** The figure shows the equilibrium portfolio of A-agents (left panel) and S-agents (right panel) as a function of the standard deviation estimate $\hat{\sigma}_t \equiv \sqrt{b_t/\nu_t} = 0.3$, where b_t is the sum of squared error and ν_t the degrees of freedom. Different lines corresponds to different values of the ambiguity aversion parameter κ. Parameters values: $\underline{\phi}$ and $\overline{\phi}$ such that $\sigma \in [0.2, 0.6]$, $\overline{\ell} = 1.05$, $r = 0.1$, and $\nu_t = \overline{n} = 20$.

explicitly illustrating how A-agents equilibrium weights depend on the observed standard deviation estimate $\hat{\sigma}_t$. Heterogeneity in ambiguity attitude is crucial for this result. In fact, in an economy in which A-agents are ambiguity neutral ($\kappa = 0$) there are no flows in equilibrium, even if agents differ in their degree of risk aversion.

Modelling time-variation in volatility through the learning process allows us to explicitly derive and analyze equilibrium flows. In Figure 8, we report the equilibrium flows $\Delta\theta_t^A$ defined in equation (15), from S-agents to A-agents, for a random path of dividend realizations. The solid line reports the time series of normalized surprises while the bars represent flows. Consistent with the intuition illustrated in Figures 6 and 7, following large positive and negative surprises, ambiguity-averse A-agents increase their holding of the risky asset by buying from ambiguity-neutral S-agents, $\Delta\theta_t^A > 0$. In contrast, periods with low surprises are characterized by A-agents selling to S-agents. The figure therefore reiterates the aggressiveness of ambiguity-averse agents' trades when faced with large dividend surprises and confirms, in an infinite horizon model with learning about mean and volatility, the main intuition developed in the simple two-period model of Section 2.

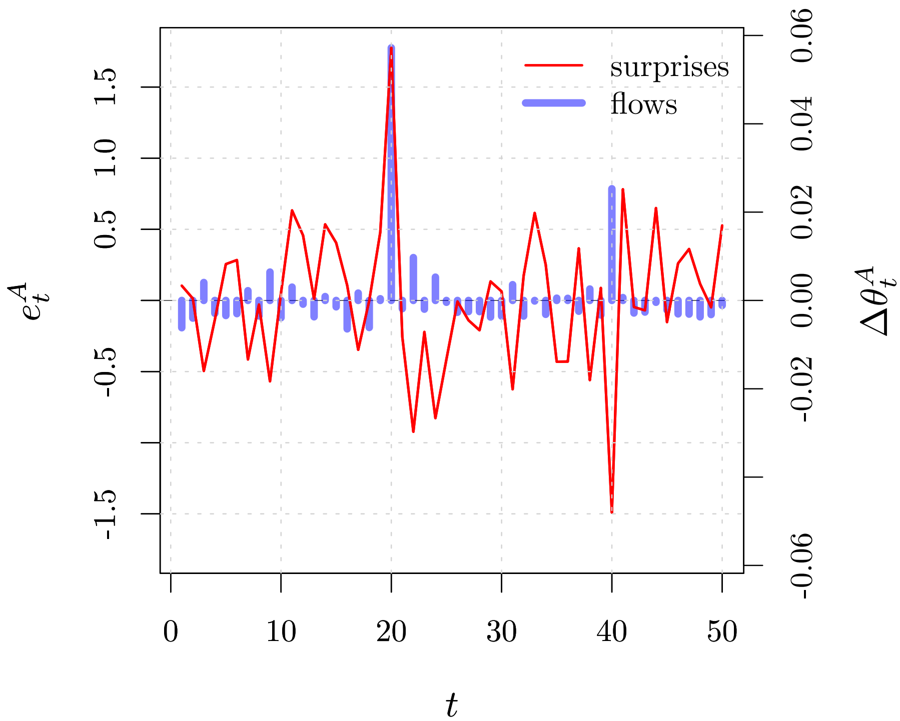

Figure 8: **Portfolio flows on a simulated path.** The figure illustrates A-agents's surprises, $e_t^A = d_{t+1} - \mu_t^A$ (left axis), and portfolio flows, $\Delta\theta_t^A$ (right axis). The state variable b_0 is chosen such that $\hat{\sigma}_0 = \sqrt{b_0/\overline{n}} = 0.3$. Parameters values: $\kappa = 1$, $\underline{\phi}$ and $\overline{\phi}$ such that the true volatility $\sigma \in [0.2, 0.6]$, $\overline{\ell} = 1.05$, $r = 0.1$, and $\nu_t = \overline{n} = 20$.

3.5 Return predictability

In this section, we explore the implications of our model for return predictability. Figure 9 shows the equilibrium risk premium as a function of the standard deviation estimate $\hat{\sigma}_t$ in the steady state of the overlapping-generation economy of Section 3.3. Not surprisingly, the risk premium is an increasing function of the estimated standard deviation. Hence, unlike the case in which variance is known and the risk premium is constant (see equation (26) in Proposition 2), the risk premium is time-varying if agents learn about the variance. This, in turn, implies that returns are predictable in our economy.

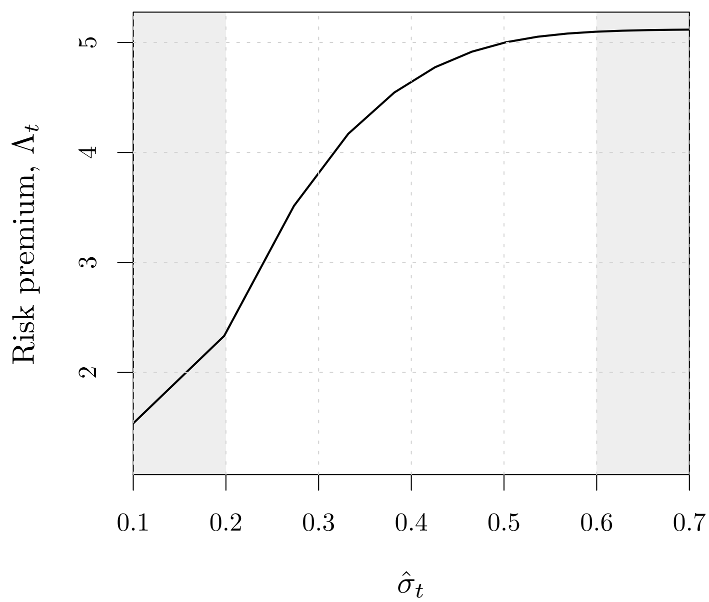

Figure 9: **Equilibrium risk premium.** The figure shows the equilibrium risk premium Λ_t as a function of the standard deviation estimate $\hat{\sigma}_t \equiv \sqrt{b_t/\nu_t}$ when both the dividend mean and variance are unknown. Parameters values: $\kappa = 1$, $\underline{\phi}$ and $\overline{\phi}$ such that the true volatility $\sigma \in [0.2, 0.6]$, $\overline{\ell} = 1.05$, $r = 0.1$, and $\nu_t = \overline{n} = 20$.

To illustrate the origin of return predictability in our model, consider first the case in which the true dividend mean is constant and unknown, while the variance is constant and known to all investors, as in Section 3.2. This case is similar to the economy studied by Lewellen and Shanken (2002). In such a setting, after positive dividend realizations, investors' estimate of the mean dividend m_t is higher than the true mean μ, and the stock is "over-priced" relative to its fundamental value. Since the true mean is lower than investors' estimate, the price will be mean reverting. An econometrician looking at the data will find that high prices predict lower returns. However, such a return predictability cannot be exploited by agents in the economy. To see why this is the case, let $\lambda_t^{\text{obj}} = r\Lambda_t^{\text{obj}} = \mu - p_t r$ denote the (per period) *objective* risk premium, and decpmpose it as follows

$$\lambda_t^{\text{obj}} = \mu - p_t r = \underbrace{\mu - \mu_t^i}_{\text{unobservable}} + \underbrace{\mu_t^i - p_t r}_{\equiv \lambda_t^i}, \tag{53}$$

with λ_t^i denoting i-agents' subjective risk premium. Using the equilibrium price p_t derived in Proposition 2, equations (25)–(26), we deduce that the subjective risk premium λ_t^i is only a deterministic function of t. Following a positive dividend realization, agents' subjective estimate of the mean μ_t^i increases and the objective expected risk premium λ_t^{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 predictability.

In contrast, when the variance is not known and there is perpetual learning, as in the model of Section 3.3, the equilibrium subjective risk premium is time-varying as it explicitly depends on the the state variable b_t. Formally, we can use equation (52) to define the objective and subjective risk premia, Λ_t^i and Λ_t^{obj}, as follows

$$p_t(m_t, b_t) = \frac{1}{r}m_t - \Lambda(b_t) = \frac{1}{r}\mu_t^i - \underbrace{\left(\frac{1}{r}(\mu_t^i - m_t) + \Lambda(b_t)\right)}_{\equiv \Lambda_t^i(b_t)}, \tag{54}$$

$$= \frac{1}{r}\mu - \underbrace{\left(\frac{1}{r}(\mu - m_t) + \Lambda(b_t)\right)}_{\equiv \Lambda_t^{\text{obj}}(b_t)}. \tag{55}$$

Because for S agents, $\mu_t^i = m_t$, equation (54) implies that $\Lambda_t(b_t)$ is agents S's subjective risk premium. As Figure 9 shows, $\Lambda_t(b_t)$ is an increasing function of $\hat{\sigma}_t \equiv \sqrt{b_t/\nu_t}$. This implies that agents expect higher returns after large positive or negative dividend surprises.[17] In a setting with unknown variance and perpetual learning, subjective risk premia are time-varying, even in the steady state. Hence, this model gives rise to return predictability that can be exploited by market participants.

Finally, note that the objective risk premium Λ_t^{obj} in equation (55) responds *asymmetrically* to dividend surprises. Positive surprises increase m_t and negative surprises decrease m_t. However, dividend surprises, regardless of their sign, increase b_t, and hence $\Lambda_t(b_t)$. Therefore, an econometrician observing ex-post dividend realization would detect a more pronounced increase in risk premia following bad dividend news than following good dividend news of equal magnitude. 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. This asymmetry, however, can only be detected ex-post

[17] 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.

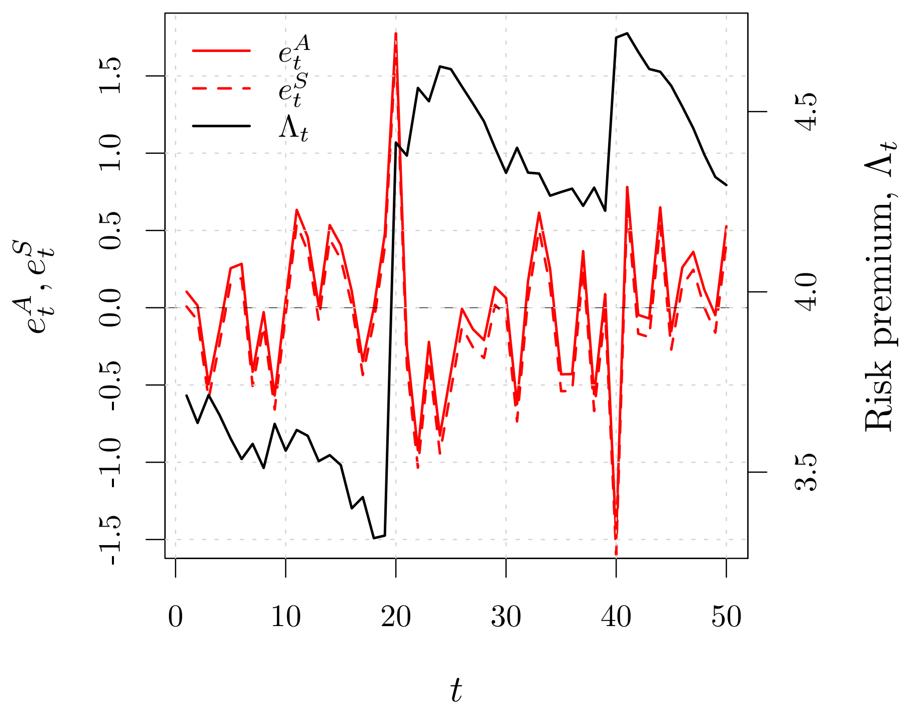

Figure 10: **Risk premium on a simulated path.** The figure illustrates surprises $e_t^i = d_{t+1} - \mu_t^i$, $i = A, S$ (left axis), with the corresponding steady state equilibrium risk premium Λ_t (right axis). The state variable b_0 is chosen such that $\hat{\sigma}_0 = \sqrt{b_0/\overline{n}} = 0.3$. Parameters values: $\kappa = 1$, $\underline{\phi}$ and $\overline{\phi}$ such that $\sigma \in [0.2, 0.6]$, $\overline{\ell} = 1.05$, $r = 0.1$, and $\nu_t = \overline{n} = 20$.

by an econometrician who knows the true mean dividend μ. Ex-ante, when agents learn from new observations, changes in the mean estimates m_t are immediately absorbed in the asset price and only the variance estimate—which is symmetric in dividend surprises—causes time variation in subjective risk premia. Figure 10 illustrates the dynamic of the risk premium $\Lambda_t(b_t)$ for a simulated random path of dividend surprises, $e_t^i = d_{t+1} - \mu_t^i$, $i = A, S$. The figure shows that the risk premium increases after large positive ($t = 20$) as well as large negative ($t = 40$) surprises, and declines gradually when dividend realizations are close to their expected value, i.e., surprises are small in magnitude.

3.6 Learning about variance vs. stochastic volatility

One might be tempted to argue that a model in which subjective variance is endogenously time-varying due to learning, as in Section 3.3, is observationally equivalent to a model with observable stochastic volatility. Although both models exhibit time-variation in volatility, they have starkly different implication for equilibrium flows. In fact, in a model with learning, a revision in the estimated variance following a new dividend observation can both increase or decrease the standard error of the mean. This is because 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, any new dividend observation can only reduce the standard error of the mean and hence its confidence interval. Because variance is known, albeit time-varying, a change in variance cannot affect the quality of past information. Therefore, in the limit with known and stochastic volatility the confidence interval of the mean collapses to a singleton and the effect of ambiguity on portfolio flows vanishes.

To illustrate this point, suppose the dividend process d_t is iid with unknown and constant mean μ and time-varying but observable variance σ_t^2. In this setting, the Generalized Least Square (GLS) estimate of the mean m_t from a history of t observations is (see, e.g., Chapter 9 in Greene, 2020)

$$m_t = \sum_{i=1}^{t} w_i d_i, \quad \text{with} \quad w_i = \frac{\frac{1}{\sigma_i^2}}{\sum_{i=1}^{t} \frac{1}{\sigma_i^2}}, \tag{56}$$

where the weight w_i represents the precision of each observation and $s_t^2 = \left(\sum_{i=1}^{t} \frac{1}{\sigma_i^2}\right)^{-1}$ the squared standard error of the mean.[18]

[18] Because dividend realizations are independent, the variance of m_t is given by

$$s_t^2 = \mathrm{var}(m_t) = \sum_{i=1}^{t} w_i^2 \underbrace{\mathrm{var}(d_i)}_{=\sigma_i^2} = \left(\frac{1}{\sum_{i=1}^{t} \frac{1}{\sigma_i^2}}\right)^2 \sum_{i=1}^{t} \left[\left(\frac{1}{\sigma_i^2}\right)^2 \sigma_i^2\right] = \frac{1}{\sum_{i=1}^{t} \frac{1}{\sigma_i^2}}.$$

At time $t+1$, the updated values of the mean and standard error, after observing the new realized dividend d_{t+1} and variance σ^2_{t+1}, are

$$m_{t+1} = (1 - w_{t+1})m_t + w_{t+1}d_{t+1}, \quad w_{t+1} = \frac{\frac{1}{\sigma^2_{t+1}}}{\frac{1}{s^2_t} + \frac{1}{\sigma^2_{t+1}}} = \frac{s^2_t}{s^2_t + \sigma^2_{t+1}} \tag{57}$$

$$\frac{1}{s^2_{t+1}} = \frac{1}{s^2_t} + \frac{1}{\sigma^2_{t+1}}. \tag{58}$$

Equation (58) shows that with stochastic but known variance, the updated standard error s_{t+1} does not depend on the new dividend realization d_{t+1} and that $s_{t+1} \leq s_t$. Hence, new observations can only *reduce* the standard error of the mean. Because the standard error controls the size of the set of priors $\mathcal{P}^\mu_t = [m - \kappa s_t, m + \kappa s_t]$ in equation (58), in a model with stochastic volatility a new dividend observation always reduces ambiguity. Dividends d_{t+1} observed in times of high volatility σ_{t+1} receive a tiny weight w_{t+1} in the updated mean m_{t+1} and only marginally reduce the standard error s_{t+1}.

In contrast, in the model of Section 3.3 where agents learn about an unknown variance, large dividend surprises ($e_{t+1} = d_{t+1} - m_t$) increase the estimated variance, directly leading to a higher estimated standard error of the mean. Therefore, the new signal affects the quality of *all* historical dividends, and agents revise their confidence interval of the mean. We conclude that a model of stochastic but know volatility would imply negligible equilibrium flows following dividend surprises, contrary to the empirical evidence.

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; and (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-S 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: First, 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. Second, 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 λ in equation (53) 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)[19] 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. The market return is taken from Kenneth French's data library.[20]

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 13-F 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.[21]

To measure surprises in the firms' future profitability, we use standardized quarterly market returns, z_r, i.e., normalized to have zero mean and unit variance using a 20-quarter rolling windows, and we 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 11 shows the relationship between changes in institutional ownership, $\Delta\theta^S$, 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

[19]Data source: https://fred.stlouisfed.org/tags/series.

[20]Data source: https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

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

OP Global Ownership database for firms listed on the NYSE, AMEX or NASDAQ exchanges and with a market capitalization 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 ownership. These changes should be interpreted relative to the substantial trend towards institutional ownership which is present since 1980, see Stambaugh (2014).[22]

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^S after positive surprises found in both data sets.

Table 1 provides details on the analysis underlying the results in Figure 11. The table shows that high as well as low standardized returns z_r are associated with low contemporaneous changes in institutional ownership, $\Delta\theta^S$, intermediate z_r-values come with an increase in institutional ownership $\Delta\theta^S$. 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).[23] The KW test confirms that $\Delta\theta^S$ differs across bins, and the post-hoc Dunn test attests that central bins have a significantly higher $\Delta\theta^S$ compared to the extreme bins (see the corresponding p-values).

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

[22]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 a market capitalization in excess to \$5 millions, and from 44% to 76% for firms with a 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).

[23]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 these results. For that purpose, a subsequent (post hoc) Dunn test (Dunn, 1964) allows for a pairwise comparison of the bins.

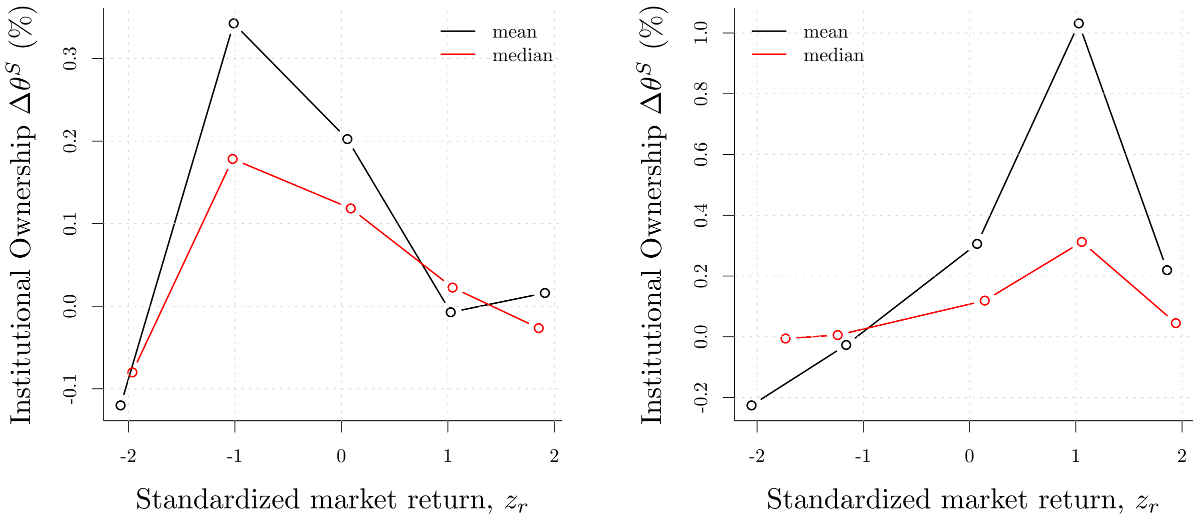

Figure 11: **Change in institutional holdings and dividend surprises.** The figure shows mean (black) and median (red) quarterly changes in institutional ownership, $\Delta\theta^S$, as a function of dividend surprises. As a proxy for surprises, we use the standardized quarterly market returns, z_r, 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, AMEX or NASDAQ with a market capitalization larger than \$5 millions. The market return is taken from Kenneth French's data library.

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 12 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 standardized 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 standardized returns. Both panels of Figure 12 show that market risk premia are higher following negative *and* positive surprises.

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 12 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).

Table 2 provides details on the analysis underlying the results in Figure 12. 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, on 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.

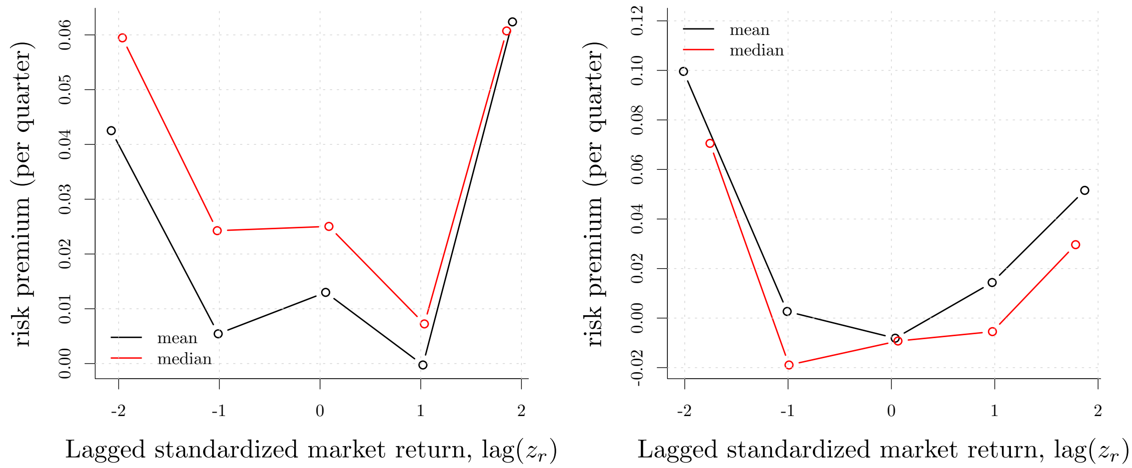

Figure 12: **Risk premia and dividend surprises.** The figure shows the mean (black) and median (red) market risk premia as a function of dividend surprises. As a proxy for surprises, we use the standardized quarterly market returns, z_r, obtained from a 20-quarter rolling window. We use z_r to group 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, AMEX or NASDAQ with a market capitalization larger than $5 millions. The market return is taken from Kenneth French's data library.

5 Conclusion

We study asset prices and portfolio flows following episodes of increased economic uncertainty through the lens of an equilibrium model in which agents learn about the mean and the volatility of the endowment process and differ in their aversion towards 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— ambiguity-neutral (subjective expected-utility) investors reduce their share in the risky asset while ambiguity-averse investors increase their share. Moreover, 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. When some of the agents are ambiguity averse, the equilibrium risk premium depends linearly on both the variance *and* the volatility of the endowment. Therefore, 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 model which is analytically tractable. We then analyze an infinite-horizon overlapping-generation model. When agents learn about the mean and the variance of the endowment process, ambiguity is time-varying and persists over time. The model highlights that three main ingredients are needed to explain flows of funds and risk premia in our setting: (i) differences in ambiguity aversion; (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.

FRED										
	mean		median			KW: $\chi^2 = 13.93$, $p = 0.0075$				
bin	z_r	$\Delta\theta^S$ (%)	z_r	$\Delta\theta^S$ (%)	n	Dunn post hoc				
1	-2.07	-0.12	-1.96	-0.08	24	bin	1	2	3	4
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	$\Delta\theta^S$ (%)	z_r	$\Delta\theta^S$ (%)	n	clustered Wilcoxon rank sum				
1	-2.05	-0.23	-1.73	-0.01	37,457	bin	1	2	3	4
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 $\Delta\theta^S$ (%) 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 $\Delta\theta^S$ (%) 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 $\Delta\theta^S$ 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, AMEX or NASDAQ and a 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 a market capitalization larger than \$5 millions. The market return is taken from Kenneth French's data library.

A Proofs

Proof of Proposition 1

The agents solve the following optimization problem

$$\max_{\theta^i} \mathbb{E}\left[-\frac{1}{\gamma} e^{-\gamma_i (W^i(1+r)+\theta^i(\tilde{d}-p(1+r))}\right]. \tag{A.1}$$

Using the normality of $\tilde{d}$, the optimal portfolio weights are:

$$\theta^i = \frac{\mu^i - p(1+r)}{\gamma\sigma^2\left(\frac{t+1}{t}\right)}, \tag{A.2}$$

where $\mu^s = m$ and $\mu^A = m - \text{sign}(\theta^A)\kappa s$. In equilibrium $\theta^A \geq 0$, therefore, imposing market clearing we obtain

$$p = \frac{1}{r+1} m - \lambda, \tag{A.3}$$

where

$$\lambda = \begin{cases} \frac{\kappa}{2}\frac{\sigma}{\sqrt{t}} + \frac{\gamma}{2}\left(\frac{t+1}{t}\right)\sigma^2 & \text{if } \kappa \leq \kappa^*, \\ \gamma\left(\frac{t+1}{t}\right)\sigma^2 & \text{if } \kappa > \kappa^*. \end{cases} \quad \text{with} \quad \kappa^* \equiv \gamma\left(\frac{t+1}{\sqrt{t}}\right)\sigma \tag{A.4}$$

∎

Lemma A.1. *Let $\bar{\lambda}^S(\sigma)$ and $\bar{\lambda}^A(\sigma)$ denote the iso-portfolios of agent S and A, respectively. For all equilibrium values λ of the risk premium in equation (11), we have that $\partial\bar{\lambda}^A(\sigma)/\partial\sigma < \partial\bar{\lambda}^S/\partial\sigma$.*

Proof. From equation (9) we derive the risk premium that A-agents require for holding a fraction θ^A of the risky asset (the iso-portfolio line) and its derivative with respect to the dividend volatility σ as

$$\bar{\lambda}^i = \frac{\kappa\sigma}{\sqrt{t}} + \gamma\theta^i\left(\frac{t+1}{t}\right)\sigma^2, \tag{A.5}$$

$$\frac{\partial\bar{\lambda}^i}{\partial\sigma} = \frac{\kappa}{\sqrt{t}} + 2\gamma\theta^i\left(\frac{t+1}{t}\right)\sigma., \quad i = S, A. \tag{A.6}$$

We prove that along the equilibrium risk premium λ in equation (11) the slope of $\bar{\lambda}^A$ is flatter than the slope of $\bar{\lambda}^S$, i.e.,

$$\frac{\partial \bar{\lambda}^A}{\partial \sigma} \;<\; \frac{\partial \bar{\lambda}^S}{\partial \sigma}. \tag{A.7}$$

Using equations (A.5)–(A.6), and market clearing, $\theta^S = 1 - \theta^A$, this is equivalent to prove

$$\frac{\kappa}{\sqrt{t}} + 2\gamma\theta^A \left(\frac{t+1}{t}\right)\sigma \;<\; 2\gamma(1-\theta^A)\left(\frac{t+1}{t}\right)\sigma, \tag{A.8}$$

or, rearranging,

$$4\gamma\theta^A \left(\frac{t+1}{t}\right)\sigma \;<\; 2\gamma\left(\frac{t+1}{t}\right)\sigma - \frac{\kappa}{\sqrt{t}}. \tag{A.9}$$

We restrict our analysis to the region where both agents are in the market, $\sigma > \frac{\sqrt{t}}{t+1}\frac{\kappa}{\gamma}$, and substitute equilibrium portfolios weights from equation (12) into the above inequality. This yields

$$2\gamma\left(\frac{t+1}{t}\right)\sigma - 2\left(\frac{\sqrt{t}\kappa}{(t+1)\sigma}\right)\left(\frac{t+1}{t}\right)\sigma \;<\; 2\gamma\left(\frac{t+1}{t}\right)\sigma - \frac{\kappa}{\sqrt{t}}, \tag{A.10}$$

$$2\frac{\kappa}{\sqrt{t}} \;>\; \frac{\kappa}{\sqrt{t}}, \tag{A.11}$$

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

Proof of Proposition 2

We first solve for the equilibrium in a fictitious finite-horizon overlapping-generation economy with horizon τ, and we then derive the equilibrium in the infinite horizon as limit for $\tau \to \infty$.

Let $p_{t,\tau}$ be the time t equilibrium price in a τ-period economy. The risky asset demand $\theta^i_{t,\tau}$, $i = A, S$ is

$$\theta^i_{t,\tau} = \frac{\mathbb{E}^i_t\left[p_{t+1,\tau-1} + d_{t+1}\right] - Rp_{t,\tau}}{\gamma \mathrm{Var}_t\left[p_{t+1,\tau-1} + d_{t+1}\right]}, \quad \tau > t,$$

where we denoted by $R \equiv (1+r)$. Using the fact that $p_{t,0} = 0$ for all t, we can construct the equilibrium in a $\tau = 1$ economy. In this economy, when both agents participate

$$\theta^i_{t,1} \quad = \quad \frac{\mathbb{E}^i_t \left[d_{t+1}\right] - Rp_{t,1}}{\gamma \mathrm{Var}_t \left[d_{t+1}\right]} = \frac{\mu^i_t - Rp_{t,1}}{\gamma\sigma^2 \left(\frac{t+1}{t}\right)},$$

where

$$d_{t+1} \sim^i \mathcal{N}\left(\mu^i_t, \sigma^2 \left(\frac{t+1}{t}\right)\right), \quad \text{with} \;\; \mu^S_t = m_t \;\; \text{and} \;\; \mu^A_t = m_t - \kappa\frac{\sigma}{\sqrt{t}}. \tag{A.12}$$

Imposing market clearing we have

$$p_{t,1} = \frac{1}{R} m_t - \Lambda_{t,1}, \tag{A.13}$$

where the risk premium $\Lambda_{t,1}$ is

$$\Lambda_{t,1} = \frac{\kappa}{2} g_{t,1}\sigma + \frac{\gamma}{2} f_{t,1}\sigma^2, \quad \text{with} \;\; g_{t,1} = \frac{1}{R}\frac{1}{\sqrt{t}}, \quad \text{and} \;\; f_{t,1} = \frac{1}{R}\left(\frac{t+1}{t}\right). \tag{A.14}$$

In a $\tau = 2$ period economy, agents demand is

$$\theta^i_{t,2} = \frac{\mathbb{E}^i_t \left[p_{t+1,1} + d_{t+1}\right] - Rp_{t,2}}{\gamma \mathrm{Var}_t \left[p_{t+1,1} + d_{t+1}\right]}, \tag{A.15}$$

where $p_{t+1,1}$ is given by equation (A.13). Because

$$m_{t+1} = \frac{t}{t+1} m_t + \frac{1}{t+1} d_{t+1},$$

using the predictive distribution (A.12) we obtain

$$\mathbb{E}^S_t[p_{t+1,1} + d_{t1}] \quad = \quad \left(1 + \frac{1}{R}\right) m_t - \Lambda_{t+1,1}, \tag{A.16}$$

$$\mathbb{E}^A_t[p_{t+1,1} + d_{t1}] \quad = \quad \left(1 + \frac{1}{R}\right) m_t - \left(1 + \frac{1}{R(t+1)}\right)\kappa\frac{\sigma}{\sqrt{t}} - \Lambda_{t+1,1}, \tag{A.17}$$

$$\mathrm{Var}_t \left[p_{t+1,1} + d_{t+1}\right] \quad = \quad \left(1 + \frac{1}{R(t+1)}\right)^2 \left(\frac{t+1}{t}\right)\sigma^2, \tag{A.18}$$

with $\Lambda_{t+1,1}$ defined in equation (A.14). Substituting in equation (A.15) and imposing market clearing we obtain

$$p_{t,2} = \left(\frac{1}{R} + \frac{1}{R^2}\right) m_t - \Lambda_{t,2}, \tag{A.19}$$

where

$$\Lambda_{t,2} = \frac{\kappa}{2} g_{t,2}\sigma + \frac{\gamma}{2} f_{t,2}\sigma^2, \tag{A.20}$$

with

$$\begin{aligned} g_{t,2} &= \frac{1}{R}\left(1 + \frac{1}{R(t+1)}\right)\frac{1}{\sqrt{t}} + \frac{1}{R^2}\frac{1}{\sqrt{t+1}}, \\ f_{t,2} &= \frac{1}{R}\left(1 + \frac{1}{R(t+1)}\right)^2\left(\frac{t+1}{t}\right) + \frac{1}{R^2}\left(\frac{t+2}{t+1}\right). \end{aligned}$$

Following similar steps, we can show that for a generic τ the equilibrium price is:

$$p_{t,\tau} = \sum_{i=1}^{\tau} \frac{1}{R^i} m_t - \Lambda_{t,\tau}, \tag{A.21}$$

where

$$\Lambda_{t,\tau} = \frac{\kappa}{2} g_{t,\tau}\sigma + \frac{\gamma}{2} f_{t,\tau}\sigma^2, \tag{A.22}$$

with

$$\begin{aligned} g_{t,\tau} &= \sum_{j=1}^{\tau} \frac{1}{R^{\tau+1-j}}\left(1 + \frac{1}{t+\tau-j+1}\sum_{i=1}^{j-1}\frac{1}{R^i}\right)\frac{1}{\sqrt{t+\tau-j}}, \\ f_{t,\tau} &= \sum_{j=1}^{\tau} \frac{1}{R^{\tau+1-j}}\left(1 + \frac{1}{t+\tau-j+1}\sum_{i=1}^{j-1}\frac{1}{R^i}\right)^2\left(\frac{t+\tau-j+1}{t+\tau-j}\right). \end{aligned}$$

Taking the limit as $\tau \to \infty$ we obtain

$$g_t = \lim_{\tau\to\infty} g_{t,\tau} = \sum_{j=1}^{\infty} \frac{1}{R^j}\left(1 + \frac{1}{r(t+j)}\right)\frac{1}{\sqrt{t+j-1}}, \tag{A.23}$$

$$f_t = \lim_{\tau\to\infty} f_{t,\tau} = \sum_{j=1}^{\infty} \frac{1}{R^j}\left(1 + \frac{1}{r(t+j)}\right)^2\frac{t+j}{t+j-1}. \tag{A.24}$$

Hence the equilibrium price in the infinite-horizon overlapping generation economy is

$$p_t = \frac{1}{r} m_t - \Lambda_t, \tag{A.25}$$

with

$$\Lambda_t \;=\; g_t \frac{\kappa}{2}\sigma + f_t \frac{\gamma}{2}\sigma^2, \tag{A.26}$$

and g_t, and f_t given in equations (A.23) and (A.24), respectively.

To determine equilibrium weights we start from the expression for the agents' optimal asset demand

$$\theta_t^i = \frac{\mathbb{E}_t^i\left[p_{t+1} + d_{t+1}\right] - (1+r)p_t}{\gamma \mathrm{Var}_t\left[p_{t+1} + d_{t+1}\right]}, \quad i = S, A. \tag{A.27}$$

Direct computation using the equilibrium price in equation (A.25) yields:

$$\begin{aligned}
\mathbb{E}_t^S[p_{t+1} + d_{t+1}] &= \left(1 + \frac{1}{r}\right) m_t - g_{t+1}\frac{\kappa}{2}\sigma - f_{t+1}\frac{\gamma}{2}\sigma^2, \\
\mathbb{E}_t^A[p_{t+1} + d_{t+1}] &= \left(1 + \frac{1}{r}\right)\left(m_t - \kappa\frac{\sigma}{\sqrt{t}}\right) - g_{t+1}\frac{\kappa}{2}\sigma - f_{t+1}\frac{\gamma}{2}\sigma^2, \\
\mathrm{Var}_t^i[p_{t+1} + d_{t+1}] &= \left(1 + \frac{1}{r(t+1)}\right)^2 \left(\frac{t+1}{t}\right)\sigma^2, \quad i = A, S.
\end{aligned}$$

Substituting these expressions in equation (A.27), we obtain the following equilibrium weights:

$$\begin{aligned}
\theta_t^A &= \frac{-\frac{1+r}{r}\frac{1}{\sqrt{t}}\kappa\sigma + \left[(1+r)g_t - g_{t+1}\right]\frac{\kappa}{2}\sigma + \left[(1+r)f_t - f_{t+1}\right]\frac{\gamma}{2}\sigma^2}{\gamma\left(1 + \frac{1}{r(t+1)}\right)^2\left(\frac{t+1}{t}\right)\sigma^2} \\
&= \frac{1}{2} - \frac{\kappa}{2\gamma}\left(\frac{r\sqrt{t}}{1 + r(t+1)}\right)\frac{1}{\sigma}
\end{aligned}$$

and

$$\begin{aligned}
\theta_t^S &= \frac{\left[(1+r)g_t - g_{t+1}\right]\frac{\kappa}{2}\sigma + \left[(1+r)f_t - f_{t+1}\right]\frac{\gamma}{2}\sigma^2}{\gamma\left(1 + \frac{1}{r(t+1)}\right)^2\left(\frac{t+1}{t}\right)\sigma^2} \\
&= \frac{1}{2} + \frac{\kappa}{2\gamma}\left(\frac{r\sqrt{t}}{1 + r(t+1)}\right)\frac{1}{\sigma}
\end{aligned}$$

$\blacksquare$

B Learning about variance: Technical appendix

In this appendix, we derive the predictive distribution of the dividend when the variance is not known (Section B.1) and provide details of data processing with information leakage (Section B.2). The principles of Bayesian data analysis can be found, e.g., in the textbook of Gelman, Carlin, Stern, Dunson, Vehtari, and Rubin (2020).

B.1 Predictive distribution of the dividend

Agents of generation t base their belief about the dividend d_{t+1}, on which their terminal wealth depends, on the prior information they receive in their first period of life. This is characterized by the state variables m_t, b_t, n_t, and ν_t. If the precision $\phi = 1/\sigma^2$ has a truncated Gamma distributed with shape parameter b and ν degrees of freedom is, its density is given by

$$p(\phi|b,\nu) = \frac{1}{C(b_t,\nu_t;\underline{\phi},\overline{\phi})}\phi^{\frac{\nu_t}{2}-1}e^{-\phi\frac{b_t}{2}}\mathbf{1}_{[\underline{\phi},\overline{\phi}]}, \quad \phi \sim \mathrm{TG}\left[\frac{\nu_t}{2},\frac{b_t}{2};\underline{\phi},\overline{\phi}\right], \quad 0<\underline{\phi}<\overline{\phi}<\infty,$$

with $\mathbf{1}$ the indicator function and

$$C(b_t,\nu_t;\underline{\phi},\overline{\phi}) \;=\; \int_{\underline{\phi}}^{\overline{\phi}} \phi^{\frac{\nu_t}{2}-1}e^{-\phi\frac{b_t}{2}}d\phi = \left(\frac{b_t}{2}\right)^{-\frac{k_t}{2}}\left[\Gamma\left(\frac{\nu_t}{2},\underline{\phi}\frac{b_t}{2}\right)-\Gamma\left(\frac{\nu_t}{2},\overline{\phi}\frac{b_t}{2}\right)\right]. \quad \text{(B.1)}$$

The function $\Gamma(x,y)$ is the upper incomplete Gamma function defined as

$$\Gamma(x,y) = \int_y^{\infty}\phi^{x-1}e^{-\phi}d\phi.$$

Definition B.1 (Dampened t-distribution). *Let ϕ be a truncated Gamma random variable,*

$$\phi \sim TG\left[\frac{\nu}{2},\frac{\nu}{2};\underline{\phi},\overline{\phi}\right], \quad 0<\underline{\phi}<\overline{\phi}\le\infty,$$

and x a conditionally Normal random variable with mean 0 and precision ϕ,

$$x \sim \mathcal{N}(0,1/\phi).$$

Then, the distribution of x is a "dampened t-distribution" with ν degrees of freedom

$$x \sim t_\nu^D[\underline{\phi}, \overline{\phi}],$$

and its density is given by

$$\begin{aligned} f(x) &= \int_{\underline{\phi}}^{\overline{\phi}} \sqrt{\frac{\phi}{2\pi}} e^{-\frac{1}{2}\phi x^2} \frac{1}{C(\nu, \nu; \underline{\phi}, \overline{\phi})} \phi^{\frac{\nu}{2}-1} e^{-\phi\frac{\nu}{2}} d\phi \\ &= \frac{1}{C(\nu, \nu; \underline{\phi}, \overline{\phi})} \sqrt{\frac{1}{2\pi}} \int_{\underline{\phi}}^{\overline{\phi}} \phi^{\frac{\nu+1}{2}-1} e^{-\phi\frac{\nu}{2} - \frac{1}{2}\phi x^2} d\phi \\ &= \sqrt{\frac{1}{2\pi}} \frac{C(\nu(1+\frac{x^2}{\nu}), \nu+1; \underline{\phi}, \overline{\phi})}{C(\nu, \nu; \underline{\phi}, \overline{\phi})}, \end{aligned}$$

with $C(\cdot, \cdot; \underline{\phi}, \overline{\phi})$ a normalizing constant defined in equation (B.1).

Since ϕ has finite support and is especially bound away from 0, the fat tails of x are dampened. As a consequence, its moment generating function is finite, and, thus, all its moments exist and are finite, see Bakshi and Skoulakis (2010) for a proof. If $\underline{\phi} \to 0$ and $\overline{\phi} \to \infty$, the distribution of x becomes a Student-t distribution. In this limit, fat tails emerge and moments of order $\geq \nu$ do not exist.

Definition B.2 (Non-standardized dampened t-distribution). *A random variable y has a non-standardized dampened t-distribution with mean m, shape b, variance scale parameter v^2, ν degrees of freedom and truncation bounds $\underline{\phi}$, $\overline{\phi}$,*

$$y \sim t_\nu^D[m, b, v^2; \underline{\phi}, \overline{\phi}]$$

if

$$\begin{aligned} y|\phi &\sim \mathcal{N}(\mu, v^2/\phi), \\ \phi &\sim TG\left[\frac{\nu}{2}, \frac{b}{2}; \underline{\phi}, \overline{\phi}\right], \quad f(\phi) = \frac{1}{C(b, \nu; \underline{\phi}, \overline{\phi})} \phi^{\frac{\nu}{2}-1} e^{-\frac{b}{2}\phi} \mathbf{1}_{[\underline{\phi}, \overline{\phi}]}. \end{aligned}$$

Then the random variable $\frac{y-\mu}{v\sqrt{b/\nu}}$ has a dampened Student-t distribution, as per Definition B.1, with truncation bounds at $\frac{b}{\nu}\underline{\phi}$ and $\frac{b}{\nu}\overline{\phi}$,

$$\frac{y-\mu}{v\sqrt{b/\nu}} \sim t_\nu^D\left[\frac{b}{\nu}\underline{\phi}, \frac{b}{\nu}\overline{\phi}\right].$$

Definition: The stochastic variable y is non-standardized dampened t-distribution with mean μ, shape b, variance scaling parameter v^2 and truncated $\phi \in [\underline{\phi}, \overline{\phi}]$ if

$$\begin{aligned}
y|\phi &\sim N(\mu, v^2/\phi), \\
\phi &\sim \mathrm{TG}[\frac{\nu}{2}, \frac{b}{2}; \underline{\phi}, \overline{\phi}], \quad f(\phi) = \frac{1}{C(b,\nu;\underline{\phi},\overline{\phi})}\phi^{\frac{\nu}{2}-1}e^{-\frac{b}{2}\phi}\mathbf{1}_{[\underline{\phi},\overline{\phi}]}, \\
y &\sim t_\nu^D\left[\mu, b, v^2; \underline{\phi}, \overline{\phi}\right]
\end{aligned}$$

Then $\frac{y-\mu}{v\sqrt{b/\nu}}$ is dampened Student t-distributed with truncation bounds at $\frac{b}{\nu}\underline{\phi}$ and $\frac{b}{\nu}\overline{\phi}$,

$$\frac{y-\mu}{v\sqrt{b/\nu}} \sim t_\nu^D\left[\frac{b}{\nu}\underline{\phi}, \frac{b}{\nu}\overline{\phi}\right].$$

Lemma B.1. *Consider a subjective Normal/inverse-Gamma prior for μ and σ with parameters μ_t^i, β_t^i, n_t, and ν_t. The predictive distribution of d_{t+1} is then a dampened Student-t,*

$$d_{t+1}|\mu^i, \beta^i, n_t, \nu_t \sim^i t_\nu^D\left[\mu^i, \beta^i, \frac{n_t+1}{n_t}; \underline{\phi}, \overline{\phi}\right].$$

Proof: With the given subjective prior, the predictive density of d_{t+1} is conditionally normal

$$\begin{aligned}
f(d_{t+1}|\phi, \mu^i, n_t) &= \int_{-\infty}^{+\infty} f(d_{t+1}|\mu,\phi)p(\mu^+|\phi,\mu^i,n_t)d\mu \\
&= \int_{-\infty}^{+\infty} \sqrt{\frac{\phi}{2\pi}}e^{-\frac{1}{2}\phi(d_{t+1}-\mu)^2}\sqrt{\frac{n_t\phi}{2\pi}}e^{-\frac{1}{2}n_t\phi(\mu-\mu^i)^2}d\mu \\
&= \sqrt{\frac{n_t\phi}{(n_t+1)2\pi}}e^{-\frac{1}{2}\frac{n_t\phi}{(n_t+1)}(d_{t+1}-\mu^i)^2}\int_{-\infty}^{+\infty}\sqrt{\frac{(n_t+1)\phi}{2\pi}}e^{-\frac{1}{2}(n_t+1)\phi\left(\mu-\frac{d_{t+1}-n_t\mu^i}{n_t+1}\right)^2}d\mu \\
&= \sqrt{\frac{n_t\phi}{(n_t+1)2\pi}}e^{-\frac{1}{2}\frac{n_t\phi}{(n_t+1)}(d_{t+1}-\mu^i)^2}, \\
d_{t+1}|\phi,\mu^i,n_t &\sim N\left(\mu^i, \frac{n_t+1}{n_t\phi}\right).
\end{aligned}$$

The unconditional density of d_{t+1} can be determined from the conditional density by integrating out the precision ϕ.

$$\begin{aligned}
f(d_{t+1}|\mu^i,\beta^i,n_t,\nu_t) &= \int_{\underline{\phi}}^{\overline{\phi}} f(d_{t+1}|\phi,\mu^i,n_t)p(\phi|\beta^i,n_t)d\phi, \\
&= \int_{\underline{\phi}}^{\overline{\phi}} \sqrt{\frac{n_t\phi}{(n_t+1)2\pi}} e^{-\frac{1}{2}\frac{n_t\phi}{(n_t+1)}(d_{t+1}-\mu^i)^2} \frac{1}{C(\beta^i,\nu_t;\underline{\phi},\overline{\phi})}\phi^{\frac{\nu_t}{2}-1}e^{-\phi\frac{\beta^i}{2}}d\phi, \\
&= \frac{1}{C(\beta^i,\nu_t;\underline{\phi},\overline{\phi})}\sqrt{\frac{n_t}{(n_t+1)2\pi}}\int_{\underline{\phi}}^{\overline{\phi}} \phi^{\frac{\nu_t+1}{2}-1}e^{-\phi\frac{\beta^i}{2}-\frac{1}{2}\frac{n_t\phi}{(n_t+1)}(d_{t+1}-\mu^i)^2}d\phi, \\
&= \sqrt{\frac{n_t}{(n_t+1)2\pi}}\frac{C(\beta^i+\frac{n_t}{(n_t+1)}(d_{t+1}-\mu^i)^2,\nu_t+1;\underline{\phi},\overline{\phi})}{C(\beta^i,\nu_t;\underline{\phi},\overline{\phi})}, \\
\left(\frac{d_{t+1}-\mu^i}{\sqrt{\frac{n_t+1}{n_t}}\sqrt{\frac{\beta^i}{\nu_t}}}\right)|\mu^i,\beta^i,n_t,\nu_t &\sim t^D_{\nu_t}[\frac{\beta^i}{\nu_t}\underline{\phi},\frac{\beta^i}{\nu_t}\overline{\phi}], \\
d_{t+1}|\mu^i,\beta^i,n_t,\nu_t &\sim t^D_{\nu_t}\left[\mu^i,\beta^i,\frac{n_t+1}{n_t};\underline{\phi},\overline{\phi}\right].
\end{aligned}$$

■

B.2 Data processing with information leakage

The information set of generation t about the unknown dividend mean μ and the precision ϕ is given by the priors in equations (29) and (30). The truncated Gamma density of ϕ has ν_t degrees of freedom and is restricted to a support within the interval $[\underline{\phi},\overline{\phi}]$ and given by

$$p(\phi|b_t,\nu_t) = \frac{1}{C(b_t,\nu_t;\underline{\phi},\overline{\phi})}\phi^{\frac{\nu_t}{2}-1}e^{-\phi\frac{b_t}{2}}\mathbf{1}_{[\underline{\phi},\overline{\phi}]}, \quad \phi \sim \mathrm{TG}\left[\frac{\nu_t}{2},\frac{b_t}{2};\underline{\phi},\overline{\phi}\right], \quad 0<\underline{\phi}<\overline{\phi}<\infty,$$

with $\mathbf{1}$ the indicator function and $C(b_t,\nu_t,\underline{\phi},\overline{\phi})$ defined in equation (B.1).

The conditional density of μ with n_t (effective) observations is

$$\sqrt{\frac{n_t\phi}{2\pi}}e^{-\frac{1}{2}n_t\phi(\mu-m_t)^2}, \quad \mu|\phi \sim \mathcal{N}\left(m_t,\frac{1}{n_t\phi}\right).$$

During the transfer of this information from generation t to generation $t+1$ —and before the next dividend d_{t+1} is observed— part of the information content is lost. For a formal modelling of the information leakage, we borrow from the so-called discount factor approach of West and Harrison (2006), the fading memory model of Nagel and Xu (2021), and the dampened t-distribution of Bakshi and Skoulakis (2010). We model information leakage in form of shocks that add noise to priors before they are updated with the information contained in the dividend d_{t+1}. Initially, with an effective number of observations n_0, the initial number of degrees of freedom v_0 is eventually set to $\nu_0 = n_0 - 1$. However, as we see below, due to information leakage, n_t and ν_t do not grow linearly with new observations but converge to a joint upper limit, which is defined by the extent of information leakage, controlled by the parameter ω below. Information leakage of the prior about μ is modelled via an additive Gaussian shock η^{μ}_{t+1} with mean 0 and variance $\left(\frac{1}{\omega} - 1\right) s_t^2$ where s_t is the standard error of the time t posterior and $\omega \in [0, 1]$. The shock η^{μ}_{t+1} is independent of the estimation error in μ. After absorbing this "leakage shock", we denote the posterior μ^+,

$$\mu^+ | \phi, m_t, n_t \sim \mathcal{N}\left(m_t, \frac{1}{\omega}\frac{1}{n_t \phi}\right), \quad \omega \in [0, 1].$$

This noisy posterior is then updated with the information contained in the dividend d_{t+1}. The following posterior is transferred to both agents of generation $t+1$ in form of the updated state variables m_{t+1} and n_{t+1}

$$
\begin{aligned}
p(\mu|d_{t+1}, \phi, m_t, n_t) &\propto f(d_{t+1}|\mu, \phi) p(\mu^+|\phi, m_t, n_t, \omega) \\
&= \sqrt{\frac{\phi}{2\pi}} e^{-\frac{1}{2}\phi(d_{t+1}-\mu)^2} \sqrt{\frac{\omega n_t \phi}{2\pi}} e^{-\frac{1}{2}\omega n_t \phi(\mu - m_t)^2}, \\
&= \sqrt{\frac{\omega n_t \phi}{(\omega n_t + 1)2\pi}} e^{-\frac{1}{2}\frac{\omega n_t \phi}{(\omega n_t+1)}(d_{t+1}-m_t)^2} \\
&\times \sqrt{\frac{(\omega n_t + 1)\phi}{2\pi}} e^{-\frac{1}{2}(\omega n_t+1)\phi\left(\mu - \frac{d_{t+1}+\omega n_t m_t}{\omega n_t + 1}\right)^2}, \\
&\propto \sqrt{\frac{(\omega n_t + 1)\phi}{2\pi}} e^{-\frac{1}{2}(\omega n_t+1)\phi\left(\mu - \frac{d_{t+1}+\omega n_t m_t}{\omega n_t + 1}\right)^2}, \\
&= \sqrt{\frac{n_{t+1}\phi}{2\pi}} e^{-\frac{1}{2}(n_{t+1})\phi(\mu - m_{t+1})^2}, \\
\mu|d_{t+1}, \phi, m_t, n_t &\sim N\left(m_{t+1}, \frac{1}{n_{t+1}\phi}\right), \\
e_{t+1} &= d_{t+1} - m_t, \\
m_{t+1} &= m_t + \frac{1}{n_{t+1}} e_{t+1}, \\
n_{t+1} &= \omega n_t + 1.
\end{aligned}
$$

To update also the ϕ prior with the $t+1$ dividend, we must first determine the distribution of d_{t+1} conditional of ϕ.

$$
\begin{aligned}
f(d_{t+1}|\phi, m_t, n_t) &= \int_{-\infty}^{+\infty} f(d_{t+1}|\mu, \phi) p(\mu^+|\phi, m_t, n_t) d\mu, \\
&= \int_{-\infty}^{+\infty} \sqrt{\frac{\phi}{2\pi}} e^{-\frac{1}{2}\phi(d_{t+1}-\mu)^2} \sqrt{\frac{\omega n_t \phi}{2\pi}} e^{-\frac{1}{2}\omega n_t \phi(\mu - m_t)^2} d\mu, \\
&= \sqrt{\frac{\omega n_t \phi}{(\omega n_t + 1)2\pi}} e^{-\frac{1}{2}\frac{\omega n_t \phi}{(\omega n_t+1)}(d_{t+1}-m_t)^2} \int_{-\infty}^{+\infty} \sqrt{\frac{(\omega n_t + 1)\phi}{2\pi}} e^{-\frac{1}{2}(\omega n_t+1)\phi\left(\mu - \frac{d_{t+1}-\omega n_t m_t}{\omega n_t + 1}\right)^2} d\mu, \\
&= \sqrt{\frac{\omega n_t \phi}{(\omega n_t + 1)2\pi}} e^{-\frac{1}{2}\frac{\omega n_t \phi}{(\omega n_t+1)}(d_{t+1}-m_t)^2}, \\
d_{t+1}|\phi, m_t, n_t &\sim N\left(m_t, \frac{\omega n_t + 1}{\omega n_t \phi}\right).
\end{aligned}
$$

The $t+1$ posterior of the precision ϕ under information leakage is subject to a multiplicative shock $1/\omega\eta^{\phi}_{t+1}$ that is generalized-beta distributed (see Bakshi and Skoulakis, 2010, equation (35)), leading to a posterior denoted ϕ^+, that is again a Gamma distribution truncated at the same bounds,

$$\phi^+|b_t, \nu_t \sim \text{TG}\left[\omega\frac{\nu_t}{2}, \omega\frac{b_t}{2}; \underline{\phi}, \overline{\phi}\right], \quad 0 < \underline{\phi} < \overline{\phi} < \infty.$$

Updating with d_{t+1} leads to the posterior which is transferred to the agents of generation $t+1$ in the form of b_{t+1} and ν_{t+1}.

$$\begin{aligned}
p(\phi|d_{t+1}, m_t, b_t, n_t, \nu_t; \underline{\phi}, \overline{\phi}) &\propto f(d_{t+1}|\phi, m_t, n_t)p(\phi^+|b_t, \nu_t; \underline{\phi}, \overline{\phi}), \\
&= \sqrt{\frac{\omega n_t \phi}{(\omega n_t + 1)2\pi}} e^{-\frac{1}{2}\frac{\omega n_t \phi}{(\omega n_t+1)}(d_{t+1}-m_t)^2} \\
&\times \frac{1}{C(b_t, \omega\nu_t; \underline{\phi}, \overline{\phi})}\phi^{\frac{\omega\nu_t}{2}-1}e^{-\phi\omega\frac{b_t}{2}}\mathbf{1}_{[\underline{\phi},\overline{\phi}]}, \\
&\propto \phi^{\frac{\omega\nu_t+1}{2}-1}e^{-\phi\left(\omega\frac{b_t}{2}+\frac{1}{2}\frac{(d_{t+1}-m_t)^2\omega n_t}{\omega n_t+1}\right)}\mathbf{1}_{[\underline{\phi},\overline{\phi}]}, \\
&\propto \frac{1}{C(b_{t+1}, \nu_{t+1}; \underline{\phi}, \overline{\phi})}\phi^{\frac{\nu_{t+1}}{2}-1}e^{-\phi\frac{b_{t+1}}{2}}\mathbf{1}_{[\underline{\phi},\overline{\phi}]},
\end{aligned}$$

$$\begin{aligned}
\phi|d_{t+1}, b_t, m_t, n_t, \nu_t; \underline{\phi}, \overline{\phi} &\sim \text{TG}\left[\frac{\nu_{t+1}}{2}, \frac{b_{t+1}}{2}; \underline{\phi}, \overline{\phi}\right], \\
e_{t+1} &= d_{t+1} - m_t, \\
b_{t+1} &= \omega b_t + \omega\frac{n_t}{n_{t+1}}e_{t+1}^2, \\
n_{t+1} &= \omega n_t + 1, \\
\nu_{t+1} &= \omega\nu_t + 1.
\end{aligned}$$

When $\omega < 1$ and t large, the effective number of observations n_t and the degrees of freedom ν_t converge to the same upper limit

$$\lim_{t\to\infty} n_t = \lim_{t\to\infty} \nu_t \equiv \overline{n} = \frac{1}{1-\omega}.$$

We implement a desired asymptotic effective number of observations $\overline{n}$ by choosing

$$\omega = \frac{\overline{n}-1}{\overline{n}}.$$

C Numerical procedure to determine the equilibrium

Both types of agents of generation t know the state variables m_t, b_t, n_t and ν_t from the "information processing" step. Agents' memory does not fade over the lifetime, i.e., they act rationally according to their utility functions. However, generation-t agents anticipate that information gets lost, when information is transferred to generation $t+1$. Since this generation buys the asset from generation-t agents, agents of generation t must anticipate the demand of generation-$t+1$ agents and the price p_{t+1}.

Let $\Delta\mu_t^i$ denote the agents i's adjustment to m_t when forming beliefs about the dividend mean, that is $\Delta\mu_t^i = m_t - \mu_t^i$. The density of the dividend d_{t+1} under the subjective prior of agents i is dampened t-distributed

$$\begin{aligned} d_{t+1} &\sim^i t_{\nu_t}^D\left[\mu_t^i, \beta_t^i, \frac{n_t+1}{n_t}\right], \\ \frac{d_{t+1} - \mu_t^i}{\sqrt{\frac{n_t+1}{n_t}}\sqrt{\frac{\beta_t^i}{\nu_t}}} &\sim^i t_{\nu_t}^D\left[\frac{\beta_t^i}{\nu_t}\underline{\phi}, \frac{\beta_t^i}{\nu_t}\overline{\phi}\right]. \end{aligned}$$

The individual surprise e_{t+1}^i is defined relative to the subjective expectation μ_t^i

$$\begin{aligned} e_{i+1}^i &= d_{t+1} - \mu_t^i, \\ \frac{e_{i+1}^i}{\sqrt{\frac{n_t+1}{n_t}}\sqrt{\frac{\beta_t^i}{\nu_t}}} &\sim^i t_{\nu_t}^D\left[\frac{\beta_t^i}{\nu_t}\underline{\phi}, \frac{\beta_t^i}{\nu_t}\overline{\phi}\right]. \end{aligned}$$

While agents have subjective beliefs about the distribution of d_{t+1}, they agree on the way information is handed over to the next generation (including the information leakage during the transition of information) and how the next generation will learn from observing d_{t+1}.

We express the mechanics of updating the state variables in terms of e^i_{t+1}

$$\begin{aligned}
n_{t+1} &= \omega n_t + 1, \\
\nu_{t+1} &= \omega \nu_t + 1, \\
m_{t+1} &= \frac{\omega n_t}{\omega n_t + 1} m_t + \frac{1}{\omega n_t + 1} d_{t+1}, \\
&= m_t - \frac{1}{n_{t+1}} \Delta \mu^i_t + \frac{1}{n_{t+1}} e^i_{t+1}, \\
b_{t+1} &= \omega b_t + \frac{1}{2} \frac{\omega n_t}{\omega n_t + 1} (m_t - d_{t+1})^2, \\
&= \omega b_t + \frac{\omega n_t}{n_{t+1}} (\Delta \mu^i_t - e^i_{t+1})^2.
\end{aligned}$$

We assume that t is large, so n_t and ν_t have already reached their asymptotic limit $\overline{n}$. We conjecture that in an economy that lasts for τ generations the price can be written as a function of the state variables $p_t = h(\tau) m_t - \Lambda(b_t, \tau)$ and all agents agree on this functional form. For $\tau \to \infty$, we can write p_{t+1} as

$$\begin{aligned}
p_{t+1} &= \frac{1}{r} m_{t+1} - \Lambda(b_{t+1}), \\
&= \left(\frac{1}{r} m_t - \frac{1}{r\overline{n}} \Delta \mu^i_t + \frac{1}{r\overline{n}} e^i_{t+1} \right) - \Lambda(b_{t+1}), \\
&= \left(\frac{1}{r} m_t - \frac{1}{r\overline{n}} \Delta \mu^i_t + \frac{1}{r\overline{n}} e^i_{t+1} \right) - \Lambda(b_{t+1}).
\end{aligned}$$

Under this conjecture, the budget constraint becomes independent of m_t and only depends on $\Lambda(b_t)$ and $\Lambda(b_{t+1})$.

$$
\begin{aligned}
W_{t+1}^i(\theta) &= (W_t^i - \theta p_t)(1+r) + \theta(d_{t+1} + p_{t+1}),\\
&= (W_t^i - \theta\left(\frac{1}{r}m_t - \Lambda(b_t)\right))(1+r)\\
&+ \theta\left(1+\frac{1}{r}\right)m_t - \theta(1+\frac{1}{r\overline{n}}\Delta\mu_t^i) + \theta\left(1+\frac{1}{r\overline{n}}e_{t+1}^i\right) - \theta\Lambda(b_{t+1}),\\
&= (W_t^i + \theta\Lambda(b_t))(1+r) - \theta\left(1+\frac{1}{r\overline{n}}\Delta\mu_t^i\right)\\
&+ \theta\left(1+\frac{1}{r\overline{n}}e_{t+1}^i\right) - \theta\Lambda(b_{t+1}).
\end{aligned}
$$

The expected utility of agents i is then

$$
\begin{aligned}
\mathbb{E}^i(u(W_{t+1}^i(\theta))) &= \frac{1}{C(b_t^i,\overline{n};\underline{\phi},\overline{\phi})}\int_{\underline{\phi}}^{\overline{\phi}} \mathbb{E}^i(u(W_{t+1}^i(\theta)|\phi,b_t)\phi^{\frac{\overline{n}}{2}-1}e^{-\phi\frac{\beta_t^i}{2}}d\phi,\\
\mathbb{E}^i(u(W_{t+1}^i(\theta))|\phi,b_t) &=\\
&= -\frac{1}{\gamma}\exp\left\{-\gamma\left[(1+r)(W_t+\theta\Lambda(b_t)) - \theta\Big(1+\frac{1}{r\overline{n}}\Big)\Delta\mu_t^i\right]\right\}\\
&\times \sqrt{\frac{\overline{n}\phi}{(\overline{n}+1)2\pi}}\int_{-\infty}^{\infty}\exp\left\{-\gamma\theta\left[\left(1+\frac{1}{r\overline{n}}\right)e_{t+1}^i - \Lambda(b_{t+1})\right]\right\}\\
&\times\ \exp\left\{-\frac{1}{2}\frac{\overline{n}\phi}{(\overline{n}+1)}{e_{t+1}^i}^2\right\}de_{t+1}^i\\
b_{t+1} &= \omega\left(b_t + (\Delta\mu_t^i - e_{t+1}^i)^2\right).
\end{aligned}
$$

Since e_{t+1}^i is dampened t, $\Lambda(b_t) \geq 0$, and $\lim_{b_t\to\infty}\Lambda(b_t) < \infty$, the expected utility is well defined.

The marginal utility is computed as

$$
\begin{aligned}
\frac{d\mathbb{E}^i(u(W^i_{t+1}(\theta))|b_t)}{d\theta} &= \frac{1}{C(b^i_t,\overline{n};\underline{\phi},\overline{\phi})} \\
&\times \int_{\underline{\phi}}^{\overline{\phi}} \frac{d\mathbb{E}(u(W^i_{t+1}(\theta))|\phi,b_t)}{d\theta}\phi^{\frac{\overline{n}}{2}-1}e^{-\phi\frac{\beta^i_t}{2}}d\phi, \\
\frac{d\mathbb{E}^i(u(W_{t+1}(\theta))|\phi,b_t)}{d\theta} &= -\gamma\left[(1+r)\Lambda(b_t)-\Big(1+\frac{1}{r\overline{n}}\Big)\Delta\mu^i_t\right]\mathbb{E}^i(u(W_{t+1}(\theta))|\phi,b_t) \\
+\frac{1}{\gamma}\exp\left\{-\gamma\left[(1+r)(W_t+\theta\Lambda(b_t))-\theta\Big(1+\frac{1}{r\overline{n}}\Big)\Delta\mu^i_t\right]\right\} \\
\times\sqrt{\frac{\overline{n}\phi}{(\overline{n}+1)2\pi}}\int_{-\infty}^{\infty}\gamma\left[\left(1+\frac{1}{r\overline{n}}\right)e^i_{t+1}-\Lambda(b_{t+1})\right] \\
&\times\exp\left\{-\gamma\theta\left[\left(1+\frac{1}{r\overline{n}}\right)e^i_{t+1}-\Lambda(b_{t+1})\right]\right\} \\
&\times\ \exp\left\{-\frac{1}{2}\frac{\overline{n}\phi}{(\overline{n}+1)}{e^i_{t+1}}^2\right\}de^i_{t+1}, \\
b_{t+1} &= \omega(b_t+\left(\Delta\mu^i_t-e^i_{t+1})^2\right)
\end{aligned}
$$

We determine the function $\Lambda(b_t)$ as a fixed point via value function iteration. When both agents invest at a given b_t, they take $\Lambda(b_t)$ as given and optimize their holding θ^i_t via the first-order condition $\frac{d\mathbb{E}^i(u)}{d\theta^i}=0$. The equilibrium risk premium $\Lambda(b_t)$ satisfies market clearing, $\theta^S(b_t)+\theta^A(b_t)=1$.

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.

Bakshi, G., and G. Skoulakis, 2010, "Do subjective expectations explain asset pricing puzzles?," *Journal of Financial Economics*, 98(3), 462–477.

Barber, B. M., X. Huang, T. Odean, and C. Schwarz, 2021, "Attention-induced trading and returns: Evidence from Robinhood users," *The Journal of Finance*, 77(6), 3141–3190.

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, 2016, "Asset pricing when 'This time is different'," *The Review of Financial Studies*, 30(2), 505–535.

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

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

De Long, J. B., A. Shleifer, L. H. Summers, and R. J. Waldmann, 1990, "Noise trader risk in financial markets," *Journal of political Economy*, 98(4), 703–738.

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.

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.

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.

Gelman, A., J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, 2020, *Bayesian Data Analysis*. third, electronic edition edn.

Geweke, J., 2001, "A note on some limitations of CRRA utility," *Economics letters*, 71(3), 341–345.

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.

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, 2016, "Learning from inflation experiences," *The Quarterly Journal of Economics*, 131(1), 53–87.

Malmendier, U., D. Pouzo, and V. Vanasco, 2020, "Investor experiences and financial market dynamics," *Journal of Financial Economics*, 136(3), 597–622.

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

Nagel, S., and Z. Xu, 2021, "Asset pricing with fading memory," *The Review of Financial Studies*.

——— , 2022, "Dynamics of Subjective Risk Premia," *NBER working paper*.

Panageas, S., 2020, "The implications of heterogeneity and inequality for asset pricing," *NBER working paper*.

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.

Weitzman, M. L., 2007, "Subjective expectations and asset-return puzzles," *American Economic Review*, 97(4), 1102–1130.

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.