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

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

August 25, 2023

*We thank Ella Dias Saraiva-Patelli, Ines Chaieb, Bernard Dumas, Philipp Illeditsch, Dongliang Lu, Albert Menkveld, Alessandro Melone, Martin Schneider, Raman Uppal, Youchang Wu, Goufu Zhou, and seminar participants at the Federal Reserve Bank of Atlanta, the University of Liechtenstein, CBS Copenhagen, SDU Odense, Tilburg University, WU Wien, the 16th Jackson Hole Finance Group Conference, the 2023 Annual Meeting of the Swiss Society for Financial Market Research and the Computational and Financial Econometrics conference in London.

†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 Portfolio Flows

Abstract

We study equilibrium asset prices and portfolio flows in a model where agents learn about economic fundamentals and differ in their aversion to parameter uncertainty. Exploiting the connection between confidence intervals from classical statistics and multi-prior sets for ambiguity-sensitive decision makers, we show that, because ambiguity-averse agents hold conservative portfolios, in equilibrium they are willing to accept a lower compensation to take on additional risk, making them natural *buyers* of risky assets when volatility rises. The model generates time-varying risk premia that are amplified by bad news and dampened by good news. We empirically document that these predictions are consistent with observed portfolio flows of agency and proprietary traders and with patterns of risk premia around information-sensitive events.

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 or macro announcements, frequently see institutional investors off-loading risky holdings in their portfolios with individual investors taking the opposite side of the trade. Common explanations for such flows include, for example, the role of informational asymmetry, portfolio constraints, or limited attention of retail investors.[1] Similarly, a considerable body of research has documented that a substantial portion of the equity risk premium is earned during information sensitive events, such as FOMC announcements.[2] Despite the abundance of literature on these topics, *jointly* explaining equilibrium dynamics of portfolio flows and risk premia in a unified framework remains an open challenge.

In this paper, we address this challenge by proposing an equilibrium asset pricing model in which agents learn about the endowment process from realized dividends, have uniform information but are heterogeneous in their aversion against ambiguity in parameter estimates. Following Bewley (2011), we interpret confidence intervals from classical statistics as sets of posterior distributions of a "Knightian" decision maker. Confidence intervals around estimated parameter values represent therefore the amount of Knightian uncertainty, or ambiguity, an agent faces. Following large unexpected dividend realizations, investors' variance estimates increase and, at the same time, confidence intervals widen. Hence, learning about parameter values generates time-varying ambiguity.

Our main result is to show that in equilibrium an increase in estimated volatility induces ambiguity-averse agents to *increase* risk in their portfolio. This seems counterintuitive, since ambiguity-averse agents hold conservative portfolios and an increase in ambiguity about expected future dividends should intuitively make them even more cautious. However, because ambiguity-averse agents hold conservative portfolios, they are willing to accept less compensation for taking on additional risk in spite of an increase in ambiguity.

An implication of this finding is that, in a model with ambiguity-sensitive agents, learning about the dividend volatility, not just the mean, has a first-order impact on portfolio flows and risk premia in equilibrium. It is typically believed that the variance parameters are relatively easy to learn and that learning about volatility has a negligible impact on portfolios and equilibrium returns (e.g., Collin-Dufresne et al., 2016b; Buss et al., 2021). While it is true that variances can be learned easily when agents have access to a continuous stream of information, the same might not be true when information arrives in "batches" such as during earnings or macroeconomic announcements. Furthermore, because in our model ambiguity-sensitive agents are not locally risk-

[1] See, among others, Frazzini and Lamont (2007), Barber and Odean (2008) and Hirshleifer et al. (2008) Kaniel et al. (2008), Kaniel et al. (2012).

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

neutral, learning about volatility becomes particularly relevant for the dynamics of equilibrium flows and risk premia. By emphasizing the importance 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".

We first illustrate the main intuition in a simple two-period heterogeneous agent model in which ambiguity about the mean dividend is naturally tied to the estimated dividend volatility and which we can fully solve analytically. We show that when both ambiguity-averse and ambiguity-neutral agents coexist, the equilibrium risky holdings of the ambiguity-averse agent, while more conservative than those of the ambiguity-neutral agent, are positively related to the level of dividend volatility. Therefore, an increase in estimated volatility implies an equilibrium flow of risky assets from ambiguity-neutral (sellers) to ambiguity-averse (buyer) agents. Because volatility directly affects equilibrium portfolio holdings, learning about volatility becomes a crucial driver of flows in a model with ambiguity averse agents. The key intuition from the simple model is that in an equilibrium where both ambiguity-averse and ambiguity-neutral agents coexist, an increase in estimated volatility leads to a new equilibrium risk premium that is "too high" (i.e., price too low) for the ambiguity-averse agents and "too low" (i.e., price too high) for ambiguity-neutral agents to justify the previously optimal portfolio weights. This difference in subjective valuations implies gains from trade in which ambiguity-averse agents increase their position in the risky asset when uncertainty increases.

We then generalize the model to an infinite-horizon setting, where overlapping generations of agents learn about the mean and the variance of the endowment process. Leaning about variance introduces an important technical challenge. In fact, when both the mean and the variance are not known, the predictive distribution of dividends is Student-t. Due to the fat tails of the Student-t distribution, expected utility is not well-defined in this case, see, e.g., Geweke (2001). We overcome this difficulty by imposing an a-priori restriction on the dividend variance, exploiting 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 distribution with thinner tails, that allows us to fully characterize the equilibrium with learning about both mean and variance. We show that learning about variance generates equilibrium returns that are left-skewed, even with an iid-normal endowment process.

Because in our model the true dividend mean and variance are time-invariant, agents eventually learn these parameters perfectly, thus making learning irrelevant in the limit. We provide a tractable way to achieve perpetual learning in an overlapping-generation model by assuming that "information leakage" occurs as generations overlap: new generations partially disregard 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 et al. (2016a), and Malmendier et al. (2020).

We empirically investigate our model predictions regarding portfolio flows and risk premia using a novel database of Euro STOXX 50 futures transactions on the Eurex, one of the largest futures and options markets in the world. We find that when uncertainty in the markets is high agency traders—acting on behalf of clients—sell and proprietary traders—acting on their own account—buy. This patter is robust and cannot be simply attributed to momentum or contrarian strategies. Instead, these findings are consistent with the predictions of our equilibrium model where agency traders are ambiguity neutral and proprietary traders are ambiguity averse.[3] Furthermore, we also show that, in line with our model, proprietary traders earn a risk premium for providing liquidity at the expense of agency traders in periods of market turmoil, e.g., Nagel (2012). This finding is consistent with Nagel and Xu (2022) who report that subjective risk premia increase with the subjective estimate of variance.

Our work relates to three strands of literature. First, we contribute to the literature that studies asset prices under parameter uncertainty and learning.[4] We show that even in a model with iid-normal dividends, time variation in the estimated variance has significant qualitative implications on the joint dynamics of equilibrium flows and asset prices that features complex patterns present in empirical data. When dealing with parameter uncertainty and learning, the vast majority of the asset pricing literature assumes that the mean of the endowment process is unknown, but its 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 meetings or corporate earning announcements, and agents cannot avoid the effort to learn about volatility.[5]

Second, we contribute to the literature on asset pricing with heterogeneous agents.[6] We differ from the work in this literature by considering learning and agents' ambiguity aversion. Chapman and Polkovnichenko (2009) study asset pricing in two-date economies with heterogeneous agents

[3]This dichotomy is inspired by the "competence hypothesis" of Heath and Tversky (1991), according to which decision makers are generally ambiguity-averse toward tasks for which they do not feel competent.

[4]Among others, key contributions are David (1997), Veronesi (1999), Pástor (2000), Barberis (2000), Xia (2001) Leippold et al. (2008), and Collin-Dufresne et al. (2016b). Pástor and Veronesi (2009) an extensive overview of learning in financial markets.

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

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

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.[7] Buss et al. (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 et al. (2021), heterogeneity in beliefs emerges endogenously in our model as a consequence of agents' different attitude towards ambiguity.

Third, our work is related to the large literature that studies asset prices and the trading behavior of agency and proprietary traders. Ample evidence indicates that proprietary traders act as liquidity providers who meet agency traders' demand for immediacy.[8] Consistent with this view, we document that agency traders tend to sell when volatility rises. Although proprietary traders might be less sophisticated (see, e.g., Menkveld and Saru, 2023), they face lower agency costs and less liquidity constraints than their agency counterparts. This advantage allows them to act as liquidity providers, especially during times of financial turmoil when liquidity is a scarce resource. These patterns of flows, together with the observed high level of risk premia are consistent with the findings of the demand-based asset pricing literature, e.g., Koijen and Yogo (2019) where price-inelastic proprietary investors buy from agency traders in periods of high uncertainty.

The rest of the paper proceeds as follows. In Section 2 we provide intuition in a simple equilibrium model. Section 3 presents an overlapping-generations model with 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 derives the predictive dividend distribution and expected utility when dividend variance is unknown; Appendix C provides technical details of Bayesian learning with unknown variance; Appendix D provides details of the numerical construction of the equilibrium; and Appendix E analyzes the implications of stochastic volatility for equilibrium portfolio flows.

[7]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 et al. (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 et al. (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.

[8]See, e.g., Nagel (2012) and Biais et al. (2016).

2 A two-period model

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

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

The dividend $\tilde{d}$ is normally distributed with unknown mean μ and known variance σ^2, $\tilde{d} \sim \mathcal{N}(\mu, \sigma^2)$. The assumption of known dividend variance will be relaxed in Section 3. We assume that agents enter the initial date having observed a history of dividend realizations from which they can calculate the time series average m and the associated standard error s with $s \propto \sigma$.[9]

Preferences. The economy is populated by two types of agents, $i = A, B$, 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 estimate of the dividend mean μ. Type-B agents are standard Bayesian (subjective expected utility) investors. They use m as their subjective dividend mean and account for its estimation error by inflating the variance σ^2 by the squared standard error s^2.[10] Therefore, the predictive distribution of the dividend $\tilde{d}$ for agent B is

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

In contrast, type-A agents are averse to uncertainty in the mean estimate. To model aversion to uncertainty we exploit the connection between classical confidence regions and "Knightian" uncertainty, or ambiguity, e.g., Bewley (2011). We assume that ambiguity is represented by "multiple priors" about the distribution of $\tilde{d}$ and that A-agents are averse to this ambiguity. Operationally, we characterize the set of priors as "confidence interval" around the mean estimate, m, whose size depends on the standard error s and the agents attitude towards ambiguity. Specifically, we characterize the ambiguity that A-agentes face as the confidence interval

$$\mathcal{P} \equiv [m - \kappa s, m + \kappa s], \tag{2}$$

[9] Assuming n_t dividend observations, these quantities are, respectively, $m = \frac{1}{n_t}\sum_{k=1}^{t} d_k$ and $s = \frac{\sigma}{\sqrt{n_t}}$.

[10] See, e.g., Section 2.5 of Gelman et al. (2020) for a proof of this result.

with $\kappa > 0$ a preference parameter that captures ambiguity aversion. When $\kappa = 0$, the set of priors $\mathcal{P}$ collapses to the singleton m, and A- and B-agents are identical. The parameter κ also has a classical statistical interpretation as a quantile of a distribution. Therefore, agents A face the following *set* of predictive distributions

$$\tilde{d} \sim^A \mathcal{N}\left(\tilde{\mu}^A, \sigma^2 + s^2\right), \quad \tilde{\mu}^A \in \mathcal{P}. \tag{3}$$

Optimal Portfolios. At the initial date, agents $i = A, B$ are initially endowed with wealth W^i and chooses a portfolio of θ^i units of the risky assets. The agents' wealth at the terminal date is

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

Agents B choose the portfolio θ^B to maximize their expected utility of terminal wealth, that is,

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

subject to the budget constraint (4).

Being averse to ambiguity, agents A choose portfolios by maximizing expected utility under the "worst-case scenario" from the set $\mathcal{P}$ in equation (2) as in Gilboa and Schmeidler (1989). That is, type-A agents solve the following problem[11]

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

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

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

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 (3) where the belief $\tilde{\mu}^A$ is selected to be either $\mu^A \equiv m - \kappa s$, if $\theta^A > 0$ or $\mu^A \equiv m + \kappa s$, if $\theta^A < 0$. When ambiguity averse agents do not

[11] For simplicity, in our analysis we rely on the "max-min" implementation of the Gilboa and Schmeidler (1989) model, as in Garlappi et al. (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, e.g., Anderson et al. (2000) and Hansen and Sargent (2008).

participate, i.e., $\theta^A = 0$, no distinct prior is selected. Optimal portfolios are, however, well-defined

$$\theta^B = \frac{m - p(1+r)}{\gamma\left(\sigma^2 + s^2\right)} \quad \text{and} \quad \theta^A = \begin{cases} \frac{m-\kappa s - p(1+r)}{\gamma(\sigma^2+s^2)} > 0 & \text{if} \quad m - p(1+r) \quad > \kappa s, \\ 0 & \text{if} \quad |m - p(1+r)| \quad < \kappa s. \\ \frac{m+\kappa s - p(1+r)}{\gamma(\sigma^2+s^2)} < 0 & \text{if} \quad m - p(1+r) \quad < -\kappa s \end{cases} \tag{8}$$

This implies that agent A's problem in equation (6) is equivalent to the problem of agent B but with a "distorted" belief about the expected dividend. While B agents use $\mathbb{E}^B(\tilde{d}) = \mu^B = m$, ambiguity averse agents A determine their demand under the prior that implies $\mathbb{E}^A(\tilde{d}) = \mu^A$ (or they do not participate, in which case no prior is selected). Figure 1 shows the optimal demands θ^i of both agents as a function of the risky asset's price p. Because of ambiguity aversion, agents A hold a more conservative portfolio than agents B, $|\theta^A| < |\theta^B|$. Moreover, ambiguity aversion may induce non-participation, that is, $\theta^A = 0$.

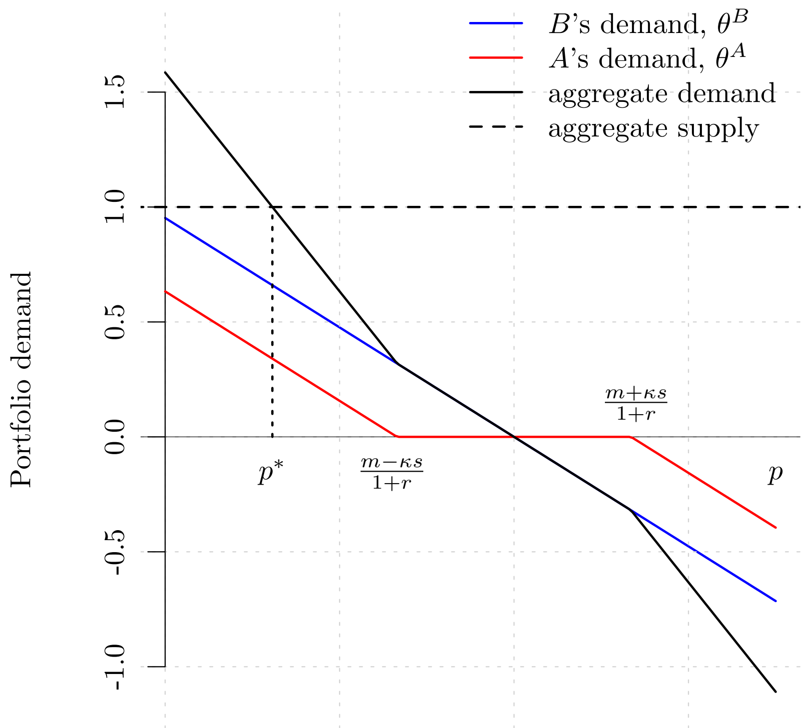

Figure 1: **Risky asset demand.** The figure shows the risky asset demand θ^B and θ^A from equation (8) as a function of the risky asset price p. The red line denotes type-A's demand; the blue line type-B's demand; the black line is aggregate demand; and the dashed line is the aggregate supply of the risky asset.

In equilibrium, when markets clear and $\theta^A + \theta^B = 1$, none of the agents hold short positions in the risky asset and either both types of agents choose positive portfolio weights or ambiguity averse agents do not participate. We obtain the following characterization of the equilibrium price p of the risky asset.

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

$$p = \frac{1}{1+r} m - \lambda, \tag{9}$$

where the risk premium λ is

$$\lambda = \begin{cases} \frac{1}{1+r}\left(\frac{\kappa}{2} s + \frac{\gamma}{2}\left(\sigma^2 + s^2\right)\right) & \text{if } \kappa \leq \kappa^*, \\ \frac{1}{1+r}\gamma\left(\sigma^2 + s^2\right)\sigma^2 & \text{if } \kappa > \kappa^*. \end{cases} \quad \text{with} \quad \kappa^* \equiv \gamma \frac{\sigma^2 + s^2}{s}. \tag{10}$$

Ambiguity averse agents participate, if their coefficient of ambiguity aversion is sufficiently low, $\kappa \leq \kappa^*$.

A proof of Proposition 1 is in Appendix A.

The demands for the risky asset in equation (8) implies that in equilibrium either both agents hold long positions or only B agents participate. Because the standard error is proportional to the dividend variance, $s \propto \sigma$, equation (A.4) shows that when both agents participate, i.e., $\kappa \leq \kappa^*$, the equilibrium risk premium is linear-quadratic in the dividend volatility σ. This is because the preferences of type-A agents exhibit "first-order" risk aversion, (see, e.g., Segal and Spivak, 1990). Intuitively, unlike B agents who are locally risk-neutral, A agents are locally risk-averse and demand a compensation for holding a vanishing amount of risk.

Substituting the equilibrium price p from equation (9) in the agents' demand functions (8) and simplifying we obtain that, when both agents participate, the equilibrium weights are

$$\theta^A = \frac{1}{2} - \frac{\kappa}{2\gamma} \underbrace{\left(\frac{s}{\sigma^2 + s^2}\right)}_{\propto \frac{1}{\sigma}}, \quad \text{and} \quad \theta^B = \frac{1}{2} + \frac{\kappa}{2\gamma} \underbrace{\left(\frac{s}{\sigma^2 + s^2}\right)}_{\propto \frac{1}{\sigma}}, \quad \kappa \leq \kappa^*. \tag{11}$$

Equation (11) shows that when both agents participate, ambiguity averse agents A increase their holdings of risky asset when volatility rises while Bayesian agents B decrease their holdings. As $\sigma \to \infty$, the portfolio holdings converge asymptotically to the constant weights $\theta^A = \theta^B = 1/2$.

Ignoring the connection between s and σ implied by the iid-normal model, the portfolio weight θ^A is negatively related to the standard error s and positively related to the variance σ^2. Hence, ambiguity-averse agents reduce exposure to the risky asset as ambiguity increase, i.e., as the confidence interval $\mathcal{P}$ in equation (2) widens, and increase exposure as variance increases.

Equation (11) shows that dividend volatility is a key variable for the determination of equilibrium portfolio weights. This feature is a general property of any model with heterogeneous asset demands. In our model, the difference in demand originates from heterogeneous aversion towards ambiguity. In general, differences in demand may emerge from a variety of reasons, such as, heterogeneous information, bounded rationality, differences in belief formation, etc. In fact, as long as these differences result in different dividend expectations across agents, $\mu^A < \mu^B$, in a CARA-normal setting, the equilibrium portfolio weights in equation (11) will take the following form

$$\theta^A = \frac{1}{2} - \frac{\Delta\mu}{\gamma(\sigma^2 + s^2)}, \quad \theta^B = \frac{1}{2} + \frac{\Delta\mu}{\gamma(\sigma^2 + s^2)}. \tag{12}$$

with $\Delta\mu \equiv \mu^B - \mu^A > 0$ denoting the difference in expectations. If $\Delta\mu$ is independent of σ, the pessimistic investors A hold a conservative portfolio and increase their risky holding following an increase in volatility. If $\Delta\mu$ depends on σ, agents A's risky holding increases (decreases) with σ depending on whether the "tilt", $\Delta\mu/(\sigma^2 + s^2)$ decreases (increases) in σ. In our model $\Delta\mu \propto \sigma$ and $s \propto \sigma$, therefore the tilt $\Delta\mu/(\sigma^2 + s^2) \propto 1/\sigma$ and the pessimistic investors A increase risky holdings as volatility increases. In general, however, the above analysis suggests that, as long as agents disagree on the mean dividend, volatility plays a key role in the determination of equilibrium holdings and flows.

Figure 2 provides an intuition for the structure of the equilibrium holdings in equation (11). 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 equation (8). Red-dashed lines refer to A-agents and blue-dashed lines refer to B-agents. The solid black line traces the intersection of complementary iso-portfolio curves, i.e., the set of volatility and risk premia (σ, λ) for which the market clears, $\theta^B + \theta^A = 1$. From equation (A.4), A-agents participate only when their ambiguity aversion $\kappa < \kappa^* \equiv \gamma\frac{\sigma^2+s^2}{s}$. If agents enter the initial date having observed a history of n_t dividend realizations, the standard error is $s = \sigma/\sqrt{n_t}$. Hence, the participation condition $\kappa < \kappa^* = \gamma\frac{n_t+1}{\sqrt{n_t}}\sigma$ can be equivalently expressed in terms of volatility as $\sigma > \sigma^* \equiv \frac{\sqrt{n_t}}{n_t+1}\frac{\kappa}{\gamma}$.

The red-shaded area in Figure 2 indicates (σ, λ) combinations for which A-agents do not participate. For values of $\sigma < \sigma^*$, the risk premium is too low for A-agents to participate. In this case, the equilibrium risk premium coincides therefore with the $\theta^B = 100\%$ iso-curve, i.e., the highest blue-dashed line. For values of $\sigma > \sigma^*$, both agents participate. Lemma A.1 in Appendix A shows that in any equilibrium in which A-agents participate, their iso-portfolio lines are always flatter than those of B-agents. Intuitively, because A-agents hold fewer units of the risky asset than B-agents, starting from an equilibrium in which both A and B participate, A-agents require relatively less compensation than B for bearing an additional unit of volatility while keeping the

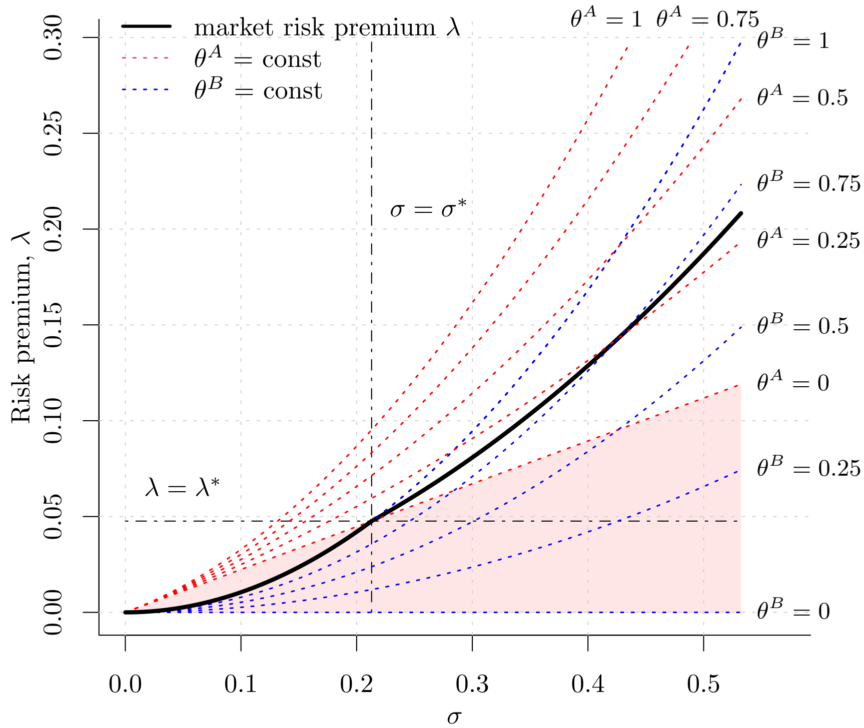

Figure 2: **Equilibrium portfolios and risk premia.** The figure shows iso-portfolios lines of type-A (red-dashed) and type B (blue-dashed) agents. These lines represent the set of volatility and risk premium values (σ, λ) that correspond to a given constant portfolio weight in equation (8), where the standard error $s = \sigma/\sqrt{n_t}$. The solid black line represents the set of points (σ, λ) at which the market clears, $\theta^A + \theta^B = 1$, and λ is the equilibrium risk premium λ. The vertical dashed-dotted line indicates the participation threshold $\kappa < \kappa^*$, or equivalently, $\sigma > \sigma^*$ with $\sigma^* \equiv \frac{\sqrt{n_t}}{n_t+1}\frac{\kappa}{\gamma}$, and the horizontal dashed-dotted line indicates the hurdle risk premium $\lambda^* \equiv \frac{\kappa}{(1+r)\sqrt{n_t}}\sigma^*$. Parameter values: $n_t = 20$, $\gamma = 1$, $\kappa = 1$.

portfolio unchanged. Hence, starting from any equilibrium with participation, a positive shock to volatility generates "gains from trade" where A-agents are willing to buy and B-agents are willing to sell.

In this section, we have assumed that the dividend variance is known. Therefore, in this model there cannot be equilibrium flows unless one is willing to assume that variance moves over time in an unexpected way so that agents are constantly surprised by shocks to volatility. We do not find such an assumption consistent with the forward-looking nature of market participants. A more realistic way to introduce time variation in volatility is to consider the case in which volatility is unknown and agents learn about it by observing dividend realizations. We develop such a model in the next section.

3 An overlapping-generations model

We consider an infinite-horizon overlapping-generations (OLG) model in which each generation consists of type-A and type-B agents in equal mass, as in Section 2, living for two periods. The setup we consider is similar to De Long et al. (1990) and Lewellen and Shanken (2002), however, unlike De Long et al. (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.

3.1 Setup

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 i.i.d. and normally distributed,

$$d_t \sim \mathcal{N}(\mu, \sigma^2), \tag{13}$$

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. They learn about μ and σ by observing dividend realizations over time.[12]

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 their assessment of expected end-of-period wealth: B-agents are Bayesian and A-agents are ambiguity averse.

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.

[12] To simplify the exposition of the model, we ignore the fact that eventually agents will learn the true parameters in the limit. In Section 3.4, we generalize the model to the case of perpetual learning.

3.2 Learning when variance is unknown

The unknown dividend variance poses a technical challenge in the definition of the agents' problem. Standard results from statistics (see, e.g. Greene, 2020), imply that when both the dividend mean and variance are unknown, the predictive distribution is Student-t. Hence, because learning about volatility generates fat-tailed dividend distributions, expected utility is not well defined (see, e.g. Geweke, 2001; Weitzman, 2007). We overcome this difficulty by adopting the approach proposed by Bakshi and Skoulakis (2010), who provide a methodology for solving asset pricing models with unknown volatility and Bayesian learning. To guarantee that expected utility is well-defined in such models, they propose to replace the standard normal-Gamma conjugate prior with a normal-*truncated* Gamma conjugate prior.[13] Their key methodological contribution is to show how to impose truncation bounds on the Gamma distribution to preserve conjugacy, that is, to ensure that the same bounds are preserved after agents update their priors. Bakshi and Skoulakis (2010) prove that the predictive distribution of dividends obtained in the normal-truncated Gamma setting is a "dampened" Student-t, that is a Student-t with thinner tails. Under this distribution, expected utility, and therefore the agents' portfolio choice problems, are well defined.

The learning problem of agents in each generation t can be decomposed into two steps: (i) information updating and (ii) belief formation. The first step is common to both A and B agents: they are equally informed and their heterogeneous preferences do not affect how they update their information in light of new dividend observations. The second step, belief formation, differs depending on agents' preferences: A and B make different use of their updated information from step (i) to form beliefs about moments of the predictive distribution of future dividends.

Figure 3 summarizes the two steps involved in agents' learning process. We now discuss each step in detail.

Information updating. While dividend variance is unknown, all generations agree that the variance is finite. Hence, the dividend precision $\phi \equiv 1/\sigma^2$ is bounded, $\phi \in [\underline{\phi}, \overline{\phi}]$ with $0 < \underline{\phi} < \overline{\phi} < \infty$. At each time t, both types of agents inherit information about μ and σ from generation $t-1$ in the form of a normal-inverse truncated Gamma posterior which they use as their t-prior, or "model", that gets updated after observing the new dividend, d_t. The generation $t-1$ posterior of ϕ that is handed over to generation t is a *truncated* Gamma with ν_{t-1} degrees of freedom and shape parameter b_{t-1}, that is,

$$\phi|(t-1) \sim^{A,B} \mathrm{TG}\left[\frac{\nu_{t-1}}{2}, \frac{b_{t-1}}{2}; \underline{\phi}, \overline{\phi}\right], \quad 0 < \underline{\phi} < \overline{\phi} < \infty, \tag{14}$$

[13] See, e.g., Gelman et al. (2020) for a discussion of Bayesian statistics in the normal-Gamma framework. The density of the *truncated* Gamma distribution is stated in equation (B.1) of Appendix B.

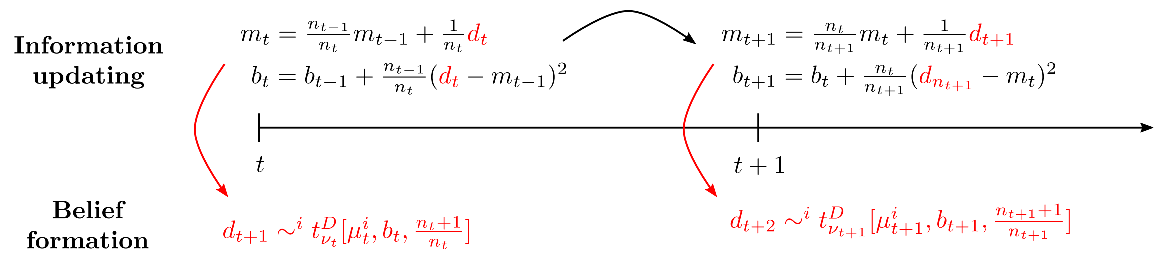

Figure 3: **Information updating and belief formation.** The top part of the figure illustrates the "information updating" step in which the state variables b_t and m_t are updated upon observing a new dividend realization d_{t+1}. The bottom part of the figure illustrates the "belief formation" step, 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. The predictive distribution $t^D_{\nu_t}[\cdot,\cdot,\cdot]$ refers to the dampened Student-t with ν_t degrees of freedom, see equation (24).

The degrees of freedom ν_{t-1} describes the precision of the ϕ-prior and is in classical statistics set to $n_t - 1$ to obtain a bias free estimate of the variance, see, e.g., Greene (2020). When we discuss learning under information leakage in Section 3.4 this close connection is no longer cawturetrue. Conditional on the precision ϕ, agents' prior about the mean μ is normally distributed, that is,

$$\mu|\phi,(t-1) \sim^{A,B} \mathcal{N}\left(m_{t-1}, \frac{1}{n_{t-1}\phi}\right), \tag{15}$$

as in the model of Section 2, with n_{t-1} the number of observations measuring the precision of the μ-prior. Again, with information leakage, the close connection between n_t and the number of observations is released. Hence, m_{t-1}, n_{t-1} and b_{t-1}, ν_{t-1} is the information that is passed on from generation $t-1$ to generation t, where m_{t-1} is the sample mean estimated from n_{t-1} observations and the shape parameter b_{t-1} is essentially the sum of historical squared errors and, hence, with ν_{t-1} degrees of freedom, b_{t-1}/ν_{t-1} captures an estimate of the variance. We can think of (m_{t-1}, b_{t-1}) as a "model of the world", agreed upon by all agents, with m_{t-1}, n_{t-1}, b_{t-1} and ν_{t-1} representing its state variables.

Generation t updates this information with the newly observed dividend d_t. The computed t-posterior is again of the normal-inverse truncated Gamma family

$$\phi|t \quad \sim^{A,B} \quad \mathrm{TG}\left[\frac{\nu_t}{2}, \frac{b_t}{2}; \underline{\phi}, \overline{\phi}\right], \quad 0 < \underline{\phi} < \overline{\phi} < \infty, \quad \nu_t = \nu_{t-1} + 1, \tag{16}$$

$$\mu|\phi, t \quad \sim^{A,B} \quad \mathcal{N}\left(m_t, \frac{1}{n_t\,\phi}\right), \quad n_t = n_{t-1} + 1, \tag{17}$$

where the state variables m_t and b_t are obtained from updating (m_{t-1}, b_{t-1}) upon observing the dividend d_t, that is,

$$m_t = \frac{n_{t-1}}{n_t} m_{t-1} + \frac{1}{n_t} d_t \tag{18}$$

$$b_t = b_{t-1} + \frac{n_{t-1}}{n_t} \underbrace{(d_t - m_{t-1})^2}_{\equiv e_t^2}, \tag{19}$$

with $e_t \equiv d_t - m_{t-1}$ denoting the time t dividend surprise relative to the mean m_{t-1}.

Belief formation. Although generation t-agents observe the same state variables, update information in the same way, and arrive at the same posteriors given in equations (16) and (17), they differ in how they use these posteriors to form their beliefs about the distribution of future dividends from which they derive their demand. Specifically, A-agents are averse against ambiguity in the expected dividend. To keep things simple, we assume that both agents are neutral to ambiguity in the variance. Assuming that A-agents are also averse towards ambiguity in the variance requires the specification of a set of priors for both μ and ϕ. This complicates the analysis without qualitatively changing the key features of the learning model [14]

Both agents use as a prior about the precision ϕ the posterior that results from information updating with d_t, i.e., the truncated Gamma distributed with shape parameter b_t from equations (16) and (19),

$$\phi \sim^{A,B} \text{TG}\left[\frac{\nu_t}{2}, \frac{b_t}{2}; \underline{\phi}, \overline{\phi}\right]. \tag{20}$$

B-agents are ambiguity neutral, hence, they use directly the time-t posterior of the information-updating step as their prior. Conditional on ϕ, their prior about the mean μ is normal with mean $\mu_t^B = m_t$,

$$\mu|\phi \sim^B \mathcal{N}\left(\mu_t^B, \frac{1}{n_t\,\phi}\right), \quad \mu_t^B = m_t. \tag{21}$$

In contrast, A-agents are uncertain about the model that drives the predictive distribution of future dividends. Unlike B-agents, who use (m_t, b_t) as their model, A-agents consider a set of models, obtained by constructing a confidence region centered around m_t, as this is done in the two-period model of Section 2. Specifically, they consider confidence intervals for μ from the t-posterior (17)

[14] Because the estimates of μ and σ are independent, confidence intervals are two-dimensional trapezoids. With max-min preferences, the ambiguity-averse agent will always elect the highest possible return variance when constructing optimal portfolios, see, e.g., Easley and O'Hara (2009). Therefore, the portfolio choice problem with ambiguity about both μ and ϕ reduces to a problem with ambiguity only about μ, where ϕ is fixed at the lowest possible value in the prior support.

and entertain the following conditionally normal priors of μ,

$$\mu|\phi \sim^A \mathcal{N}\left(\tilde{\mu}_t^A, \frac{1}{n_t\,\phi}\right), \quad \tilde{\mu}_t^A \in \mathcal{P}_t^m \tag{22}$$

where $\mathcal{P}_t^m$ denotes the ambiguity sets for the mean, μ,

$$\mathcal{P}_t \equiv [m_t - \kappa s_t, m_t + \kappa s_t]\,, \tag{23}$$

with $s_t = \frac{\kappa}{\sqrt{n_t \mathbb{E}_t[\phi|b_t]}}$ representing agent A's subjective estimate of the standard error of the dividend mean. Following Bakshi and Skoulakis (2010), in Lemma B.1 we show, under these assumptions, agents i's predictive distribution of future dividends is a "dampened" Student-t[15]

$$d_{t+1} \sim^A t_{\nu_t}^D\left[\tilde{\mu}_t^A, b_t, \frac{n_t+1}{n_t}\right], \quad d_{t+1} \sim^B t_{\nu_t}^D\left[\mu_t^B, b_t, \frac{n_t+1}{n_t}\right], \tag{24}$$

where $\mu_t^B = m_t$ and $\tilde{\mu}_t^A \in \mathcal{P}_t^m$. Because the dampened Student-t distribution has thinner tails than the Student-t, the agents' expected utility is well-defined, and we can solve for their optimal portfolios.

3.3 Equilibrium

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

$$W_{t+1}^i = W_t^i(1+r) + \theta_t^i\left(p_{t+1} + d_{t+1} - p_t(1+r)\right), \quad i = A, B. \tag{25}$$

Hence, to determine their portfolios at time t, agents have to form expectations about future dividends d_{t+1} and prices p_{t+1}. They choose their optimal portfolio by solving, respectively, the following maximization problems,

$$\max_{\theta_t^B} \mathbb{E}_t[u(W_{t+1}^B)], \tag{26}$$

and

$$\max_{\theta_t^A} \min_{\tilde{\mu}^A \in \mathcal{P}_t} \mathbb{E}_t[u(W_{t+1}^A)], \tag{27}$$

where $\mathcal{P}_t$ is the ambiguity set defined in equation (23). Proposition B.1 characterizes the time-t generation expected utility $\mathbb{E}_t^i[u(W_{t+1}^i)]$ for a given portfolio θ_t^i.

Unfortunately, there is no closed-form solution for agents' demand in this setting. However, we show in Appendix B.1 that the OLG model with unknown variance features an equilibrium price

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

of the form

$$p(m_t, b_t) = \frac{1}{r} m_t - \Lambda(b_t), \tag{28}$$

with the consequence that A-agents again select either the prior $\mu_t^A = m_t - \kappa s_t$ or they do not participate (as it is the case in the two-period model in Section 2). Hence, participating ambiguity-averse A-agents select portfolios as if they were purely risk averse with a distorted belief about the expected dividend, $\mathbb{E}_t^A(\tilde{d}_{t+1}) = \mu_t^A = m_t - \kappa s_t$, relative to the expectation of the Bayesian B-agents, who make portfolio decisions according to $\mathbb{E}_t^B(\tilde{d}_{t+1}) = m_t$.

As Proposition B.1 shows, the determination of (participating) agents' expect utility in each generation only requires the numerical evaluation of a well-behaved integral. Agent i's demand θ_t^i is obtained by solving the first-order conditions

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

Appendix D provides details of the numerical procedure we use to construct the equilibrium.

As a special case, we solve the case with known variance, for which we can obtain a closed-form solution. The following proposition characterizes the equilibrium price and the portfolio weights when variance is known.

Proposition 2. *Assume $d_t \sim \mathcal{N}(\mu, \sigma^2)$, with μ unobservable and σ observable and that agents $i = A, B$ form beliefs μ_t^i about μ as described in equations (21) and (22) with $\phi = 1/\sigma^2$ known. Then, the equilibrium price of the risky asset when both agents participate is*

$$p_t = \frac{1}{r} m_t - \Lambda_t, \tag{29}$$

where the risk premium Λ_t is given by

$$\Lambda_t = g_t \frac{\kappa}{2} \sigma + f_t \frac{\gamma}{2} \sigma^2, \tag{30}$$

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

$$\theta_t^A = \frac{1}{2} - \frac{\kappa}{2\gamma} \left(\frac{r\sqrt{n_t}}{1 + r(n_t + 1)} \right) \frac{1}{\sigma} \quad \text{and} \quad \theta_t^B = \frac{1}{2} + \frac{\kappa}{2\gamma} \left(\frac{r\sqrt{n_t}}{1 + r(n_t + 1)} \right) \frac{1}{\sigma}. \tag{31}$$

A proof of Proposition 2 is in Appendix A.

The equilibrium weights (31) are the infinite-horizon OLG equivalent of the equilibrium weights in equation (11) in the two-period model of Section 2. As in the simple model of Section 2, A-agents hold conservative portfolios, $\theta_t^A < \theta_t^B$, but as volatility increase, they increase the weight in the risky asset, i.e., $\partial\theta_t^A/\partial\sigma > 0$.[16]

Portfolio flows: As in the simple model of Section 2, after obtaining the equilibrium prices, we can derive the equilibrium portfolio weights of both agents. We denote by $\theta^i(m_t, b_t)$ the resulting equilibrium demand of type-i agents in generation t and define portfolio flows as

$$\Delta\theta^i(m_t, b_t) = \theta^i(m_t, b_t) - \theta^i(m_{t-1}, b_{t-1}), \quad i = A, B. \tag{32}$$

The equilibrium portfolio weights depend on the state variables m_t and b_t. Therefore, unlike the static model of Section 2 or the OLG model with known variance of Proposition 2, learning about the dividend variance generates flows across agents in equilibrium. Specifically, a positive flow, $\Delta\theta^i(m_t, b_t) > 0$, 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.

3.4 Perpetual learning

In the analysis so far, we have ignored the fact that, because the unknown true dividend mean μ and variance σ^2 are constant, agents eventually learn the true parameters. This problem is particularly relevant in the context of an OLG economy, where there is an implicit assumption that generations overlap forever. To address this common shortcoming of learning models, we modify the analysis of Section 3.2 by assuming that some 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), Collin-Dufresne et al. (2016a), and Ehling et al. (2018).

To model perpetual learning in a tractable fashion in an OLG model, we assume that when generation overlap, a shock occurs that reduces the informativeness of the posteriors of ϕ and $\mu|\phi$ in equations (14) and (15) as they are handed-over from one generation to the next. The effect of this shock is to introduce a "time-distortion" by a factor $\omega \in (0, 1)$.[17] In particular, with information leakage, the number of observations, n_t, and the degrees of freedom, ν_t, are no longer incremented by one with each new dividend observation, see equations (16) and (17). The iteration schemes are now $n_t = \omega n_{t-1} + 1$ and $\nu_t = \omega\nu_{t-1} + 1$, i.e., the information content in the $t-1$ prior is discounted before learning from the new observation is done. In the context of information leakage, we refer to

[16] As $t \to \infty$ agents eventually learn the true mean μ and the weights in equation (31) converge to 1/2. In Section 3.4 we extend the model to allow for perpetual learning.

[17] Appendix C 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.

n_t as the "effective number of observations" as it is used, e.g., in Nagel and Xu (2021). Intuitively, the quantity n_t represents a new "clock" that runs slower than t. This device allows us to achieve perpetual learning in our model while keeping most of the analysis of Section 3.2 unaffected.

In the case of perpetual learning, the prior of $\phi|t$ in equation (16) becomes

$$\phi \sim^{A,B} \mathrm{TG}\left[\frac{\nu_t}{2}, \frac{b_t}{2}; \underline{\phi}, \overline{\phi}\right], \quad 0 < \underline{\phi} < \overline{\phi} < \infty, \tag{33}$$

where the degrees of freedom are ν_t instead of $t-1$; the prior of $\mu|\phi, t$ in equation (17) becomes

$$\mu|\phi \sim^{A,B} \mathcal{N}\left(m_t, \frac{1}{n_t \phi}\right), \tag{34}$$

where the precision is $n_t\phi$ instead of $t\phi$; the updating equations of the state variables m_t and b_t in equations (18)–(19) become

$$m_{t+1} = \frac{\omega n_t}{\omega n_t + 1} m_t + \frac{1}{\omega n_t + 1} d_{t+1}, \tag{35}$$

$$b_{t+1} = \omega b_t + \frac{\omega n_t}{\omega n_t + 1}(d_{t+1} - m_t)^2; \tag{36}$$

and the updating equation of the additional two new state variables n_t and ν_t is

$$n_{t+1} = \omega n_t + 1, \quad \nu_{t+1} = \omega \nu_t + 1. \tag{37}$$

In these equations, the parameter $\omega \in [0, 1]$ controls the amount information leakage. By setting $\omega = 1$, we can recover the case described in Section 3.2. Notice that, as t increases, n_t approaches the asymptotic value $\overline{n} = \frac{1}{1-\omega}$. Therefore, in the steady state the problem can be described by just two state variables, m_t and b_t.

3.5 Results

We numerically solve for an equilibrium with unknown mean and variance with perpetual learning, following the procedure described in Appendix D. Figure 4 shows the equilibrium portfolio weights. The red line refers to the ambiguity averse agents' portfolios and the blue line to the Bayesian agents' portfolios. Consistent with the weights derived in equation (11) in the two-period model of Section 2 and with those for an OLG economy with known variance in Proposition 2, the weights of the ambiguity averse agent A is *increasing* and those of the Bayesian agent B are decreasing in $\sqrt{b_t/\nu_t}$, which is an indicator of the estimated dividend volatility.[18] In the figure we set the

[18] If dividend precision $\phi = 1/\sigma^2$ is non-truncated Gamma distributed, $\mathbb{E}(\phi) = \nu_t/b_t$. This motivates the usage of b_t/ν_t, i.e., the sum of squared errors divided by the number of degrees of freedom as an indicator of dividend variance and the square root of this expression as an indicator of dividend volatility.

truncation bounds $\underline{\phi}$ and $\overline{\phi}$ such that the true volatility $\sigma \in [0.2, 0.6]$. When the $\sqrt{b_t/\nu_t}$ 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 shaded areas in the figure show, the portfolio weights are less sensitive to changes in dividend volatility for values outside the range of $[0.1, 0.6]$.

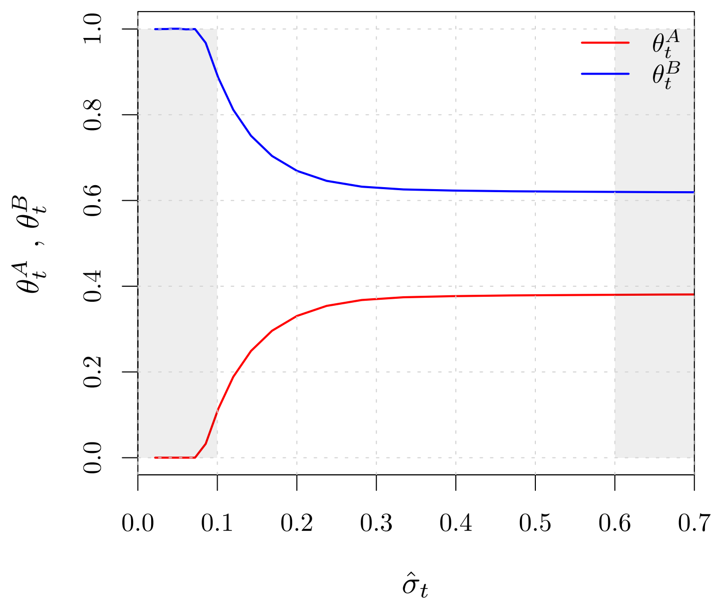

Figure 4: **Equilibrium portfolios and volatility.** The figure shows A and B agents' equilibrium portfolios as a function of the standard deviation estimate $\hat{\sigma}_t \equiv \sqrt{b_t/\nu_t}$, where b_t is the sum of squared error and ν_t the degrees of freedom. Parameters values: $\kappa = 1$, $\underline{\phi}$ and $\overline{\phi}$ such that the true volatility $\sigma \in [0.1, 0.6]$, $r = 0.1$, and $\nu_t = \overline{n} = 20$.

The dependence of the portfolio weights θ_t^A and θ_t^B on $\sqrt{b_t/\nu_t}$, shown in Figure 4, has a direct counterpart in terms of dividend "surprises", i.e., deviations of the realized dividend from the historical mean m_t. Figure 5 shows the equilibrium risky asset holdings of A-agents (left panel) and B-agents (right panel) as a function of this surprise. Different lines correspond 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 (B-agents). Large dividend surprises are associated with large subjective values of volatility. The U-shape nature of the equilibrium portfolios in the left panel of Figure 5 implies that A-agents are more aggressive than B-agents in their trades. After large positive and negative surprises A-agents *increase* their risky asset holdings. Heterogeneity in ambiguity attitude is crucial for this result. In fact, if A-agents were ambiguity neutral, $\kappa = 0$, there will be no flows in equilibrium, even if A and B were to differ in their degree of

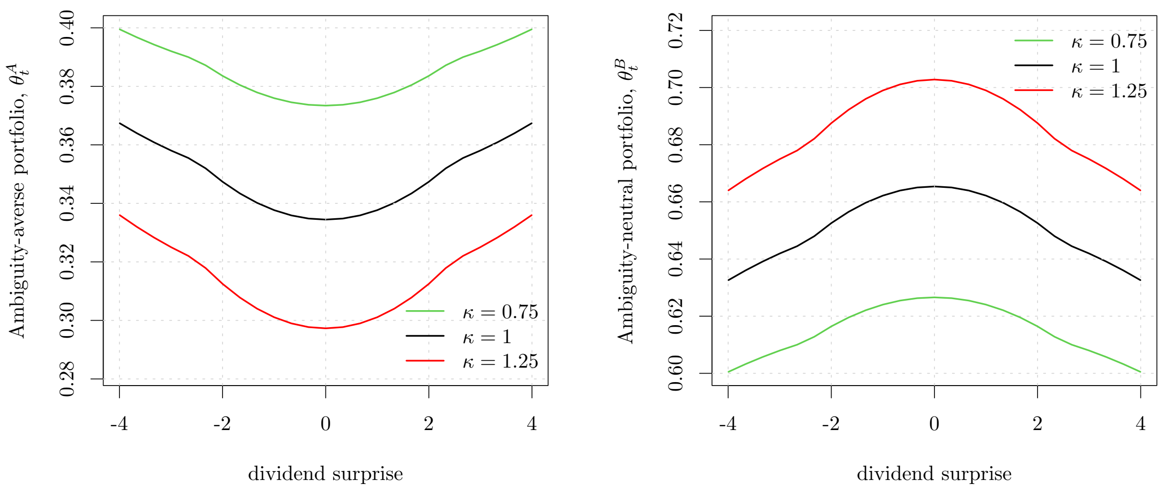

Figure 5: **Equilibrium portfolios and dividend surprise.** The figure shows the equilibrium portfolio of A-agents (left panel) and B-agents (right panel) as a function of the standard deviation estimate $\hat{\sigma}_t \equiv \sqrt{b_t/\nu_t} = 0.3$, where b_t is the sum of squared error and ν_t the degrees of freedom. Different lines corresponds to different values of the ambiguity aversion parameter κ. Parameters values: $\underline{\phi}$ and $\overline{\phi}$ such that $\sigma \in [0.1, 0.6]$, $r = 0.1$, and $\nu_t = \overline{n} = \frac{1}{1-\omega} = 20$.

risk aversion. As discussed earlier, heterogeneity in ambiguity attitude provides a micro-foundation for difference in beliefs across agents, which is at the root of the trading motive in our model.

Figure 6 shows equilibrium flows for a given random path of dividend realizations. The solid red line reports the time series of normalized surprises, while the blue bars represent flows of risky asset from B to A, defined in equation (32). A positive value of flows means that A is buying from B in equilibrium, and vice versa for negative values. Consistent with results shown in Figures 4 and 5, following large positive and negative surprises, ambiguity-averse A-agents increase their holding of the risky asset by buying from ambiguity-neutral B-agents, $\Delta\theta_t^A > 0$. In contrast, periods with low surprises are characterized by A-agents selling to B-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.

3.6 Implications for return predictability

As the expression of the equilibrium price in equation (28) shows, when variance is unknown the risk premium is time-varying and depends on the state variable b_t. In contrast, when variance is

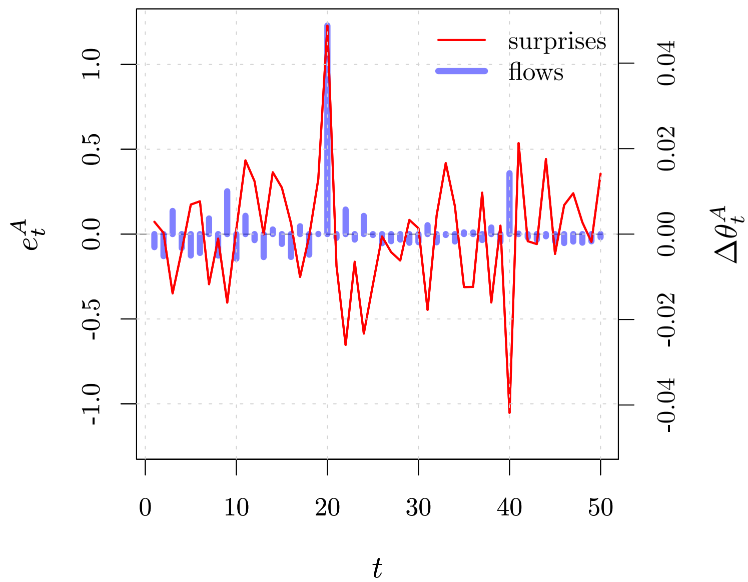

Figure 6: **Portfolio flows on a simulated path.** The figure illustrates A-agents's surprises, $e_t^A = d_{t+1} - \mu_t^A$ (left axis), and portfolio flows, $\Delta\theta_t^A$ (right axis). The state variable b_0 is chosen such that $\hat{\sigma}_0 = \sqrt{b_0/\overline{n}} = 0.3$. Parameters values: $\kappa = 1$, $\underline{\phi}$ and $\overline{\phi}$ such that the true volatility $\sigma \in [0.1, 0.6]$, $r = 0.1$, and $\nu_t = \overline{n} = \frac{1}{1-\omega} = 20$.

known, as Proposition 2 shows, the risk premium is a deterministic function of time. Therefore, learning about volatility *qualitatively* changes the nature of equilibrium asset prices and, by making risk-premia endogenously time-varying, generates return predictability in the model.

To explore the implications of our model for return predictability, Figure 7 shows the equilibrium risk premium as a function $\sqrt{b_t/\nu_t}$, an estimate of the standard deviation, in the steady state of the OLG economy with perpetual learning. The figure shows that the risk premium is an increasing function of the estimated standard deviation. Combined with the properties of portfolio flows shown in Figure 4, the risk premium pattern in Figure 7 implies that when B agents sell to A agents, a phenomenon that occurs following surprising dividend realizations, asset prices are lower. This feature is consistent with the findings from the demand-based asset pricing literature (e.g., Koijen and Yogo, 2019) where price-inelastic A-agents are willing to absorb flows from selling B-agents, whose demand shows higher price elasticity.

To understand 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 Proposition 2. This case is similar to the economy studied by Lewellen and

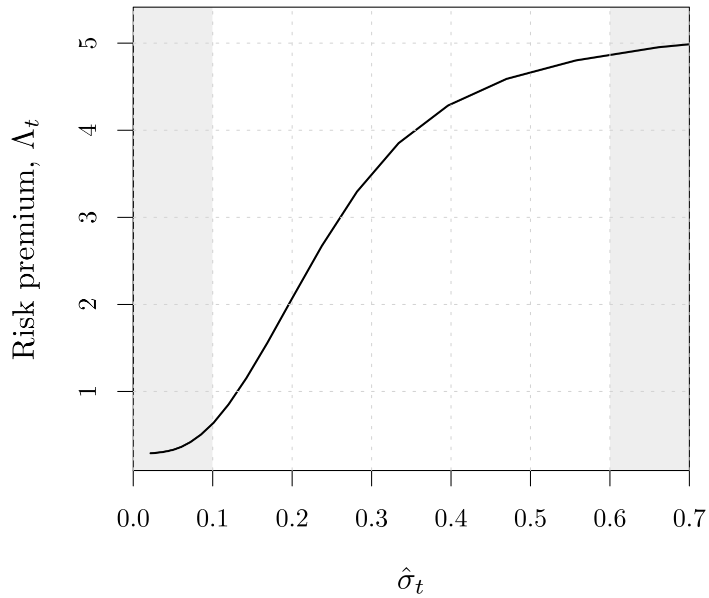

Figure 7: **Equilibrium risk premium.** The figure shows the equilibrium risk premium Λ_t as a function of the standard deviation estimate $\hat{\sigma}_t \equiv \sqrt{b_t/\nu_t}$ when both the dividend mean and variance are unknown. Parameters values: $\kappa = 1$, $\underline{\phi}$ and $\overline{\phi}$ such that the true volatility $\sigma \in [0.1, 0.6]$, $r = 0.1$, and $\nu_t = \overline{n} = \frac{1}{1-\omega} = 20$.

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 decompose it as follows

$$\lambda_t^{\text{obj}} = \mu - p_t r = \underbrace{\mu - \mu_t^i}_{\text{unobservable}} + \underbrace{\mu_t^i - p_t r}_{\equiv \lambda_t^i}, \tag{38}$$

with λ_t^i denoting i-agents' subjective risk premium. Using the equilibrium price p_t derived in Proposition 2, 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, which leads to an increase in the price p_t. As a consequence, see equation (38), 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, the equilibrium subjective risk premium is time-varying as it explicitly depends on the state variable b_t. Formally, from equation (28) we can define the objective and subjective risk premia, Λ_t^{obj} and Λ_t^i as follows

$$p_t(m_t, b_t) = \frac{1}{r} m_t - \Lambda(b_t) = \frac{1}{r}\mu_t^i - \underbrace{\left(\frac{1}{r}(\mu_t^i - m_t) + \Lambda(b_t)\right)}_{\equiv \Lambda_t^i(b_t)}, \tag{39}$$

$$= \frac{1}{r}\mu - \underbrace{\left(\frac{1}{r}(\mu - m_t) + \Lambda(b_t)\right)}_{\equiv \Lambda_t^{\text{obj}}(b_t)}. \tag{40}$$

Because for B agents, $\mu_t^i = m_t$, equation (39) implies that $\Lambda_t(b_t)$ is agents B's subjective risk premium. As Figure 7 shows, $\Lambda_t(b_t)$ is an increasing function of $\sqrt{b_t/\nu_t}$. Therefore, agents expect higher returns after large positive or negative dividend surprises.[19] Time-varying subjective risk premia implies return predictability.

Finally, note that the objective risk premium Λ_t^{obj} in equation (40) 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 realizations would detect risk premia that are amplified by bad news and dampened by good new. Figure 8 illustrates the dynamic of the risk premium $\Lambda_t(b_t)$ for a simulated random path of dividend surprises, $e_t^i = d_{t+1} - \mu_t^i$, $i = A, B$. 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.

4 Empirical analysis

In this section, we provide evidence in support of our model predictions using a novel database of trading activity on Euro STOXX 50 futures. Before presenting our main empirical results we provide a brief description of the data.

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

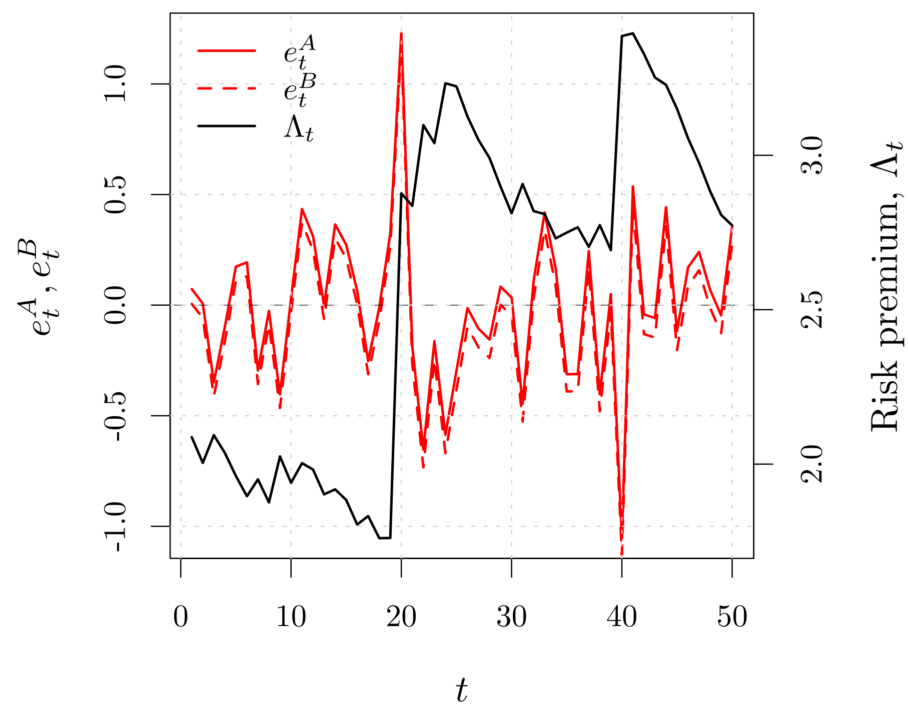

Figure 8: **Risk premium on a simulated path.** The figure illustrates surprises $e_t^i = d_{t+1} - \mu_t^i$, $i = A, B$ (left axis), with the corresponding steady state equilibrium risk premium Λ_t (right axis). The state variable b_0 is chosen such that $\hat{\sigma}_0 = \sqrt{b_0/\overline{n}} = 0.3$. Parameters values: $\kappa = 1$, $\underline{\phi}$ and $\overline{\phi}$ such that $\sigma \in [0.1, 0.6]$, $r = 0.1$, and $\nu_t = \overline{n} = \frac{1}{1-\omega} = 20$.

4.1 Data

Our main data source consists of Euro STOXX 50 futures transactions on the Eurex, one of the most active futures and options markets in the world.[20] The Euro STOXX 50 is the index for the largest and most liquid stocks in the Eurozone. The sample period spans from January 2002 to December 2020 and the data contain information on order flows of three different trader types: agency traders, market makers, and proprietary traders. Agency traders are market participants who trade for a client, while market makers and proprietary traders act on their own account. Trading takes place in an electronic limit order book, and trading flows are recorded at a frequency

[20] This data set was also used for, e.g., the Finance Crowd Analysis Project, see `https://fincap.academy` and Menkveld et al. (2023). See, e.g., Menkveld and Saru (2023) for further background information and institutional details. We thank Deutsche Boerse for providing us these data.

of milliseconds. In total, we have 824 million trades from the three different trader types at 188 million different timestamps. For the purpose of our analysis, we aggregate buy and sell orders at a daily frequency Flows of market makers and agency traders show a correlation coefficient of −0.8. Those of proprietary traders and agency traders are correlated with a coefficient of −0.86. Therefore, in line with the Finance Crowd Analysis Project, we analyze flows as the weight changes in agency traders' portfolio $\Delta\theta_t^{\text{agency}} = \theta_t^{\text{agency}} - \theta_{t-1}^{\text{agency}}$ and do not discriminate between flows originating from trades with proprietary traders and from trades with market makers.

Two challenges arise when bringing our 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 measures of dividend surprises. To address the first challenge, we rely on the microstructure literature that has modeled market makers as ambiguity-averse agents, arguing that ambiguity aversion may emerge as a natural response of market makers to inventory risk and to the fear of trading against informed traders, see, e.g., Routledge and Zin (2009), Easley and O'Hara (2010), and Zhou (2021).[21] This modelling choice is in line with experimental studies documenting that ambiguity aversion is influenced by the perceived competence of decision makers ("competence hypothesis", see Heath and Tversky, 1991), or "by a comparison with less ambiguous events or with more knowledgeable individuals" ("comparison hypothesis", see Fox and Tversky, 1995). Graham et al. (2009) argue that investors who perceive themselves competent are likely to have less parameter uncertainty about their subjective distribution of future asset returns. Menkveld and Saru (2023) document that for low frequency trades (greater than five seconds), agency traders are better informed than proprietary traders. To the extent that ambiguity aversion is a reaction to missing information, this evidence indicates that for low frequency trades proprietary traders have a stronger desire for robustness than agency traders. Hence, we use these arguments to identify agency traders as the least ambiguity averse agents, type-B in our model. Proprietary traders and market maker together constitute ambiguity averse agents, type-A in our model.

Because daily observations of dividends are not available, we address the second challenge by using the volatility index VSTOXX directly and interpret changes in VSTOXX as resulting from surprises about corporate profitability. The VSTOXX is the 30-day implied volatility calculated from options on the Euro STOXX 50 index. As such, this measure captures general market conditions and can be mapped to a measure of aggregate surprise in the market, consistent with the learning framework of our model.

[21]These papers study the 2007-2009 financial crisis and explain "market freezes" in markets for certain credit products with market makers' ambiguity aversion. While our model also features non-participation by ambiguity-averse agents, our focus is not on the micro-structure determinants of the bid-ask spread and on market freezes. Instead, our goal is to study the sensitivity of traders' asset flows to uncertainty shocks.

4.2 Equilibrium flows

Our model predicts that changes in the estimated volatility induces trades between market participants. To assess this channel empirically, at each day t we sort the aggregate daily flow data on the STOXX 50 index futures according to the change in the volatility index (ΔVSTOXX) from day $t-1$ to day t. Table 1 shows median values within quintile bins of volatility changes. Because of market clearing, we have $\Delta\theta^{\text{agency}} = -\left[\Delta\theta^{\text{proprietary}} + \Delta\theta^{\text{market maker}}\right]$.

Q	return	$\Delta\theta^{\text{agency}}$	VSTOXX	ΔVSTOXX
1	1.22	12019.00	24.75	-1.68
2	0.50	4623.00	19.18	-0.64
3	0.06	612.00	17.91	-0.09
4	-0.30	-4636.50	19.43	0.45
5	-1.38	-17224.00	25.11	1.71

Table 1: **Flows and volatility.** The table shows agency traders' flows $\Delta\theta^{\text{agency}}$ on the Euro Stoxx 50 futures, returns on the Euro Stoxx 50 and VSTOXX levels depended on changes in the volatility index ΔVSTOXX. The data is from Eurex Exchange, and flow data are aggregated on a daily frequency. Observations are grouped in quintile bins for ΔVSTOXX, and median values are reported for each bin.

The results in Table 1 are consistent with our model predictions regarding equilibrium flows. High level of surprises, captured by high levels of ΔVSTOXX, are associated with negative flows from agency to proprietary traders: Proprietary traders buy and agency traders sell in response to surprises in the market. As our model indicates, such a pattern would emerge when all traders learn and proprietary traders have a stronger desire for robustness than agency traders.[22] This evidence is also consistent with the literature that identifies proprietary traders as liquidity providers in times of market turmoil (see, e.g., Nagel, 2012).

Columns 2 and 3 in Table 1 show a strong and negative correlation between realized returns of the Euro Stoxx 50 and ΔVSTOXX. This empirical fact is known as the "asymmetric volatility phenomenon".[23] Given the highly negative relationship between changes in implied volatility and contemporaneous realized returns, one may argue that the flow patterns we observe are not driven by traders' reaction to surprises, as our model predicts, but by reaction to returns. For example, if agency traders follow "momentum" strategies and proprietary traders are "contrarian", we would observe negative agency traders' flows following high volatility simply because high volatility is

[22] These findings are not due to a specific chosen time frame and/or to a specific number of bins. A Kruskal-Wallis test shows that the median trading volumes of agency traders in the different bins are significantly different (p-value $<$ 2.2e-16), and a post-hoc Dunn test confirms the monotonic decline of aggregated daily agency trades with increasing volatility.

[23] For the US market, Dennis et al. (2006) estimate a negative correlation of -0.679 between returns and daily changes in implied volatility (see, e.g, also Wu, 2001; Bekaert and Wu, 2000).

associated with low returns. To isolate the effect of volatility and returns on flows, we analyze flow patterns *conditional* on realized returns. Table 2 shows the agency traders' flows at time t for different levels of volatility change (by row), conditional on realized returns at time $t-1$ (by column). The results in Table 2 show a very robust pattern across each column. In sum, even after conditioning for return levels, agency traders sell and proprietary traders buy when volatility increases. We take this evidence as support of the mechanism identified by our model where the more ambiguity averse (proprietary traders) buy from the less averse (agency traders) when uncertainty rises.

	lagged return				
ΔVSTOXX	1	2	3	4	5
1	9686	13148	13162	19416	11159
2	-123	5263	6949	6536	7499
3	-546	-450	990	-1178	4662
4	-9827	-6904	-3071	-3918	1143
5	-22138	-18283	-18279	-18439	-11196

Table 2: **Flows, volatility, and lagged returns.** The table shows agency traders' flows $\Delta\theta^{\text{agency}}$ on the Euro Stoxx 50 futures depended on changes in the volatility index ΔVSTOXX (rows) and lagged returns of the Euro Stoxx 50 (columns). The data is from Eurex Exchange, and flow data are aggregated on a daily frequency. Observations are grouped in quintile bins for ΔVSTOXX and lagged returns (double sort), and median values of $\Delta\theta^{\text{agency}}$ are reported for each combination.

4.3 Equilibrium risk premium

We conclude our empirical analysis by studying the relationship between equilibrium risk premia and trading activity. In our model, ambiguity-averse agents buy when uncertainty rises and ambiguity-neutral agency traders sell. Because risk premia increase with uncertainty, ambiguity-averse traders would on average earn the risk premium at the expense of ambiguity-neutral traders. To test this prediction, we regress realized returns on lagged flows of agents and on lagged variance (VSTOXX^2). We expect that higher variance at time t predicts higher $t+1$ returns (as a risk compensation), but that, on average, agency traders' flows at time t are associated with low realized returns at time $t+1$. The results reported column (1) of Table 3 confirm this conjecture. The coefficient of lagged agency traders' flows is negative and highly significant. The coefficient of VSTOXX^2 is positive and significant at the 5% level (Newey-West corrected).

Short-term return reversal is a natural alternative explanation for time-varying risk premia. To assess whether low time t returns predict high $t+1$ returns, we include lagged returns as a control variable. Column (2) in Table 3 shows that time t returns have no explanatory power in explaining $t+1$ returns. This result further emphasizes our model's prediction that time-varying variance is of first order relevance and that flows to ambiguity-averse investor groups predict a high future risk premia rather than historical returns do.[24]

	Dependent variable:	
	return	
	(1)	(2)
Constant	-3.165e-04	-3.126e-04
	(2.767e-04)	(2.762e-04)
lagged return		-2.615e-03
		(2.378e-02)
lagged $\Delta\theta^{\text{agency}}$	-1.837e-08**	-1.771e-08*
	(7.966e-09)	(1.027e-08)
lagged VSTOXX^2	8.397e-07*	8.350e-07*
	(4.712e-07)	(4.670e-07)
Observations	4,833	4,833
R^2	0.003	0.003
Adjusted R^2	0.003	0.003
Residual Std. Error	0.014 (df = 4830)	0.014 (df = 4829)
F Statistic	7.695*** (df = 2; 4830)	5.137*** (df = 3; 4829)
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 3: **Risk premia and portfolio flows.** The table shows results for Euro Stoxx 50 returns regressed on lagged Euro Stoxx 50 returns, lagged agents' flows $\Delta\theta^{\text{agendy}}$ on Euro Stoxx 50 futures and lagged variance VSTOXX^2. The data is from Eurex Exchange, and flow data are aggregated on a daily frequency. Standard errors are Newey-West corrected.

5 Conclusion

This study contributes to the literature on equilibrium asset prices and portfolio flows within an overlapping-generation model by incorporating agents' learning behavior regarding economic fundamentals and their differential aversion to parameter uncertainty. Our analysis reveals that in equilibrium, investors with ambiguity-averse preferences exhibit a propensity for more conservative portfolios while displaying heightened trading activity in response to unanticipated variations in

[24]This result complements the findings of Nagel (2012) that VIX is a much more powerful predictor of future returns than past returns are and of Nagel and Xu (2022) that subjective risk premia depend on investors' risk perception rather than following a simple anti-cyclical pattern.

corporate profitability. Notably, our model generates subjective risk premia that increases with uncertainty, alongside objective risk premia that are accentuated by adverse news and attenuated by favorable news.

The core findings of our study highlight the critical role played by two key factors in the determination of equilibrium portfolio flows and risk premia: distinctions in investors' ambiguity attitudes and the learning process pertaining to fundamental variance. Under scenarios where all agents exhibit ambiguity neutrality, equilibrium portfolio flows cease to exist. Moreover, in the context of a known variance setting, the equilibrium subjective risk premium assumes a constant value, leading to the persistence of static portfolios for agents. Unlike the case in which variance is known, we show that learning about variance generates return predictability that can be exploited by forward-looking investors.

To support our theoretical framework, we conduct an empirical analysis employing flow data on Euro Stoxx 50 futures contracts. Our empirical results provide evidence indicating that agency traders sell, and proprietary traders buy under increasing uncertainty. Additionally, we establish that the expected risk premium increases with uncertainty. These empirical findings provide support for the theoretical predictions of our model, which posits that ambiguity-neutral market participants engage in trading activities with ambiguity-averse, price-inelastic investors. Our research underscores the importance of accounting for the effects of ambiguity aversion and learning when jointly studying equilibrium asset prices and portfolio flows.

A Proofs

Proof of Proposition 1

The portfolio problem of agents can be written as follows

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

where expectations are taken with respect to agents prior distributions of $\tilde{d}$ from (1) and (3) plus the prior selection criterion in equation (7). Using the normality of $\tilde{d}$, the optimal portfolio weights are:

$$\theta^A = \begin{cases} 0 & \text{if } \mu^A = \mathcal{P} \\ \frac{\mu^A - p(1+r)}{\gamma(\sigma^2+s^2)} & \text{otherwise.} \end{cases}, \quad \theta^B = \frac{\mu^B - p(1+r)}{\gamma\left(\sigma^2+s^2\right)}, \tag{A.2}$$

Imposing market clearing, we obtain that in equilibrium $\theta^A \geq 0$ and furthermore

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

where

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

■

Lemma A.1. *Let $\bar{\lambda}^A(\sigma)$ and $\bar{\lambda}^B(\sigma)$ denote the iso-portfolios of agent A and B, respectively. For all equilibrium values λ of the risk premium in equation (A.4), we have that $\partial\bar{\lambda}^A(\sigma)/\partial\sigma < \partial\bar{\lambda}^B/\partial\sigma$.*

Proof. Let $s = \sigma/\sqrt{n_t}$, with n_t denoting the number of observations used to compute the dividend mean m and its standard error s. From agents' demand for the risky asset stated in equation (8) and the definition of the risk premium $\lambda = \frac{1}{1+r}m - p$, we derive the risk premium that agents require for holding a fraction θ^i of the risky asset (the iso-portfolio line) and its derivative with respect to the dividend volatility σ as

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

$$\frac{\partial\bar{\lambda}^A}{\partial\sigma} = \frac{\kappa}{\sqrt{n_t}} + 2\gamma\theta^A\left(\frac{n_t+1}{n_t}\right)\sigma, \qquad \frac{\partial\bar{\lambda}^B}{\partial\sigma} = 2\gamma\theta^B\left(\frac{n_t+1}{n_t}\right)\sigma. \tag{A.6}$$

We prove that along the equilibrium risk premium λ in equation (A.4) the slope of $\bar{\lambda}^A$ is flatter than the slope of $\bar{\lambda}^B$, i.e.,

$$\frac{\partial \bar{\lambda}^A}{\partial \sigma} \quad < \quad \frac{\partial \bar{\lambda}^B}{\partial \sigma}. \tag{A.7}$$

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

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

or, rearranging,

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

We restrict our analysis to the region where both agents are in the market, $\sigma > \frac{\sqrt{n_t}}{n_t+1}\frac{\kappa}{\gamma}$, and substitute equilibrium portfolios weights from equation (11) into the above inequality. This yields

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

$$2\frac{\kappa}{\sqrt{n_t}} \quad > \quad \frac{\kappa}{\sqrt{n_t}}, \tag{A.11}$$

which is true for $\kappa > 0$ and $n_t < \infty$ 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. As we see in what follows, the equilibrium price is linear in m_t independent of τ. When ambiguity-averse agents participate, they will in equilibrium take long positions in the risky asset and the selected prior from $\mathcal{P}_t$ which satisfies the max-min criterion (27) is always $\mu_t^A = m_t - \kappa s_t$. Hence, participating A-agents form portfolios according to the belief $\mu_t^A = m_t - \kappa s_t$ while Bayesian B-agents use the belief $\mu_t^B = m_t$. For ease of notation we use $\mathbb{E}_t^i$ to denote agent i's conditional expectations at time t. The risky asset demand $\theta_{t,\tau}^i$, $i = A, B$ is

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

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

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

where

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

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

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

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

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

using the predictive distribution (A.12) we obtain

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

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

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

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

$$g_{t,2} = \frac{1}{R}\left(1+\frac{1}{R(n_t+1)}\right)\frac{1}{\sqrt{n_t}} + \frac{1}{R^2}\frac{1}{\sqrt{n_t+1}},$$

$$f_{t,2} = \frac{1}{R}\left(1+\frac{1}{R(n_t+1)}\right)^2\left(\frac{n_t+1}{n_t}\right) + \frac{1}{R^2}\left(\frac{n_t+2}{n_t+1}\right).$$

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

$$g_{t,\tau} = \sum_{j=1}^{\tau}\frac{1}{R^{\tau+1-j}}\left(1+\frac{1}{n_t+\tau-j+1}\sum_{i=1}^{j-1}\frac{1}{R^i}\right)\frac{1}{\sqrt{n_t+\tau-j}},$$

$$f_{t,\tau} = \sum_{j=1}^{\tau}\frac{1}{R^{\tau+1-j}}\left(1+\frac{1}{n_t+\tau-j+1}\sum_{i=1}^{j-1}\frac{1}{R^i}\right)^2\left(\frac{n_t+\tau-j+1}{n_t+\tau-j}\right).$$

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

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

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

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

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

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

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

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

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

and

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

∎

B Predictive distribution of the dividend when the variance is unknonw

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

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

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

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

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,*

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

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.2).

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* $\frac{b}{\nu}\overline{\phi}$,

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

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

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

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

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

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

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

∎

B.1 Expected utility when variance is unknown

In the main text we argue that the equilibrium price in an infinite horizon OLG with unknown variance is linear in m_t and, hence, when both types of generation-t agents participate, their priors about μ are characterized by $\mu_t^A = m_t - \kappa s_t$ and $\mu_t^B = m_t$. In this section of the appendix we show that the model actually features an equilibrium price linear in m_t. The selected A-prior is consequently at the lower bound of the confidence interval $\mathcal{P}_t$.

In Lemma B.1 above, we show that the distribution of d_{t+1} follows a non-standardized dampened Student-t with state variables m_t, b_t and n_t. The expected utility of agents of type i with μ prior μ_t^i is then obtained by integration of the agents' CARA utility of $t+1$ wealth over the dampened t-density of the dividend

$$
\mathbb{E}_t^i[u(W_{t+1}^i)] = \int_{-\infty}^{\infty} \mathbb{E}_t^i[u(W_{t+1}^i)|\phi] p(\phi|b_t, \nu_t) d\phi, \tag{B.3}
$$

with $p(\phi|b_t, \nu_t)$ the density of the truncated Gamma distribution stated in equation (B.1). The expected utility conditional on ϕ is calculated by integrating agents CARA utility over the conditionally normal density of the μ prior μ_t^i

$$\mathbb{E}_t^i[u(W_{t+1}^i)|\phi] = -\frac{1}{\gamma}e^{-\gamma(1+r)(W_t-\theta p_t)}\sqrt{\frac{\phi}{2\left(\frac{n_t+1}{n_t}\right)\pi}}\int_{-\infty}^{\infty} e^{-\gamma\theta_t^i\left(\mu_t^i+e_{t+1}^i+p_{t+1}\right)-\frac{1}{2}\frac{\phi}{\left(\frac{n_t+1}{n_t}\right)}(e_{t+1}^i)^2} de_{t+1}^i. \tag{B.4}$$

with $e_{i+1}^i = d_{t+1} - \mu_t^i$ denotes agents i's dividend surprise.

Proposition B.1. *Suppose the equilibrium price p_t is linear in m_t with*

$$p_t = \frac{1}{r}m_t - \lambda(b_t, n_t, \nu_t),$$

and let $\Delta\mu_t^i = m_t - \mu_t^i$. Then expected utility of agents i is given by

$$\mathbb{E}_t^i(u(W_{t+1})) = \frac{1}{C(b_t,\nu_t;\underline{\phi},\overline{\phi})}\int_{\underline{\phi}}^{\overline{\phi}} \mathbb{E}_t^i(u(W_{t+1})|\phi,m_t,n_t,\nu_t)\phi^{\frac{\nu_t}{2}-1}e^{-\phi\frac{b_t}{2}}d\phi \tag{B.5}$$

$$\begin{aligned}
\mathbb{E}_t^i(u(W_{t+1})|\phi,m_t,n_t,\nu_t) &= -\frac{1}{\gamma}\sqrt{\frac{n_t\phi}{(n_t+1)2\pi}}\int_{-\infty}^{\infty} e^{-\gamma W_{t+1}^i}e^{-\frac{1}{2}\frac{n_t\phi}{(n_t+1)}(d_{t+1}-\mu_t^i)^2}dd_{t+1} \\
&= -\frac{1}{\gamma}e^{-\gamma(1+r)(W_t^i-\theta p_t)}\sqrt{\frac{n_t\phi}{(n_t+1)2\pi}}\int_{-\infty}^{\infty} e^{-\gamma\theta(d_{t+1}+p_{t+1})}e^{-\frac{1}{2}\frac{n_t\phi}{(n_t+1)}(d_{t+1}-\mu_t^i)^2}dd_{t+1} \\
&= -\frac{1}{\gamma}\exp\left\{-\gamma\left[(1+r)(W_t+\theta\Lambda(b_t,n_t,\nu_t))-\theta\Big(1+\frac{1}{rn_{t+1}}\Big)\Delta\mu_t^i\right]\right\} \\
&\times\sqrt{\frac{n_t\phi}{(n_t+1)2\pi}}\int_{-\infty}^{\infty}\exp\left\{-\gamma\theta\left[\left(1+\frac{1}{rn_{t+1}}\right)e_{t+1}^i-\Lambda(b_{t+1},n_{t+1},\nu_{t+1})\right]\right\} \\
&\times\exp\left\{-\frac{1}{2}\frac{n_t\phi}{(n_t+1)}(e_{t+1}^i)^2\right\}de_{t+1}^i \\
e_{t+1}^i &= d_{t+1}-\mu_t^i \\
b_{t+1} &= b_t+\frac{n_t}{n_{t+1}}(\Delta\mu_t^i-e_{t+1}^i)^2 \\
n_{t+1} &= n_t+1 \\
\nu_{t+1} &= \nu_t+1
\end{aligned}$$

Expected utility is independent of m_t. In particular, for ambiguity averse agents A, the largest possible value for $\Delta\mu_t^A$ minimizes expected utility, thus, the selected prior is $\mu_t^A = m_t - \kappa s_t$.

Proof: Substitute the respective densities, the linear price function and the budget constraint (25) in the agents' CARA utility and simplify the integral. Regarding $(1+r)\frac{1}{r}m_t = \frac{1}{r}m_t$, the contribution of m_t to p_t and its contribution to p_{t+1} cancel.

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 (B.5) is well defined.

C Information leakage and learning about mean and variance: Technical details

In this section of the appendix we describe how information leakage is modeled, how generation $t-1$ posteriors are first affected by leakage and how they are then updated with the observed dividend d_t. Throughout the appendix we make use of basic principles of Bayesian data analysis, see e.g., the textbook of Gelman et al. (2020).

The information set of generation $t-1$ about the unknown dividend mean μ and the precision ϕ is given by their posteriors in equations (14) and (15). When these posteriors are handed over to generation t, information is lost. We model information leakage in form of shocks that add noise to these posteriors before they are updated by generation t upon observation of the dividend d_t.

The generation $t-1$ posterior of μ with n_{t-1} (effective) observations is

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

This posterior is shocked with an additive Gaussian shock with mean 0 and variance $(\frac{1}{\omega}-1)\frac{1}{n_{t-1}\phi}$ with $\omega \in (0,1]$. This shock is independent of the estimation error in μ and increases the variance of the posterior by a factor $\frac{1}{\omega} \geq 1$. The post leakage μ-prior, which is passed to generation t is then

$$\mu^+|\phi, m_{t-1}, n_{t-1} \sim \mathcal{N}\left(m_{t-1}, \frac{1}{\omega}\frac{1}{n_{t-1}\phi}\right), \quad \omega \in (0,1].$$

This noisy posterior is then updated with the information contained in the dividend d_t. The following posterior is transferred to both types of agents of generation t in form of the updated state variables m_t and n_t

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

where

$$
\begin{aligned}
\mu|d_t,\phi,m_{t-1},n_{t-1} &\sim \mathcal{N}\left(m_t,\frac{1}{n_t\phi}\right), \\
e_t &= d_t-m_{t-1}, \\
m_t &= m_{t-1}+\frac{1}{n_t}e_t, \\
n_t &= \omega n_{t-1}+1.
\end{aligned}
$$

To update also the ϕ prior with the t dividend, we must first determine the distribution of d_t conditional of ϕ.

$$
\begin{aligned}
f(d_t|\phi,m_{t-1},n_{t-1}) &= \int_{-\infty}^{+\infty} f(d_t|\mu,\phi)p(\mu^+|\phi,m_{t-1},n_{t-1})d\mu \\
&= \int_{-\infty}^{+\infty}\sqrt{\frac{\phi}{2\pi}}e^{-\frac{1}{2}\phi(d_t-\mu)^2}\sqrt{\frac{\omega n_{t-1}\phi}{2\pi}}e^{-\frac{1}{2}\omega n_{t-1}\phi(\mu-m_{t-1})^2}d\mu \\
&= \sqrt{\frac{\omega n_{t-1}\phi}{(\omega n_{t-1}+1)2\pi}}e^{-\frac{1}{2}\frac{\omega n_{t-1}\phi}{(\omega n_{t-1}+1)}(d_t-m_{t-1})^2}\int_{-\infty}^{+\infty}\sqrt{\frac{(\omega n_{t-1}+1)\phi}{2\pi}}e^{-\frac{1}{2}(\omega n_{t-1}+1)\phi\left(\mu-\frac{d_t-\omega n_{t-1}m_{t-1}}{\omega n_{t-1}+1}\right)^2}d\mu \\
&= \sqrt{\frac{\omega n_{t-1}\phi}{(\omega n_{t-1}+1)2\pi}}e^{-\frac{1}{2}\frac{\omega n_{t-1}\phi}{(\omega n_{t-1}+1)}(d_t-m_{t-1})^2},
\end{aligned}
$$

where $d_t|\phi,m_{t-1},n_{t-1}\sim\mathcal{N}\left(m_{t-1},\frac{\omega n_{t-1}+1}{\omega n_{t-1}\phi}\right)$.

Information leakage regarding the posterior of the precision ϕ of generation $t-1$ is modeled as 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,

$$\phi^+|b_{t-1},\nu_{t-1} \sim \text{TG}\left[\omega\frac{\nu_{t-1}}{2},\omega\frac{b_{t-1}}{2};\underline{\phi},\overline{\phi}\right], \quad 0<\underline{\phi}<\overline{\phi}<\infty.$$

Updating with d_t leads to the posterior which is transferred to the agents of generation t in the form of b_t and ν_t.

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

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

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

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

D 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. 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

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

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

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

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_{t+1}^i

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

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

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

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

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

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

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

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, $\theta^B(b_t)+\theta^A(b_t)=1$.

E 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 the model of Section 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)

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

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.[25]

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_{t+1}^2, are

$$m_{t+1} = (1 - w_{t+1})m_t + w_{t+1}d_{t+1}, \quad w_{t+1} = \frac{\frac{1}{\sigma_{t+1}^2}}{\frac{1}{s_t^2} + \frac{1}{\sigma_{t+1}^2}} = \frac{s_t^2}{s_t^2 + \sigma_{t+1}^2} \tag{E.2}$$

$$\frac{1}{s_{t+1}^2} = \frac{1}{s_t^2} + \frac{1}{\sigma_{t+1}^2}. \tag{E.3}$$

Equation (E.3) 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 $\mathcal{P}_t^\mu = [m - \kappa s_t, m + \kappa s_t]$ in equation (E.3), 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}.

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

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

References

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

Anderson, E., Hansen, L. P., and Sargent, T. (2000). Robustness, detection and the price of risk. Technical report, mimeo.

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

Barber, B. M. and Odean, T. (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.

Bekaert, G. and Wu, G. (2000). Asymmetric volatility and risk in equity markets. *The Review of Financial Studies*, 13(1):1–42.

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

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

Biais, B., Declerck, F., and Moinas, S. (2016). Who supplies liquidity, how and when? *BIS Working Papers*, 563.

Buss, A., Uppal, R., and Vilkov, G. (2021). Dynamics of asset demands with confidence heterogeneity. working paper.

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

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

Collin-Dufresne, P., Johannes, M., and Lochstoer, L. A. (2016a). Asset pricing when 'This time is different'. *The Review of Financial Studies*, 30(2):505–535.

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

Constantinides, G. M. and Duffie, D. (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., Shleifer, A., Summers, L. H., and Waldmann, R. J. (1990). Noise trader risk in financial markets. *Journal of political Economy*, 98(4):703–738.

Dennis, P., Mayhew, S., and Stivers, C. (2006). Stock returns, implied volatility innovations, and the asymmetric volatility phenomenon. *Journal of Financial and Quantitative Analysis*, 41(2):381–406.

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

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

Easley, D. and O'Hara, M. (2009). Ambiguity and nonparticipation: The role of regulation. *The Review of Financial Studies*, 22(5):1817–1843.

Easley, D. and O'Hara, M. (2010). Liquidity and valuation in an uncertain world. *Journal of Financial Economics*, 97(1):1–11.

Ehling, P., Graniero, A., and Heyerdahl-Larsen, C. (2018). Asset prices and portfolio choice with learning from experience. *The Review of Economic Studies*, 85(3):1752–1780.

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

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

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

Gârleanu, N. and Panageas, S. (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., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2020). *Bayesian Data Analysis*. third, electronic edition edition.

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

Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility theory with non-unique prior. *Journal of Mathematical Economics*, 18:141—153.

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

Greene, W. H. (2020). *Econometric Analysis*. Pearson, 8 edition.

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

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

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

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

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

Ilut, C. L. and Schneider, M. (2022). Modeling uncertainty as ambiguity. *NBER working paper.*

Kaniel, R., Liu, S., Saar, G., and Titman, S. (2012). Individual investor trading and return patterns around earnings announcements. *The Journal of Finance*, 67(2):639–680.

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

Koijen, R. S. and Yogo, M. (2019). A demand system approach to asset pricing. *Journal of Political Economy*, 127(4):1475–1515.

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

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

Malmendier, U. and Nagel, S. (2016). Learning from inflation experiences. *The Quarterly Journal of Economics*, 131(1):53–87.

Malmendier, U., Pouzo, D., and Vanasco, V. (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.

Menkveld, A. J., Dreber, A., Holzmeister, F., Huber, J., Johannesson, M., Kirchler, M., Razen, M., Weitzel, U., Abad, D., Abudy, M. M., Adrian, T., Ait-Sahalia, Y., Akmansoy, O., Alcock, J., Alexeev, V., Aloosh, A., Amato, L., Amaya, D., Angel, J., Bach, A., Baidoo, E., Bakalli, G., Barbon, A., Bashchenko, O., Bindra, P. C., Bjonnes, G. H., Black, J., Black, B. S., Bohorquez, S., Bondarenko, O., Bos, C. S., Bosch-Rosa, C., Bouri, E., Brownlees, C. T., Calamia, A., Cao, V. N., Capelle-Blancard, G., Capera, L., Caporin, M., Carrion, A., Caskurlu, T., Chakrabarty, B., Chernov, M., Cheung, W. M., Chincarini, L. B., Chordia, T., Chow, S. C., Clapham, B., Colliard, J.-E., Comerton-Forde, C., Curran, E., Dao, T., Dare, W., Davies, R. J., De Blasis, R., De Nard, G., Declerck, F., Deev, O., Degryse, H., Deku, S., Desagre, C., van Dijk, M. A., Dim, C., Dimpfl, T., Dong, Y., Drummond, P., Dudda, T. L., Dumitrescu, A., Dyakov, T., Dyhrberg, A. H., Dzieliński, M., Eksi, A., El Kalak, I., ter Ellen, S., Eugster, N., Evans, M. D. D., Farrell, M., Félez-Viñas, E., Ferrara, G., Ferrouhi, E. M., Flori, A., Fluharty-Jaidee, J., Foley, S., Fong, K. Y. L., Foucault, T., Franus, T., Franzoni, F. A., Frijns, B., Frömmel, M., Fu, S., Füllbrunn, S., Gan, B., Gehrig, T., Gerritsen, D., Gil-Bazo, J., Glosten, L. R., Gomez, T., Gorbenko, A., Güçbilmez, U., Grammig, J., Gregoire, V., Hagströmer, B., Hambuckers, J., Hapnes, E., Harris, J. H., Harris, L., Hartmann, S., Hasse, J.-B., Hautsch, N., He, X., Heath, D., Hediger, S., Hendershott, T., Hibbert, A. M., Hjalmarsson, E., Hoelscher, S. A., Hoffmann, P., Holden, C. W., Horenstein, A. R., Huang, W., Huang, D., Hurlin, C., Ivashchenko, A., Iyer, S. R., Jahanshahloo, H., Jalkh, N., Jones, C. M., Jurkatis, S., Jylha, P., Kaeck, A., Kaiser, G., Karam, A., Karmaziene, E., Kassner, B., Kaustia, M., Kazak, E., Kearney, F., van Kervel, V., Khan, S., Khomyn, M., Klein, T., Klein, O., Klos, A., Koetter, M., Krahnen, J. P., Kolokolov, A., Korajczyk, R. A., Kozhan, R., Kwan, A., Lajaunie, Q., Lam, F. E., Lambert, M., Langlois, H., Lausen, J., Lauter, T., Leippold, M., Levin, V., Li, Y., Li, M. H., Liew, C. Y., Lindner, T., Linton, O. B., Liu, J., Liu, A., Llorente, G., Lof, M., Lohr, A., Longstaff, F. A., Lopez-Lira, A., Mankad, S., Mano, N., Marchal, A., Martineau, C., Mazzola, F., Meloso, D., Mihet, R., Mohan, V., Moinas, S., Moore, D., Mu, L., Muravyev, D., Murphy, D., Neszveda, G., Neumeier, C., Nielsson, U., Nimalendran, M., Nolte, S., Norden, L. L., O'Neill, P., Obaid, K., Ødegaard, B. A., Östberg, P., Painter, M., Palan, S., Palit, I., Park, A., Pascual, R., Pasquariello, P., Pastor, L., Patel, V., Patton, A. J., Pearson, N. D., Pelizzon, L., Pelster, M., Pérignon, C., Pfiffer, C., Philip, R., Plíhal, T., Prakash, P., Press, O.-A., Prodromou, T., Putniņš, T. J., Raizada, G., Rakowski, D. A., Ranaldo, A., Regis, L., Reitz, S., Renault, T., Renjie, R. W., Renò, R., Riddiough, S., Rinne, K., Rintamäki, P., Riordan, R., Rittmannsberger, T., Rodríguez-Longarela, I., Rösch, D., Rognone, L., Roseman, B., Rosu, I., Roy, S., Rudolf, N., Rush, S., Rzayev, K., Rzeźnik, A., Sanford, A., Sankaran, H., Sarkar, A., Sarno, L., Scaillet, O., Scharnowski, S., Schenk-Hoppé, K. R., Schertler, A., Schneider, M., Schroeder, F., Schuerhoff, N., Schuster, P., Schwarz, M. A., Seasholes, M. S., Seeger, N., Shachar, O., Shkilko, A., Shui, J., Sikic, M., Simion, G., Smales, L. A., Söderlind, P., Sojli, E., Sokolov, K., Spokeviciute, L., Stefanova, D., Subrahmanyam, M. G., Neusüss, S., Szaszi, B., Talavera, O., Tang, Y., Taylor, N., Tham, W. W., Theissen, E., Thimme, J., Tonks, I., Tran, H., Trapin, L., Trolle, A. B., Valente, G., Van Ness, R. A., Vasquez, A., Verousis, T., Verwijmeren, P., Vilhelmsson, A., Vilkov, G., Vladimirov, V., Vogel, S., Voigt, S., Wagner, W., Walther, T., Weiss, P., van der Wel, M., Werner, I. M., Westerholm, P. J., Westheide, C., Wipplinger, E., Wolf, M., Wolff, C. C. P., Wolk, L., Wong, W.-K., Wrampelmeyer, J., Xia, S., Xiu, D., Xu, K., Xu, C., Yadav, P. K., Yagüe, J., Yan, C., Yang, A., Yoo, W., Yu, W., Yu, S., Yueshen, B. Z., Yuferova, D., Zamojski, M., Zareei, A., Zeisberger, S., Zhang, S. S., Zhang, X., Zhong, Z., Zhou, Z. I., Zhou, C., Zhu, S., Zoican, M., Zwinkels, R. C. J., Chen, J., Duevski, T., Gao, G., Gemayel, R., Gilder, D., Kuhle, P., Pagnotta, E., Pelli, M., Sönksen, J., Zhang, L., Ilczuk, K., Bogoev, D., Qian, Y., Wika, H. C., Yu, Y., Zhao, L., Mi, M., Bao, L., Vaduva, A., Prokopczuk, M., Avetikian, A., and Wu, Z.-X. (2023). Non-standard errors. *The*

Journal of Finance, (forthcoming).

Menkveld, A. J. and Saru, I. L. (2023). Who knows? Information differences between trader types. *Working Paper*.

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

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

Nagel, S. and Xu, Z. (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 Veronesi, P. (2009). Learning in financial markets. *Annual Review of Financial Economics*, 1:361–381.

Routledge, B. R. and Zin, S. E. (2009). Model uncertainty and liquidity. *Review of Economic Dynamics*, 12(4):543–566.

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

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

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.

Wu, G. (2001). The determinants of asymmetric volatility. *The Review of Financial Studies*, 14(3):837–859.

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

Zhou, T. (2021). Ambiguity, asset illiquidity, and price variability. *Journal of Economic Behavior & Organization*, 191:280–292.