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

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

December 14, 2023

*We thank Adelina Barbalau, Ella Dias Saraiva-Patelli, Ines Chaieb, Bernard Dumas, Philipp Illeditsch, Yehuda Izhakian, Dongliang Lu, Albert Menkveld, Alessandro Melone, Hannes Mohrschladt, Martin Schneider, Raman Uppal, Youchang Wu, Goufu Zhou, and seminar participants at the Federal Reserve Bank of Atlanta, the 2023 European Meeting of the Econometric Society, the 29th Annual Meeting of the German Finance Association 2023, 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, the 2022 Computational and Financial Econometrics conference, and the 2023 Northern Finance Association Conference.

†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 have more risk-bearing capacity, 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 provide empirical support of our model predictions using a novel dataset of trading activity on Euro Stoxx 50 futures contracts.

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 parameters governing the endowment process by observing its realizations over time. While they all have access to the same information, agents differ in their confidence in the quality of the parameter estimates. We capture this difference by assuming that some agents are more averse than others to the uncertainty, or ambiguity, surrounding parameter estimates. Following insights from the decision theory literature, e.g., Bewley (2011), we model ambiguity as a confidence interval that represents the set of posterior distributions of a "Knightian" decision maker. Large and unexpected dividend realizations, by increasing investors' volatility estimates, widen their confidence interval. Hence, in this setting, agents' learning about volatility gives rise to time-varying ambiguity.

Our main result is to show that in equilibrium an increase in estimated volatility induces ambiguity-averse agents to *increase* the weight of risky assets in their portfolio. This result seems counterintuitive, because an increase in perceived volatility leads to less precisely estimated model parameters, increases the amount of ambiguity, and thus makes the risky asset less attractive to an ambiguity-averse agent. This logic, while valid in partial equilibrium, ignores that in general equilibrium prices adjust to reflect agents' perceived volatility and ambiguity. As a result of these equilibrium adjustments, ambiguity-averse agents are relatively more willing to take on risk, making them natural buyers of risky assets when perceived volatility rises.

[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.

We first illustrate the main intuition for this result in a two-period model that we can solve analytically. In the model agents do not know the mean of the dividend, but know perfectly its volatility. Because their choices are based on an estimate of the mean, we assume that agents differ in their aversion to the uncertainty of this estimate. This heterogeneity could be capturing differences in information, inventory or agency costs, or behavioral difference among agents. We show that, in equilibrium, while aversion to ambiguity steers the ambiguity-averse agents towards more conservative portfolios, their risky holdings increase with the dividend volatility. Therefore, an increase in volatility is associated with an equilibrium *flow* of risky assets from ambiguity-neutral sellers to ambiguity-averse buyers. Intuitively, an increase in volatility raises the risk premium in equilibrium. However, because ambiguity-averse agents hold a conservative portfolio, they have more risk-bearing capacity and therefore tend to perceive the new risk premium as "too high" for their existing portfolio. Hence, they desire to increase their risk exposure. Similarly, ambiguity-neutral agents perceive the new risk premium as "too low" and are willing to sell risky assets. These volatility-induced flows emerge because of heterogeneity in agents' aversion to parameter uncertainty. They would not emerge, for example, in a model with heterogeneity in agents' risk aversion.

An implication of this finding is that, in a model with ambiguity-sensitive agents, time-variation in agents' perceived, or subjective, volatility becomes particularly relevant for understanding the dynamics of equilibrium flows and risk premia. Learning about the dividend volatility, not just the mean, is an important source of time-variation in agents' perceived volatility. It is a commonly accepted principle that volatility is relatively easier to estimate than the mean and therefore learning about it has little impact on portfolios and equilibrium returns (e.g., Collin-Dufresne et al., 2016b; Buss et al., 2021). While this is true when agents have access to a continuous stream of information, the same might not hold when information arrives in "bursts" such as during earnings or macroeconomic announcements.

To illustrate the effect of learning and ambiguity on equilibrium flows and risk premia, we generalize the two-period model to an infinite-horizon setting, where overlapping generations of heterogeneous agents learn about the mean and the variance of the endowment process. Learning about variance introduces an important technical challenge. In fact, it is well-known that when both the mean and the variance of normally distributed dividends are unknown, their predictive distribution is a Student-t. Because of the fat tails of this distribution, expected utility is not well-defined for a wide class of commonly used utility functions, see, e.g., Geweke (2001). We overcome this difficulty by imposing a restriction on the true dividend variance and applying 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. Under this assumption, we show that the predictive distribution of dividends is a "dampened Student-t", i.e., a Student-t distribution with thinner tails. This allows us to fully characterize the equilibrium with learning about both mean and variance. We show that, consistent with the intuition from the two-period model, learning about variance generates equilibrium flows where, as perceived volatility increases, ambiguity-averse agents buy the risky asset and ambiguity-neutral agents sell. Furthermore, we show that learning about variance generates returns that are left-skewed, even with an iid-normal endowment process, consistent with the early findings of French et al. (1987) and, more recently, Chen et al. (2001) and Robins and Smith (2023).

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 introducing a "time distortion" to capture possible "information leakages" that occur 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 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—who act on behalf of clients—sell, and proprietary traders—who act on their own account—buy. This pattern 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.

[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.

Literature. 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. 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 analytical tractability and the fact that it is easy to learn the variance from a large sample of continuous observations. 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] 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".

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 endowed with non-expected utility preferences. We focus on one form of deviation from expected utility, that is, 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.

[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.

[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.

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 and return 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.

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 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

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

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 $\gamma > 0$.[10] 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.[11] 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. We model aversion to uncertainty by exploiting the connection between classical confidence interval and "Knightian" uncertainty, or ambiguity, as in Bewley (2011). Specifically, we assume that ambiguity is represented by "multiple priors" about the distribution of $\tilde{d}$ and that A-agents are averse to this ambiguity. We characterize the set of priors as a "confidence interval" around the mean estimate, m, whose size depends on the standard error s and the agents' degree of ambiguity aversion, that is,

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

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

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

[9]If agents enter the initial date having observed a time series of t dividend realizations, these quantities are, respectively, $m = \frac{1}{t}\sum_{k=1}^{t} d_k$ and $s = \sigma/\sqrt{t}$.

[10]The model can be easily extended to the case of heterogeneous risk aversion. However, none of the results we present are qualitatively affected by difference in risk aversion. Therefore, for ease of exposition, we assume that both agents have identical absolute risk aversion γ.

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

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). A-agents 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[12]

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

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 participate, i.e., $\theta^A = 0$, no distinct prior is selected. Optimal portfolios are given by

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

From the expression of θ^A in the above equation, we infer that agent A's problem in equation (6) is equivalent to B's problem but with a "distorted" belief about the expected dividend. While B

[12] 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).

agents' subjective mean dividend is $\mathbb{E}^B(\tilde{d}) = \mu^B = m$, A agents' subjective mean is $\mathbb{E}^A(\tilde{d}) = \mu^A$, given in equation (7).

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, $|\theta^A| < |\theta^B|$, when the price p is sufficiently high or low, and do not participate, $\theta^A = 0$, when the price falls in the range $p \in \left[\frac{m-\kappa s}{1+r}, \frac{m+\kappa s}{1+r}\right]$. In equilibrium, $\theta^A + \theta^B = 1$ and there cannot be short positions in

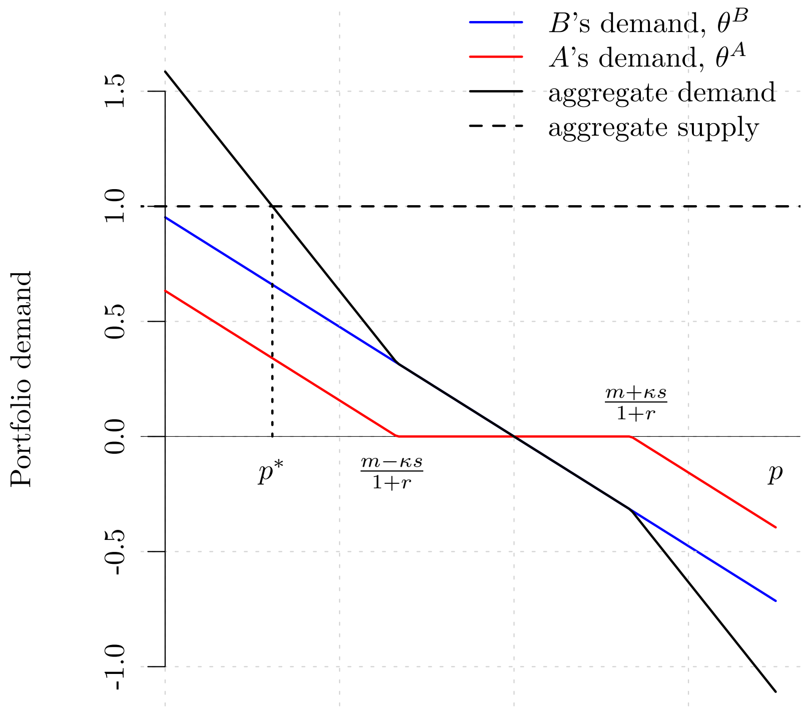

Figure 1: **Risky asset demand.** The figure shows the risky asset demand θ^B and θ^A from equation (8) as a function of the risky asset price p. The red line denotes type-A's demand; the blue line type-B's demand; the black line is aggregate demand; and the dashed line is the aggregate supply of the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

the risky asset. The following proposition characterizes the equilibrium price p of the risky asset.

Proposition 1. *The equilibrium price p of the risky asset is given by*

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

where the risk premium λ *is*

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

Ambiguity averse agents participate if their coefficient of ambiguity aversion $\kappa \le \kappa^*$*.*

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, s, is proportional to the dividend variance, $s \propto \sigma$, equation (10) shows that when both agents participate, i.e., $\kappa \le \kappa^*$, 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), that is, 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 portfolios are

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

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 $\theta^A = \theta^B = 1/2$.

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

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

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 $\Delta\mu = \kappa\sigma$ 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. More generally, as long as agents disagree on the mean dividend, volatility plays a key role in the determination of equilibrium holdings and flows.

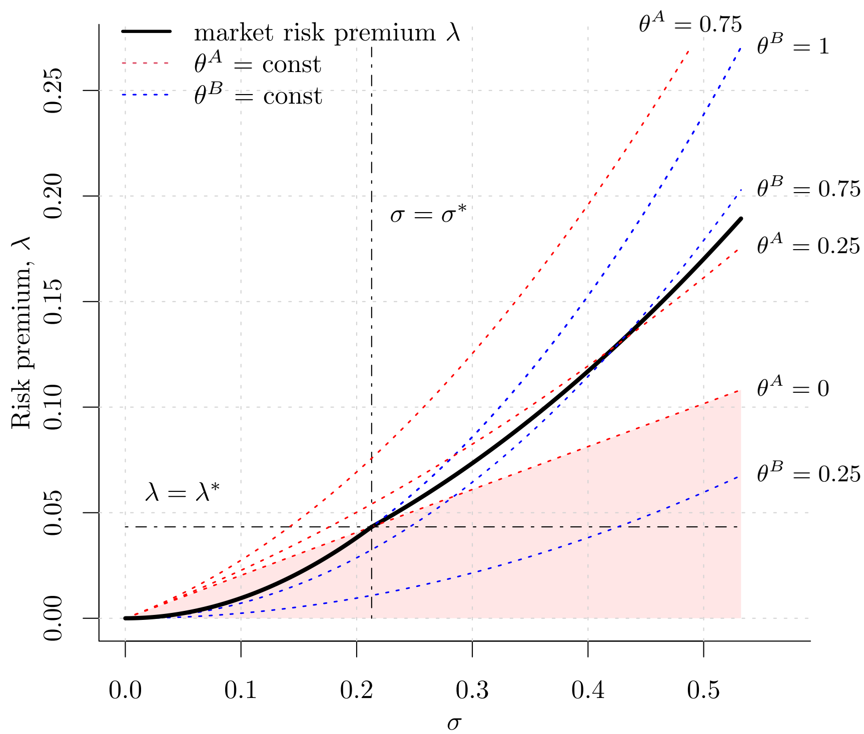

Figure 2: **Equilibrium portfolios and risk premia.** The figure shows iso-portfolios lines of type-A (red-dashed) and type B (blue-dashed) agents. These lines represent the set of volatility and risk premium values (σ, λ) that correspond to a given constant portfolio weight in equation (8), where the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

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 (10), A-agents participate only when their ambiguity aversion $\kappa < \kappa^* \equiv \gamma \frac{\sigma^2+s^2}{s}$, or, equivalently, if σ is sufficiently high, $\sigma > \sigma^*$ in Figure 2.[13]

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 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 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.

So far, 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.

[13] If agents enter the initial date having observed a history of t dividend realizations, the standard error is $s = \sigma/\sqrt{t}$. Hence, the participation condition $\kappa < \kappa^* = \gamma \frac{t+1}{\sqrt{t}} \sigma$ can be equivalently expressed in terms of volatility as $\sigma > \sigma^* \equiv \frac{\sqrt{t}}{t+1} \frac{\kappa}{\gamma}$.

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 $t = 1, \dots, \infty$ 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.[14]

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 assets. 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

[14]To simplify the exposition of the model, in this section 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.

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.

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.[15] 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 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 which we now discuss in detail.

Information updating. While dividend variance is unknown, all generations agree that the dividend variance is finite, or equivalently, that its precision $\phi \equiv 1/\sigma^2$ is bounded, $\phi \in [\underline{\phi}, \overline{\phi}]$ where the bounds $0 < \underline{\phi} < \overline{\phi} < \infty$ are common across all generations. At each time t, both types of agents

[15] 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.

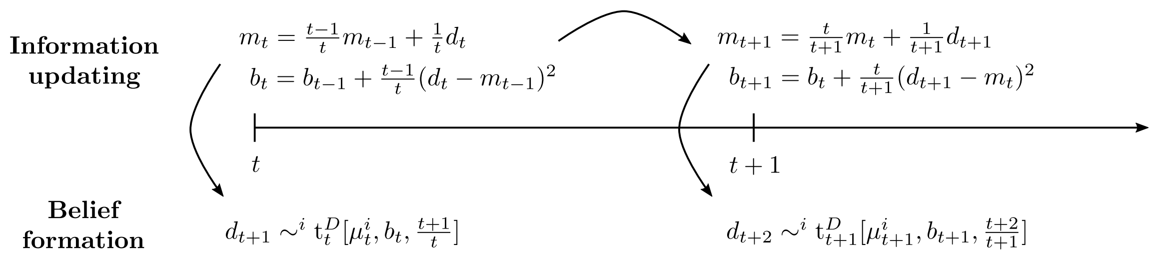

Figure 3: **Information updating and belief formation.** The top part of the figure illustrates the "information updating" step in which agents update the state variables m_t and b_t after observing a new dividend realization d_{t+1}. The bottom part of the figure illustrates the "belief formation" step, and shows how agents at time t use their updated information to form predictive distributions about the future dividend. Agents' beliefs are represented by the predictive distribution $t_{\nu_t}^D[\cdot,\cdot,\cdot]$ which is a dampened Student-t with ν_t degrees of freedom, as defined in equation (23).

inherit information about μ and σ from the previous generation. This information consists of a normal/inverse-Gamma pair of posteriors. Agents use these posteriors as their prior and update them after observing the new dividend, d_t. Specifically, we assume that generation $t-1$, having observed a time series of $t-1$ dividend realizations, hands over to generation t a prior for ϕ that is a *truncated* Gamma with $t-1$ degrees of freedom[16] and shape parameter b_{t-1}, i.e.,

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

The shape parameter b_{t-1} is a weighted sum of historical squared errors (see, equation (19) below) and, hence, $b_{t-1}/(t-2)$ can be thought of an estimate of the variance.

Conditional on the precision ϕ, agents' prior about the mean μ is normally distributed, that is,

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

with $t-1$ the number of observations, measuring the precision of the μ-prior. In sum, the information that generation $t-1$ hands over to generation t consists of the estimates of the sample mean, m_{t-1}, and of the sample variance, $\hat{\sigma}_{t-1}^2 = b_{t-1}/(t-2)$. We can think of (m_{t-1}, b_{t-1}) as a "model of the world", agreed upon by all agents.

[16]The degrees of freedom describe the precision of the ϕ-prior. In classical statistics, when estimating mean and variance of a normal distribution, setting the degrees of freedom of the Gamma distribution to the number of observation minus one guarantees a bias free estimate of the variance, see, e.g., Greene (2020).

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,

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

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

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,

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

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

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 (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 to 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 to 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.[17]

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). B-agents are ambiguity neutral and use as their unique prior the time-t posterior of the information-updating step given in equation (17). Conditional on ϕ, their prior about the mean μ

[17]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.

is normal with mean $\mu_t^B = m_t$,

$$\mu|\phi, t \sim^B \mathcal{N}\left(\mu_t^B, \frac{1}{t\,\phi}\right), \quad \mu_t^B = m_t. \tag{20}$$

In contrast, A-agents are ambiguity averse and entertain the following set of priors of μ,

$$\mu|\phi, t \sim^A \mathcal{N}\left(\tilde{\mu}_t^A, \frac{1}{t\,\phi}\right), \quad \tilde{\mu}_t^A \in \mathcal{P}_t \tag{21}$$

where $\mathcal{P}_t$ is the confidence interval for μ

$$\mathcal{P}_t \equiv [m_t - \kappa s_t, m_t + \kappa s_t], \tag{22}$$

with $s_t = \sqrt{\frac{1}{t}\mathbb{E}[\frac{1}{\phi|b_t}]}$ an estimate of the standard error of the dividend mean. Following Bakshi and Skoulakis (2010), in Lemma B.1 we show that, under these assumptions, agents i's predictive distribution of future dividends is a "dampened" Student-t, whose density is formally defined in Appendix B,

$$d_{t+1} \sim^A \mathrm{t}_t^D\left[\tilde{\mu}_t^A, b_t, \frac{t+1}{t}; \underline{\phi}, \overline{\phi}\right], \quad d_{t+1} \sim^B \mathrm{t}_t^D\left[\mu_t^B, b_t, \frac{t+1}{t}; \underline{\phi}, \overline{\phi}\right], \tag{23}$$

where $\mu_t^B = m_t$ and $\tilde{\mu}_t^A \in \mathcal{P}_t$. Because the dampened Student-t distribution has thinner tails than the Student-t, the agents' expected CARA utility is well-defined. This allows us to solve for optimal portfolios and construct an equilibrium.

3.3 Equilibrium

Investors in the model have rational expectations, that is, the pricing function they *perceive* corresponds to the *true* pricing function. Because parameters need to be learned, rational expectations do not imply that investor's subjective beliefs about the distribution of dividends correspond to the true distribution.

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

$$W_{t+1}^i = W_t^i(1+r) + \theta_t^i\left(p_{t+1} + d_{t+1} - p_t(1+r)\right), \quad i = A, B. \tag{24}$$

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,

$$\max_{\theta_t^B} \mathbb{E}_t[u(W_{t+1}^B)], \tag{25}$$

and

$$\max_{\theta_t^A} \min_{\tilde{\mu}^A \in \mathcal{P}_t} \mathbb{E}_t[u(W_{t+1}^A)], \tag{26}$$

where $\mathcal{P}_t$ is the ambiguity set defined in equation (22). As in the two-period model of Section 2, A-agents select either the prior $\tilde{\mu}_t^A = \mu_t^A \equiv m_t - \kappa s_t$ or they do not participate. Hence, participating ambiguity-averse A-agents select portfolios using a distorted belief, $\mathbb{E}_t^A(\tilde{d}_{t+1}) = \mu_t^A = m_t - \kappa s_t$, relative to that of the B-agents, $\mathbb{E}_t^B(\tilde{d}_{t+1}) = m_t$. 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. We show that the determination of participating agents' expected utility in each generation only requires the evaluation of a well-behaved integral. Agent i's demand θ_t^i is obtained by solving the first-order conditions

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

Unfortunately, there is no closed-form solution for agents' demand in this setting. To gain intuition, we solve the special case of 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 (20) and (21) with $\phi = 1/\sigma^2$ known. Then, the equilibrium price of the risky asset when both agents participate is*

$$p_t = \frac{1}{r} m_t - \Lambda_t, \tag{27}$$

where the risk premium Λ_t is given by

$$\Lambda_t = g_t \frac{\kappa}{2} \sigma + f_t \frac{\gamma}{2} \sigma^2, \tag{28}$$

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

$$\theta_t^A = \frac{1}{2} - \frac{\kappa}{2\gamma}\left(\frac{r\sqrt{t}}{1+r(t+1)}\right)\frac{1}{\sigma} \quad \textit{and} \quad \theta_t^B = \frac{1}{2} + \frac{\kappa}{2\gamma}\left(\frac{r\sqrt{t}}{1+r(t+1)}\right)\frac{1}{\sigma}, \tag{29}$$

with t denoting the number of dividend realization observed by agents in generation t.

The equilibrium weights (29) 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., $\partial\theta_t^A/\partial\sigma > 0$.[18]

For the case in which variance is not known, we cannot obtain an analytic expression for the equilibrium price. Aided by the structure of the equilibrium with known variance in Proposition 2, in Appendix D we numerically construct an equilibrium with unknown variance in which the price is of the form

$$p(m_t, b_t) = \frac{1}{r}m_t - \Lambda(b_t). \tag{30}$$

The key feature of this equilibrium is that, unlike the case of known variance in Proposition 2 where the risk premium Λ_t is a deterministic function of time, here it is a function of b_t, that is, it depends explicitly on the agent's perceived variance. Appendix D provides details of the numerical procedure we use to construct the equilibrium.

Portfolio flows. From the equilibrium prices $p(m_t, b_t)$, we can derive the equilibrium portfolio weights $\theta^i(m_t, b_t)$ for type-i agents in generation t. We define portfolio flows as

$$\Delta\theta^i(m_t, b_t) = \theta^i(m_t, b_t) - \theta^i(m_{t-1}, b_{t-1}), \quad i = A, B. \tag{31}$$

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.

[18] As $t \to \infty$ agents eventually learn the true mean μ and the weights in equation (29) converge to $1/2$. In Section 3.4 we extend the model to allow for perpetual learning.

3.4 Perpetual learning

The analysis so far has 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 acquired in the past is lost as generations overlap. We show that such an information leakage is a tractable way to model "fading memory", e.g., Nagel and Xu (2021), or "age-related experiential learning", e.g., Malmendier and Nagel (2016), Collin-Dufresne et al. (2016a), and Ehling et al. (2018). Our approach is similar to the "discount rate approach" of West and Harrison (2006) and consists of discounting the information content of moments estimates as time evolves.

To model perpetual learning, we assume that when generations 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 $\omega \in (0, 1)$.[19] In particular, time distortion introduces a "new clock": at time t the number of observation that is relevant for information updating, is not t but n_t. We refer to n_t as the "effective number of observations." Because of time distortion, n_t does not increment by one each time a dividend is observed but instead gets distorted first by ω. Hence,

$$n_{t+1} = \omega n_t + 1. \tag{32}$$

Intuitively, the quantity n_t runs slower than t. In fact, it runs slower and slower as it approaches its asymptotic value $\bar{n} = \frac{1}{1-\omega}$. This device allows us to achieve perpetual learning in our model while keeping most of the analysis of Section 3.2 unaffected.

With perpetual learning, the prior of the precision $\phi|t$ in equation (16) changes to

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

[19] 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.

where ν_t denotes the degrees of freedom, which, in this case, are no longer equal to $t-1$. The prior of $\mu|\phi, t$ in equation (17) becomes

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

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

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

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

and the updating equation for the degrees of freedom ν_t is

$$\nu_{t+1} = \omega \nu_t + 1. \tag{37}$$

By setting the time distortion parameter $\omega = 1$, we can recover the case described in Section 3.2. Notice that, as t increases, n_t and ν_t approach the asymptotic value $\overline{n} = \overline{\nu} = \frac{1}{1-\omega}$. 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. Appendix D describes the details of the equilibrium construction. 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 are *increasing* and those of the Bayesian agent B are decreasing in $\hat{\sigma}_t = \sqrt{b_t/\nu_t}$, which is an indicator of the estimated dividend volatility.[20] In the figure we set the truncation bounds $\underline{\phi}$ and $\overline{\phi}$ such that the true volatility $\sigma \in [0.1, 0.6]$. When $\hat{\sigma}_t$ takes values outside this range, agents become very confident that the true σ is either close to 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 b_t.

[20] If the dividend precision $\phi = 1/\sigma^2$ is Gamma distributed, $\mathbb{E}(\phi) = \nu_t/b_t$. In our application, the truncation bounds $\underline{\phi}$, $\overline{\phi}$ influence expected precision, however, b_t/ν_t is a useful measure of dividend variance.

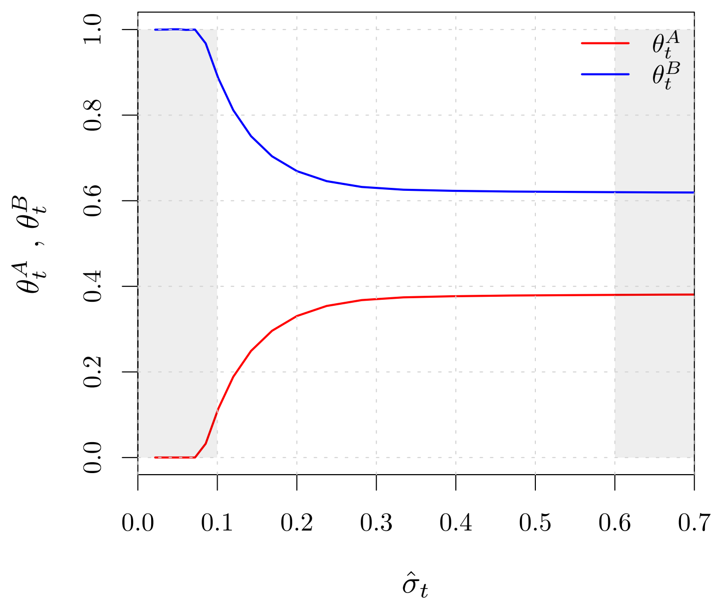

Figure 4: **Equilibrium portfolios and volatility.** The figure shows A and B agents' equilibrium portfolios as a function of $\hat{\sigma}_t = \sqrt{b_t/\nu_t}$, where b_t is the sum of squared error and ν_t the degrees of freedom. Parameters values: $\kappa = 1$, $\underline{\phi}$ and $\overline{\phi}$ 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 $\hat{\sigma}_t$, shown in Figure 4, has a direct counterpart in terms of dividend "surprises", i.e., deviations of the realized dividend from the historical average $e_t = d_t - m_{t-1}$. Figure 5 shows the equilibrium risky asset holdings of A-agents (left panel) and B-agents (right panel) as a function of this surprise. We assume $\underline{\phi}$ and $\overline{\phi}$ such that $\sigma \in [0.1, 0.6]$ and start with $t-1$-volatility estimate $\hat{\sigma}_{t-1} = \sqrt{b_{t-1}/\nu_{t-1}} = 0.3$, which lies within the interval of possible values of σ. Observations of the surprise e_t lead to an update of b_t according to the recurrence (36) which in turn leads to updates of the equilibrium portfolio weights (solid lines) relative to the time $t-1$ weights (flat dotted lines). Different colors 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 (large positive or negative values of e_t) are associated with an upward revision of b_t relative to b_{t-1}. The U-shape nature of the equilibrium portfolios in the left panel of Figure 5 implies that after large positive and negative surprises A-agents *increase* their risky asset holdings, i.e., surprises generate flows from ambiguity neutral B-agents (sellers) to ambiguity averse A-agents (buyers). If the dividend

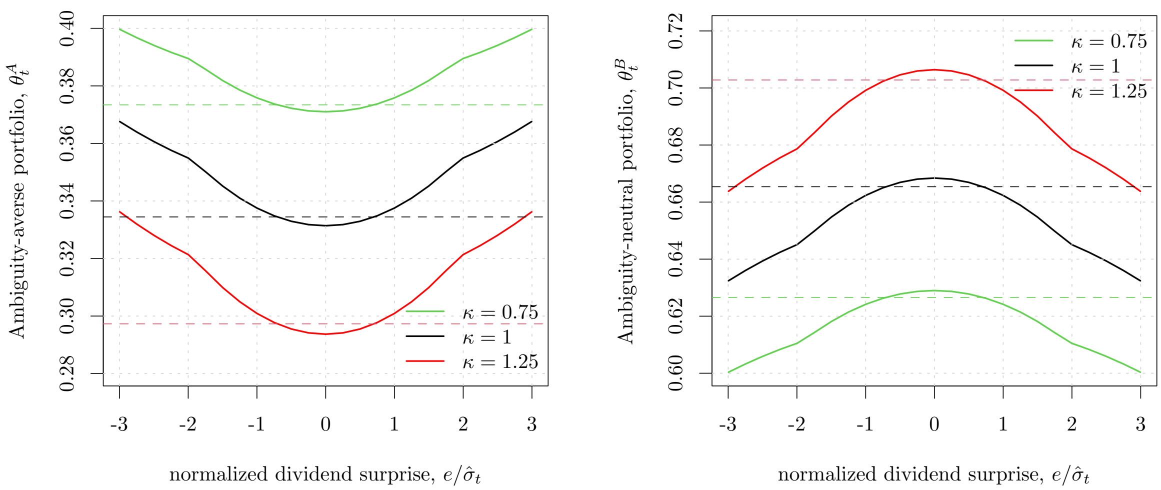

Figure 5: **Equilibrium portfolios and dividend surprise $\mathbf{e_t = d_t - m_{t-1}}$.** The figure shows the equilibrium portfolio of A-agents (left panel) and B-agents (right panel) as functions of the normalized dividend surprise $e_t/\hat{\sigma}_{t-1}$ after b_{t-1} is updated with observed e_t. Different lines corresponds to different values of the ambiguity aversion parameter κ. Large positive and negative dividend surprises lead to high values of b_t with increased risky holdings of A-agents relative to the time $t-1$ holdings – indicated by dotted horizontal lines. Parameter values: $\underline{\phi}$ and $\overline{\phi}$ such that $\sigma \in [0.1, 0.6]$, $\hat{\sigma}_{t-1} = \sqrt{b_{t-1}/\nu_{t-1}} = 0.3$, $r = 0.1$, and $\nu_{t-1} = \overline{n} = \frac{1}{1-\omega} = 20$.

realization d_t is close to the estimated mean m_{t-1}, the estimated volatility decreases, giving rise to portfolio flows from A-agents (sellers) to B-agents (buyers). 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 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 (31). 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 dividend mean and volatility, the main intuition developed in the simple two-period model of Section 2.

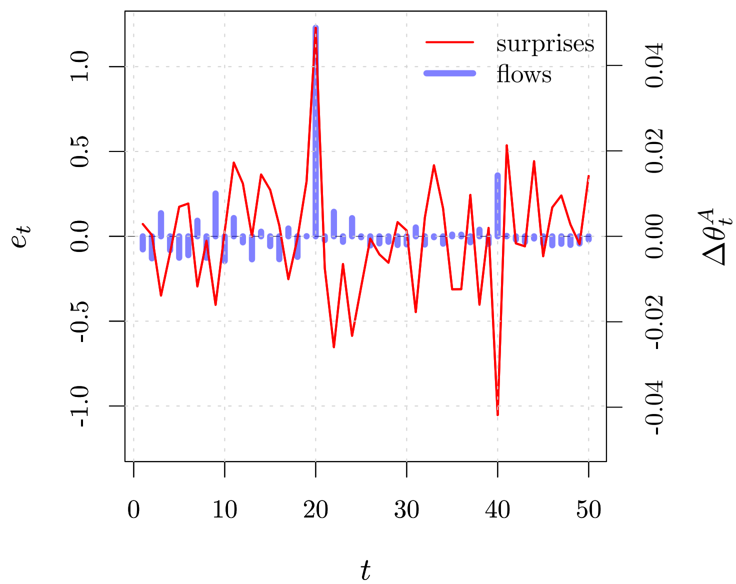

Figure 6: **Portfolio flows on a simulated path.** The figure illustrates dividend surprises, $e_t = d_t - m_{t-1}$ (left axis), and portfolio flows, $\Delta\theta_t^A$ (right axis) on a dividend path that is computed using true volatility. The starting value b_0 is chosen such that $\hat{\sigma}_0 = \sqrt{b_0/\overline{n}}$ equals the true dividend volatility, $\hat{\sigma}_0 = \sigma$. Parameters values: $\kappa = 1$, $\sigma = 0.3$, $\underline{\phi}$ and $\overline{\phi}$ such that the true volatility $\sigma \in [0.1, 0.6]$, $r = 0.1$, and $\nu_t = \overline{n} = \frac{1}{1-\omega} = 20$.

3.6 Implications for return predictability

As the expression for the equilibrium price in equation (30) shows, when variance is unknown the risk premium $\Lambda(b_t)$ depends on the state variable b_t and is therefore time-varying. This is in contrast to the case of known variance shown in Proposition 2 where 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.

Figure 7 plots the equilibrium risk premium in the steady state of the OLG economy with perpetual learning, as a function of $\hat{\sigma}_t \equiv \sqrt{b_t/\nu_t}$, an estimate of volatility. The figure shows that the risk premium is an increasing function of the estimated volatility. Therefore, agents expect higher returns after large positive or negative dividend surprises. This finding is consistent with Nagel and Xu (2022) who, based on an analysis of CFO survey data, show that the subjective risk premium is positively related to subjective estimates of variance and that CFOs' subjective return expectations strongly depend on realized variance. In states of the world with surprising dividend realizations, as shown in Figure 4, B-agents sell and A-agents buy and asset prices are low (risk-premia are high). This implies that ambiguity-averse A-agents are willing to absorb flows from selling B-agents at a price that is reflected in higher risk-premia. Ambiguity aversion of A agent can then provide a micro-foundation of asset demand elasticity (e.g., Koijen and Yogo, 2019).

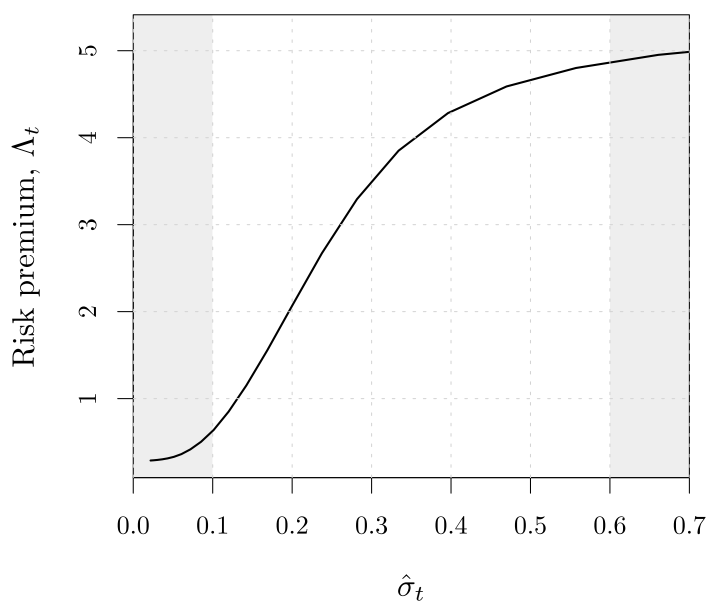

Figure 7: **Equilibrium risk premium.** The figure shows the equilibrium risk premium Λ_t as a function of $\hat{\sigma}_t = \sqrt{b_t/\nu_t}$ when both the dividend mean and variance are unknown. Parameters values: $\kappa = 1$, $\underline{\phi}$ and $\overline{\phi}$ such that the true volatility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

Figure 8 illustrates the dynamic of the risk premium $\Lambda_t(b_t)$ for a simulated random path of dividend surprises, $e_t = d_t - m_{t-1}$. 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. In sum, time-varying subjective risk premia is a source of return predictability in our model. Note that e_t is the dividend surprise that time t agents experience after observing the dividend realization at the beginning of their life. Because all agents inherit the state variable m_{t-1}, the surprise e_t is therefore independent of the agent's type. Time t agents, however, form different beliefs about the dividend at the end of their life at $t+1$. Hence, end-of-life dividend surprises vary with agent type. While these differences must be properly accounted for when calculating equilibrium prices and portfolio weights, for illustrative purposes, in the figure we only show the path-dependency of the risk premium as a function of the beginning-of-life surprise e_t.

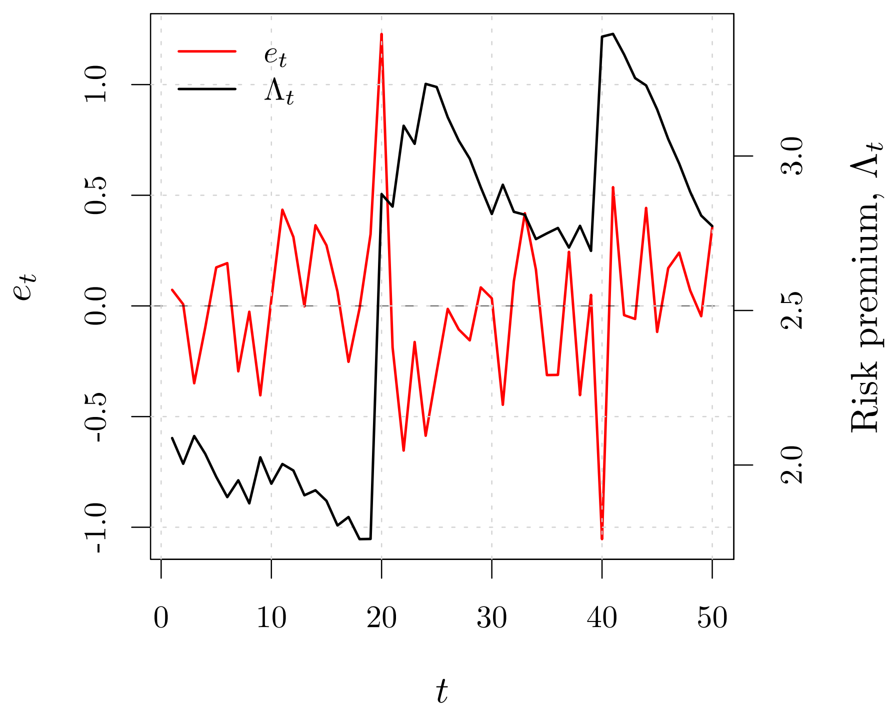

Figure 8: **Risk premium on a simulated path.** The figure illustrates surprises $e_t = d_t - m_{t-1}$ (left axis), with the corresponding equilibrium risk premium Λ_t (right axis) on a dividend path that is computed using the true volatility. The state variable b_0 is chosen such that $\hat{\sigma}_0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

As Lewellen and Shanken (2002) emphasize, to understand predictability in the presence of parameter uncertainty, it is useful to distinguish between agents' beliefs (subjective expectations) and the true data generating process (objective expectations). The two coincide in the absence of parameter uncertainty. With parameter uncertainty, investors' perceived properties of returns can differ substantially from the empirical properties detected by an "objective" observer, e.g., the econometrician. The next proposition formalizes the objective and subjective expected returns in our model.

Proposition 3. *Consider an equilibrium in the OLG model with perpetual learning of Section 3.4, where the price in the steady state is given by*

$$p(m_t, b_t) = \frac{1}{r} m_t - \Lambda(b_t), \tag{38}$$

with $d\Lambda(b_t)/db_t > 0$. Then, the one-period expected return perceived by an objective observer is

$$\mathbb{E}_t^{obj}[p_{t+1} + d_{t+1}] - p_t \approx \mu - \frac{1}{r} \underbrace{(1-\omega)(m_t - \mu)}_{\text{Learning about the mean } \mu} + \underbrace{\left((1-\omega)b_t - \omega(\sigma^2 + (m_t - \mu)^2\right) \frac{d\Lambda(b_t)}{db_t}}_{\text{Learning about the variance } \sigma^2}, \tag{39}$$

with μ and σ denoting the true mean and volatility of the dividend process, and $\omega \in (0,1)$ denoting the degree of information leakage across generations. In contrast, the subjective one-period expected return perceived by ambiguity-averse and Bayesian investors are, respectively,

$$\mathbb{E}_t^{A}[p_{t+1} + d_{t+1}] - p_t \approx m_t - \kappa(1-\omega)\left(1 + \frac{1-\omega}{r}\right)\sqrt{b_t}, \tag{40}$$

$$\mathbb{E}_t^{B}[p_{t+1} + d_{t+1}] - p_t \approx m_t. \tag{41}$$

Equation (39) shows that learning about the mean and the variance of the dividend implies a countercyclical behavior in objective expected returns. Specifically, the second term on the right-hand-side of equation (39) shows that when the estimate m_t deviates from the true mean μ, objective observers predict a reversion towards the mean. This is the objective predictability discussed by Lewellen and Shanken (2002). In addition, the third term on the right-hand-side of equation (39) characterizes predictability resulting from learning about variance. Whenever investors overestimate the return variance, that is, they estimate a large value of b_t, the objective observer expects a lower future estimate of the variance. Because the risk premium $\Lambda(b_t)$ is positively related to the estimate of the variance b_t, when estimates of the return variance is high, the objective observer anticipates a lower future risk premium and hence higher future returns. Conversely, when the current estimate of the variance b_t is low, the objective observer expects higher future volatility, and hence lower future returns. Note that, while positive and negative surprises raise b_t, the mean

m_t increases with positive dividend realizations and decreases with negative dividend realizations. Therefore, an econometrician that observes ex-post dividend realizations would detect one-period returns that are amplified by bad news and dampened by good news, i.e., left-skewed.

The subjective expected return of the ambiguity-averse agent in equations (40) is always lower than that of the Bayesian agent in equation (41). Ambiguity aversion affects agents A's return expectations by reducing their subjective expectation of the future dividend d_{t+1} and their expectation of the future price p_{t+1}. When there is no information leakage, $\omega = 1$, the subjective expectation of A and B agents coincide. That is, in an OLG economy without perpetual learning, ambiguity vanishes and ambiguity aversion will have no impact on equilibrium prices.

Finally, note that, although dividends are iid-normal, equilibrium returns from the OLG model are left-skewed, a feature of stock-market indexes documented by, e.g., Chen et al. (2001) and, more recently, Robins and Smith (2023) who extend the results of French et al. (1987). Left skewness in returns originates from the dynamics of the state variable b_t which, from equation (36) can be thought of as the weighted sum of squared dividend surprises e_t^2. Because the dividend surprises e_t have a dampened-t distribution with truncation bounds $\underline{\phi}$ and $\overline{\phi}$, if these bounds are sufficiently far apart, the distribution of b_{t+1}, conditional on b_t, is approximately F with 1 and $\overline{\nu} = 1/(1-\omega)$ degrees of freedom.[21] This distribution is bounded from below and right skewed. The risk premium $\Lambda(b_{t+1})$ is an increasing function of b_{t+1} and has therefore also a right-skewed distribution. Because prices and risk premia are inversely related, a right-skewed distribution of risk premia implies a left-skewed distribution for prices and returns.

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.

[21] This result follows from the fact that the square of a t-distribution with n degrees of freedom has an F distribution with 1 and n degrees of freedom, $F(1, n)$. In our case e is dampened-t, and b_t is a weighted sum of squares of e, as per equation (36). Therefore, the convergence result to $F(1, \overline{\nu})$ only hold approximately when the truncation bounds $\underline{\phi}$ and $\overline{\phi}$ are sufficiently far apart. As $\overline{\nu} \to \infty$, i.e., $\omega \to 1$, the distribution $F(1, \overline{\nu})$ converges to a chi-squared distribution with one degree of freedom, which is also right skewed.

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.[22] 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 (in number of contracts) 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 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. Because of market clearing, we have $\Delta\theta^{\text{agency}} = -\left[\Delta\theta^{\text{proprietary}} + \Delta\theta^{\text{market maker}}\right]$. In the rest of the analysis, we use the label 'flows' to indicate the portfolio flows of agency traders, $\Delta\theta^{\text{agency}}$.

We take prices of Euro STOXX 50 and implied volatility (VSTOXX) from Refinitiv Datastream. In the rest of the analysis, we use the label IV (implied volatility) to refer to the volatility index VSTOXX.

4.2 Mapping the model to the data

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 obtain good empirical proxies for the dividend mean and variance. 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

[22] This data set was also used for 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 the data to us.

fear of trading against informed traders, see, e.g., Routledge and Zin (2009), Easley and O'Hara (2010), and Zhou (2021).[23] 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 dividends are not observed at a daily frequency, it is challenging to construct empirical proxies for the dividend mean and variance. To address this challenge, we rely on the structural link between prices, dividends, and portfolio flows in our model. Specifically, we use the historical mean returns as proxy for the dividend mean m_t. As proxies for the dividend variance b_t we use the historical return variance and the implied volatility VSTOXX, calculated from options on the Euro STOXX 50 index. A key implication of our model is that an increase in agents' perceived volatility is associated with portfolio flows from agency to proprietary traders. Because we can observe these flows in the data, we include them as a proxy for dividend volatility in our analysis. We formally test this link in Section 4.3.

4.3 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, ΔIV, from day $t-1$ to day t. Table 1 shows median values within quintile bins of volatility changes.

[23] 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.

ΔIV quintile	flows (1)	returns (2)	IV (3)	ΔIV (4)
1	12510.50	1.21	24.62	-1.68
2	5041.50	0.49	18.79	-0.63
3	533.00	0.05	17.66	-0.09
4	-5088.50	-0.30	19.26	0.45
5	-17533.00	-1.40	25.11	1.73

Table 1: **Flows and volatility.** The table shows agency traders' daily flows on the Euro Stoxx 50 futures (number of contracts), daily log excess returns (in %) of the Euro Stoxx 50, and the level, IV, and changes, ΔIV, of the implied volatility index VSTOXX. The data is from Eurex Exchange, and flow data are aggregated on a daily frequency. Observations are grouped in quintile bins for ΔIV, 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 ΔIV, 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.[24] 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).

ΔIV quintile	lagged return				
	1	2	3	4	5
1	9722	143148	14392	18922	11478
2	110	5350	7276	7285	7521
3	-544	-361	220	-1242	4888
4	-9933	-7255	-3065	-4689	815
5	-22149	-18471	-18926	-18544	-11478

Table 2: **Flows, volatility, and lagged returns.** The table shows agency traders' flows on the Euro Stoxx 50 futures depended on changes in the volatility index ΔIV (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 ΔIV and lagged returns (double sort), and median values of agents' flows are reported for each combination.

[24] 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.

Columns (2) and (4) in Table 1 show a strong and negative correlation between daily realized excess returns of the Euro Stoxx 50 and ΔIV. This empirical fact is known as the "asymmetric volatility phenomenon".[25] 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 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.

4.4 Return predictability

In this section we study the implication of our model for return predictability. The analysis in Section 3.6 highlights that learning about the mean and variance of the dividend process provides two separate channels through which return predictability might be detected in the data by an objective observer, e.g., an econometrician. Specifically, from equation (39) in Proposition 3 our model predicts that an objective observer would detect a negative relation between the mean estimate m_t and future returns. This effect is attributed to the process of learning about the mean, and it is the channel identified by Lewellen and Shanken (2002). In addition, our model shows that learning about variance introduces a new channel for return predictability. To understand this channel, note that the corresponding term in equation (39) implies that an objective observer who estimates a high variance (that is, a high value of b_t) would expect lower future variance. Because risk premia are positively related to estimates of b_t, lower future variance implies higher future prices. Hence, when the current estimate of the return variance is high, the objective observer

[25] 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 Bekaert and Wu, 2000; Wu, 2001).

expects higher future returns. In sum, the model predicts that future returns are negatively related to the current mean estimate m_t and positively related to the current variance estimate b_t.

	Dependent variable: future excess returns							
	1-day	2-day	5-day	21-day	1-day	2-day	5-day	21-day
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
lagged avg. return	−0.080	−0.077	−0.091	−0.026	−0.085	−0.085	−0.098	−0.029
	(0.189)	(0.180)	(0.150)	(0.181)	(0.189)	(0.181)	(0.152)	(0.182)
lagged variance	−0.011	−0.009	−0.005	−0.003	−0.011	−0.008	−0.004	−0.003
	(0.010)	(0.009)	(0.007)	(0.008)	(0.010)	(0.009)	(0.007)	(0.008)
lagged flows	−0.016	−0.019***	−0.008*	−0.001	−0.015*	−0.018***	−0.007*	−0.0002
	(0.010)	(0.007)	(0.005)	(0.003)	(0.009)	(0.007)	(0.004)	(0.003)
lagged IV^2	0.013	0.010	0.006	0.004	0.012	0.009	0.004	0.003
	(0.010)	(0.009)	(0.006)	(0.006)	(0.010)	(0.009)	(0.006)	(0.006)
avg. return × flows					3.941	6.583*	5.837	2.060*
					(4.849)	(3.477)	(3.572)	(1.155)
lagged return	0.003	0.008	−0.009	−0.004	0.004	0.010	−0.008	−0.003
	(0.023)	(0.016)	(0.009)	(0.004)	(0.023)	(0.016)	(0.010)	(0.004)
Constant	−0.0002	−0.0001	0.00001	0.00002	−0.0002	−0.0001	0.00002	0.00002
	(0.0003)	(0.0003)	(0.0003)	(0.0004)	(0.0003)	(0.0003)	(0.0003)	(0.0004)
Observations	4,771	4,770	4,767	4,751	4,771	4,770	4,767	4,751
R^2	0.003	0.006	0.007	0.005	0.003	0.007	0.009	0.006
Adjusted R^2	0.002	0.005	0.006	0.004	0.002	0.006	0.008	0.005
Residual Std. Error	0.015	0.010	0.006	0.003	0.015	0.010	0.006	0.003
F Statistic	3.083***	5.489***	6.507***	4.755***	2.735**	5.527***	7.421***	5.141***
Note:						*p<0.1; **p<0.05; ***p<0.01		

Table 3: **Risk premia and portfolio flows.** The table shows results for (average) Euro Stoxx 50 excess returns regressed on lagged excess returns, lagged mean excess returns and lagged variance (both over past three months), lagged agents' flows, and lagged implied variance IV^2. The data is from Eurex Exchange, and flow data are aggregated on a daily frequency. Standard errors are Newey-West corrected.

To test these predictions, we use the historical mean returns and variance, both computed over 63 days, as proxies for m_t ('avg. return') and b_t ('variance'), respectively. We also include a forward-looking measure of variance, which we take to be the square of the implied volatility, IV^2. Moreover, exploiting the structural link between flows and volatility from our model, we use portfolio flows ('flows', in million contracts) as a measure of shocks to perceived volatility in our predictability regressions. Our theory predicts a decline in volatility after a positive shock to perceived volatility, and hence a negative coefficient for agents' flows. Finally, we include lagged daily returns ('lagged return') in order to control for short-term reversals.

We report the results in Table 3 where each column represents continuous compounded average excess returns over the next 1, 2, 5, and 21 trading days. In columns (5)–(8), we introduce an interaction term, lagged avg. return × flows, to investigate potential variations in predictability strength as underlying uncertainty shifts over time. The evidence presented in Columns (1)–(4) of Table 3 shows that the variable with the strongest predictive power with respect to future returns is the lagged flow from agency to proprietary traders. According to our model, these flows are triggered by an increase in perceived variance. Therefore, this finding suggests that observable portfolio flows can serve as a useful signal for perceived variance. The estimates in Table 3 show that lagged portfolio flows significantly predict future returns up to a horizon of 5 trading days. In contrast, both the historical and implied volatility measures do not have predictive power for future returns, although the coefficient of implied volatility has a positive sign, consistent with the mean reversion implied by learning about variance.

According to equation (39), return predictability is the strongest when the "learning-about-the-mean" and "learning-about-the-variance" channels reinforce each other. This happens when either (i) the historical average return is low and the estimated variance is high at the same time or if (ii) the historical return is high and estimated variance is low. In the former case (low mean, high variance), we expect higher future returns while in the latter case (high mean and low variance), we expect lower future returns. We test this prediction in columns (5)–(8) where we add an interaction term between lagged average return and lagged agents' flows. The positive and significant coefficient estimate of the interaction term confirms our model predictions and provides evidence of return predictability up to 5 trading days. Predictive regressions of the 2-day return and the 5-day return show a high adjusted R^2 of up to 0.8%.

The negative, albeit insignificant, coefficient of the lagged average return confirms the mechanism of Lewellen and Shanken (2002) implied by learning about the mean. Finally, low and insignificant coefficients for lagged returns rule out that this predictability is triggered by short-term reversals in daily returns. In general, we observe a decline in predictive power at longer horizons. This finding is in line with our model, in which learning induces reversion in the dividend mean and variance.

5 Conclusion

We study equilibrium portfolio flows and asset prices in a model where agents learn about economic fundamentals and differ in their aversion to parameter uncertainty. We show that because ambiguity averse agents hold conservative portfolios, they have more risk-bearing capacity that make them natural buyers of risky assets when perceived volatility rises.

Our model highlights two key factors in the determination of equilibrium portfolio flows and risk premia: heterogeneity in investors' ambiguity attitudes and learning about variance. When all agents are ambiguity neutral or when variance is known to all investors, there are no flows in equilibrium, even if investors have heterogeneous risk aversion. Learning about variance introduces time-variation in risk premia and is therefore a source of return predictability. Because portfolio flows to ambiguity averse agents are closely associated with positive shocks to perceived variance, our model identifies portfolio flows as a potential explanatory variable in return predictability regressions.

We empirically investigate the predictions of the model using flow data on Euro Stoxx 50 futures contracts. We provide evidence that agency traders sell, and proprietary traders buy during periods of market uncertainty. Furthermore, portfolio flows significantly explain future returns in predictive regressions. 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

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

∎

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 (10), we have that $\partial\bar{\lambda}^A(\sigma)/\partial\sigma < \partial\bar{\lambda}^B/\partial\sigma$.*

Proof. Let $s = \sigma/\sqrt{t}$, with 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

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

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

We prove that along the equilibrium risk premium λ in equation (10) 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

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

or, rearranging,

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

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

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

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

which is true for $\kappa > 0$ and $t < \infty$ 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 (26) 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 $\theta^i_{t,\tau}$, $i = A, B$ is

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

where we denoted by $R \equiv (1 + r)$.

Using the fact that $p_{t,0} = 0$ for all t, we can construct the equilibrium in a $\tau = 1$ economy. In this economy, when both agents participate

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

where

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

Imposing market clearing we have

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

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

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

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

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

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

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

using the predictive distribution (A.12) we obtain

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

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

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

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

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

where

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

with

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

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

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

where

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

with

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

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

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

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

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

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

with

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

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:

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

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

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

and

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

∎

Proof of Proposition 3

Using the recursive structure of the state variables m_t and b_t in equations (35) and (36) in the definition of the expected return, we get

$$\begin{aligned}\mathbb{E}_t[p_{t+1}+d_{t+1}-p_t] &= \mathbb{E}_t[d_{t+1}]+\mathbb{E}_t\left[\frac{1}{r}m_{t+1}-\Lambda(b_{t+1})\right]-\left[\frac{1}{r}m_t-\Lambda(b_t)\right] \\ &\approx \mathbb{E}_t[d_{t+1}]+\frac{1}{r}\mathbb{E}_t[m_{t+1}-m_t]-\mathbb{E}_t[b_{t+1}-b_t]\frac{d\Lambda}{db}. \qquad \text{(A.28)}\end{aligned}$$

Prom the perspective of the objective observer these expectations are

$$\begin{aligned}\mathbb{E}_t^{\text{obj}}[d_{t+1}] &= \mu, \\ \mathbb{E}_t^{\text{obj}}[m_{t+1}] &= \omega m_t+(1-\omega)\mu, \\ \mathbb{E}_t^{\text{obj}}[b_{t+1}] &= \omega b_t+\omega\mathbb{E}_t^{\text{obj}}\left[(d_{t+1}-m_t)^2\right] \\ &= \omega b_t+\mathbb{E}_t^{\text{obj}}\left[(d_{t+1}-\mu+\mu-m_t)^2\right] \\ &= \omega b_t+\omega\Big[\underbrace{\mathbb{E}_t^{\text{obj}}[(d_{t+1}-\mu)^2]}_{=\sigma^2}+2(\mu-m_t)\underbrace{\mathbb{E}_t^{\text{obj}}[d_{t+1}-\mu]}_{=0}+(m_t-\mu)^2\Big] \\ &= \omega b_t+\omega\left[\sigma^2+(m_t-\mu)^2\right].\end{aligned}$$

Substituting these objective expectations into equation (A.28) gives equation (39). To derive the subjective expectations of the one-period return note first that

$$\begin{aligned}\mathbb{E}_t^{\text{subj}}[d_{t+1}] &= \mu_t^i, \\ \mathbb{E}_t^{\text{subj}}[m_{t+1}] &= \omega m_t+(1-\omega)\mu_t^i, \\ \mathbb{E}_t^{\text{subj}}[b_{t+1}] &= \omega b_t+\omega\mathbb{E}_t^{\text{subj}}\left[(d_{t+1}-\mu_t^i)^2\right],\end{aligned}$$

where B-agents set $\mu_t^B=m_t$, while ambiguity-averse A-agents set $\mu_t^A=m_t-\kappa s_t$, with, for large t, $s_t^2=(1-\omega)^2 b_t$. The subjective expectation of $(d_{t+1}-\mu_t^i)^2$ can be determined by integration over the subjective, predictive dividend density (B.3). The conditional expectation for given precision is

$$\mathbb{E}_t^{\text{subj}}\left[(d_{t+1}-\mu_t^i)^2|\phi\right] = \frac{(n_t+1)}{n_t}\frac{1}{\phi}.$$

Because in the steady state, $t \to \infty$, $n_t \to \overline{n} = 1/(1-\omega)$, we have that

$$\lim_{t\to\infty} \mathbb{E}_t^{\text{subj}}\left[(d_{t+1} - \mu_t^i)^2|\phi\right] = \frac{2-\omega}{\phi}.$$

If we now assume that the truncation bounds $\underline{\phi}$ and $\overline{\phi}$ are wide and integrate this conditional expectation over the Gamma density of ϕ, we get

$$\mathbb{E}_t^{\text{subj}}\left[(d_{t+1} - \mu_t^i)^2\right] = \frac{(2-\omega)(1-\omega)}{2\omega - 1},$$

and, thus,

$$\mathbb{E}_t^{\text{subj}}[b_{t+1}] = \omega b_t + (1-\omega) b_t \frac{2-\omega}{2-\frac{1}{\omega}} \approx b_t$$

for ω sufficiently close 1. ∎

B Dampened Student-t distribution

In this appendix we derive the predictive distribution of dividends when the variance is unknown but bounded on a finite interval. We rely on early work by Bakshi and Skoulakis (2010) to show that such a distribution takes the form of a "dampened" Student-t which we formally define below.

Suppose the precision $\phi = 1/\sigma^2$ has a truncated Gamma distributed with support $[\underline{\phi}, \overline{\phi}]$, $\underline{\phi} > 0$, shape parameter b and ν degrees of freedom is. The density of the truncated Gamma is given by

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

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

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

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*

$$x \sim \mathrm{t}_\nu^D[\underline{\phi}, \overline{\phi}],$$

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).

Because ϕ has finite support and is bound away from 0, the density $f(x)$ in Definition B.1 are thinner than those of a Student-t distribution. Bakshi and Skoulakis (2010) show that the moment generating function of the dampened is finite, and, thus, all its moments exist and are finite. For $\underline{\phi} \to 0$ and $\overline{\phi} \to \infty$, the distribution $f(x)$ converges to 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**). *Suppose the random variable y is such that*

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

then, y has a non-standardized dampened t-distribution

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

with ν degrees of freedom, mean m, shape b, variance scale parameter v^2, and truncation bounds $\underline{\phi}$, $\overline{\phi}$. Furthermore, 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 $\frac{b}{\nu}\overline{\phi}$, that is,

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

In the model of Section 3.2, agents of generation t need to compute expected utility of their wealth, a quantity that depends on their their belief about the dividend d_{t+1}. As discussed in the main text, agents form beliefs after observing the state variables m_t, b_t, n_t, and ν_t, where n_t denotes the number of dividend realization and ν_t the precision of the ϕ-prior.[26]

Lemma B.1 (Predictive dividend distribution with unknown variance). *Consider a subjective Normal/truncated-Gamma prior for μ and ϕ with parameters μ_t^i, b_t, n_t, and ν_t. The predictive distribution of d_{t+1} is a non-standardized dampened Student-t, as per Definition B.2, that is*

$$d_{t+1}|\mu_t^i, b_t, n_t, \nu_t \sim^i \mathrm{t}_{\nu_t}^D\left[\mu_t^i, b_t, \frac{n_t+1}{n_t}; \underline{\phi}, \overline{\phi}\right]. \tag{B.3}$$

Proof: Conditional on the precision ϕ, the prior of μ, $p(\mu|\phi, \mu_t^i, n_t)$ is normal. We first prove that, conditional on the precision ϕ, the predictive density of d_{t+1}, is also normal. To see this, note that,

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

[26] Note that in the main text we use $n_t = t$ and $\nu_t = t$. In this appendix we use the more general notation n_t and ν_t that will apply also for the case of perpetual learning, introduced in Section 3.4.

Hence, $d_{t+1}|\phi, \mu_t^i, n_t \sim \mathcal{N}\left(\mu_t^i, \frac{n_t+1}{n_t\phi}\right)$.

The unconditional density of d_{t+1} can now be determined from the conditional density by integrating out the precision ϕ using the truncated inverse-Gamma distribution, that is,

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

The last equality shows that the standardized dividend $\frac{d_{t+1}-\mu_t^i}{\sqrt{\frac{n_t+1}{n_t}}\sqrt{\frac{b_t}{\nu_t}}}$ has a dampened-t distribution, as per Definition B.1, that is,

$$\left.\frac{d_{t+1}-\mu_t^i}{\sqrt{\frac{n_t+1}{n_t}}\sqrt{\frac{b_t}{\nu_t}}}\right| \mu_t^i, b_t, n_t, \nu_t \sim \mathrm{t}_{\nu_t}^D\left[\frac{b_t}{\nu_t}\underline{\phi}, \frac{b_t}{\nu_t}\overline{\phi}\right]$$

Therefore, the predictive distribution of the dividend d_{t+1} has a non-standardized dampened-t distribution, that is,

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

∎

Lemma B.1 shows 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

$$\mathbb{E}_t^i[u(W_{t+1}^i)] = \int_{-\infty}^{\infty} \mathbb{E}_t^i[u(W_{t+1}^i)|\phi]p(\phi|b_t, \nu_t)d\phi, \tag{B.4}$$

where $\mathbb{E}_t^i[]$ refers to the expectation relative to agent i's prior about μ, that is $\mu_t^A = m_t - \kappa s_t$ and $\mu_t^B = m_t$, and $p(\phi|b_t, \nu_t)$ denotes the density of the truncated Gamma distribution stated in equation (B.1). The expected utility conditional on ϕ in equation (B.4) is obtained by integrating over the conditionally normal density of the μ prior μ_t^i.

B.1 Expected utility when variance is unknown

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 $\mathcal{P}_t$.

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

$$p_t = \frac{1}{r} m_t - \Lambda(b_t), \tag{B.5}$$

and let $\Delta\mu_t^i = m_t - \mu_t^i$, with $\mu_t^B = m_t$ and $\mu_t^i = m_t - \kappa s_t$ and s_t an estimate of the standard error of the mean m_t. Then the expected utility of wealth for each agent i is given by

$$\mathbb{E}_t^i[u(W_{t+1}^i)] = \frac{1}{C(b_t,\nu_t;\underline{\phi},\overline{\phi})} \int_{\underline{\phi}}^{\overline{\phi}} \mathbb{E}_t^i[u(W_{t+1}^i)|\phi,m_t,n_t,\nu_t]\phi^{\frac{\nu_t}{2}-1}e^{-\phi\frac{b_t}{2}}d\phi, \tag{B.6}$$

where

$$\begin{aligned}\mathbb{E}_t^i[u(W_{t+1}^i)|\phi,m_t,n_t,\nu_t] &= -\frac{1}{\gamma}e^{-\gamma\left((1+r)(W_t^i+\theta\Lambda(b_t))-\theta\left(1+\frac{1}{rn_{t+1}}\right)\Delta\mu_t^i\right)} \\ &\times \sqrt{\frac{n_t\phi}{(n_t+1)2\pi}}\int_{-\infty}^{\infty} e^{-\gamma\theta\left(\left(1+\frac{1}{rn_{t+1}}\right)e_{t+1}^i-\Lambda(b_{t+1})\right)-\frac{1}{2}\frac{n_t\phi}{(n_t+1)}(e_{t+1}^i)^2}de_{t+1}^i,\end{aligned} \tag{B.7}$$

with $e_{t+1}^i = d_{t+1}-\mu_t^i$, $b_{t+1} = b_t + \frac{n_t}{n_{t+1}}(\Delta\mu_t^i - e_{t+1}^i)^2$, $n_{t+1} = n_t+1$, and $\nu_{t+1} = \nu_t+1$. Hence under the conjectured price (B.5), the expected utility of agent i is independent of m_t.

Proof: Integrating the time $t+1$ utility over the conditional density of d_{t+1} and using the budget constraint in equation (24), we obtain

$$\begin{aligned}\mathbb{E}_t^i[u(W_{t+1})|\phi,m_t,n_t,\nu_t] &= -\frac{1}{\gamma}\sqrt{\frac{n_t\phi}{(n_t+1)2\pi}}\int_{-\infty}^{\infty} e^{-\gamma W_{t+1}^i}e^{-\frac{1}{2}\frac{n_t\phi}{(n_t+1)}(d_{t+1}-\mu_t^i)^2}dd_{t+1} \\ &= -\frac{1}{\gamma}e^{-\gamma(1+r)(W_t^i-\theta p_t)}\sqrt{\frac{n_t\phi}{(n_t+1)2\pi}}\int_{-\infty}^{\infty} e^{-\gamma\theta(d_{t+1}+p_{t+1})}e^{-\frac{1}{2}\frac{n_t\phi}{(n_t+1)}(d_{t+1}-\mu_t^i)^2}dd_{t+1}.\end{aligned} \tag{B.8}$$

Substituting the linear price function from equation (B.5) into equation (B.8) we obtain equation (B.7). As equation (B.7) shows, the expected utility does not depend on the estimate of the dividend mean, m_t. ∎

The boundedness of the 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.6) is well defined.

C Perpetual learning about mean and variance: Technical details

In this appendix we describe how we model perpetual learning. The main idea consists of allowing for "leakage" in the information content that is passed on between generations. 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). We assume that when these posteriors are handed over to generation t, some information is lost. We model this information leakage as shocks that add noise to these posteriors before they are updated by generation t upon observation of the dividend d_t.

Updating of the mean μ. The generation $t-1$ posterior of μ with n_{t-1} observations is normal, conditional on the precision ϕ, that is,

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

We model information leakage in this posterior as an additive normal shock with mean 0 and variance

$$\left(\frac{1}{\omega}-1\right)\frac{1}{n_{t-1}\phi}, \quad \omega \in (0,1].$$

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

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

This noisy posterior is then updated with the information contained in the dividend d_t. The resulting 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 in the last equality we used

$$
m_t = m_{t-1} + \frac{1}{n_t}e_t, \tag{C.2}
$$

$$
n_t = \omega n_{t-1} + 1 \tag{C.3}
$$

$$
e_t = d_t - m_{t-1}. \tag{C.4}
$$

Therefore, the post-leakage μ-posterior that is inherited by generation t as their prior is

$$
\mu|\phi \sim \mathcal{N}\left(m_t, \frac{1}{n_t\phi}\right), \tag{C.5}
$$

where the updated value of the state variables m_t and n_t from equations (C.2) and (C.3) incorporate the information leakage ω.

Updating of the precision ϕ. The generation $t-1$ posterior of ϕ with n_{t-1} observations is truncated gamma with ν_{t-1} degrees of freedom, that is,

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

We model information leakage in this posterior as multiplicative shock

$$\frac{1}{\omega}\eta_{t-1} \tag{C.6}$$

where η_{t-1} has a generalized-Beta distribution, (see Bakshi and Skoulakis, 2010, equation (35)). Using Theorem 1 in Bakshi and Skoulakis (2010) we obtain that the post-leakage ϕ-prior that is passed to generation t is again a Gamma distribution truncated at the same bounds, that is

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

To update this prior upon observation of the the time t dividend, we must first determine the distribution of d_t conditional of ϕ. Using the post leakage μ-prior derived in equation (C.1) we obtain

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

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

Updating with d_t leads to the posterior of ϕ that is handed over to the agents of generation t in the form of the updated state variables b_t and ν_t. Formally,

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

Hence the post-leakage ϕ-posterior that is inherited by generation t as their prior is

$$\phi|d_t, b_{t-1}, m_{t-1}, n_{t-1}, \nu_{t-1}; \underline{\phi}, \overline{\phi} \sim \text{TG}\left[\frac{\nu_t}{2}, \frac{b_t}{2}; \underline{\phi}, \overline{\phi}\right],$$

where

$$\begin{aligned}
b_t &= \omega b_{t-1} + \omega \frac{n_{t-1}}{n_t} e_t^2, && \text{(C.7)} \\
n_t &= \omega n_{t-1} + 1, && \text{(C.8)} \\
\nu_t &= \omega \nu_{t-1} + 1, && \text{(C.9)} \\
e_t &= d_t - m_{t-1}, && \text{(C.10)}
\end{aligned}$$

the updated value of the state variables b_t, n_t, and ν_t from equations (C.7), (C.8), and (C.9) incorporate the information leakage ω. When $\omega < 1$ and t is large, the effective number of observations n_t and the degrees of freedom ν_t converge to the same upper limit

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

Therefore, by choosing the effective number of observations $\overline{n}$ we can determine the amount of information leakage ω in the model, that is,

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

D Construction of the equilibrium with perpetual learning

In this appendix, we describe the numerical procedure we use to construct the equilibrium of an OLG economy with heterogeneous agents and perpetual learning, described in Section 3.4.

Both types of generation-t agents know the state variables m_t, b_t, n_t and ν_t from the "information processing" step described in Section 3.2. Generation-t agents anticipate that information gets lost as generations overlap. We construct an equilibrium that satisfies rational expectations, that is, the pricing function perceived by agents corresponds to the true pricing function, even if, because of learning, agents' beliefs about return distribution do not coincide with the true distribution. This implies that generation-t agents must anticipate the demand of generation-$t+1$ agents and the future price function 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 a non-standardized dampened t-distribution (see Definition B.2)

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

Therefore, the individual dividend surprise $e_{t+1}^i = d_{t+1} - \mu_t^i$ is also a non-standardized dampened-t. 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}. Expressing the dynamics of the state variables (n_t, m_t, b_t, ν_t) from equations (32), (35), (36), and (37) in terms of e_{t+1}^i we obtain,

$$n_{t+1} \quad = \quad \omega n_t + 1, \tag{D.1}$$

$$m_{t+1} \quad = \quad \frac{\omega n_t}{\omega n_t + 1} m_t + \frac{1}{\omega n_t + 1} d_{t+1} \quad = \quad m_t - \frac{1}{n_{t+1}}\Delta\mu_t^i + \frac{1}{n_{t+1}} e_{t+1}^i, \tag{D.2}$$

$$b_{t+1} \quad = \quad \omega b_t + \frac{\omega n_t}{\omega n_t + 1}(m_t - d_{t+1})^2 \quad = \quad \omega b_t + \frac{\omega n_t}{n_{t+1}}(\Delta\mu_t^i - e_{t+1}^i)^2, \tag{D.3}$$

$$\nu_{t+1} \quad = \quad \omega \nu_t + 1. \tag{D.4}$$

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 $\tau \to \infty$, using equation (D.2) we can write p_{t+1} as follows,

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

Under this conjecture, the budget constraint is

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

Hence under the conjectured equilibrium price, the budget constraint only depends on $\Lambda(b_t)$ and $\Lambda(b_{t+1})$. The expected utility of agents i is then

$$\mathbb{E}^i[u(W_{t+1}^i(\theta))] = \frac{1}{C(b_t^i, \overline{n}; \underline{\phi}, \overline{\phi})}\int_{\underline{\phi}}^{\overline{\phi}} \mathbb{E}^i(u(W_{t+1}^i(\theta)|\phi, b_t)\phi^{\frac{\overline{n}}{2}-1}e^{-\phi\frac{b_t}{2}}\,d\phi,$$

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

with $b_{t+1} = \omega\left(b_t + (\Delta\mu_t^i - e_{t+1}^i)^2\right)$. Since e_{t+1}^i is dampened-t distributed, $\Lambda(b_t) \geq 0$, $\lim_{b_t\to\infty}\Lambda(b_t) < \infty$ and expected utility is well defined.

The first-order condition with respect to the portfolio weight θ is

$$\begin{aligned} \frac{d\mathbb{E}^i[u(W_{t+1}^i(\theta))|b_t]}{d\theta} &= \frac{1}{C(b_t^i, \overline{n}; \underline{\phi}, \overline{\phi})} \\ &\times \int_{\underline{\phi}}^{\overline{\phi}} \frac{d\mathbb{E}(u(W_{t+1}^i(\theta))|\phi, b_t)}{d\theta}\phi^{\frac{\overline{n}}{2}-1}e^{-\phi\frac{b_t}{2}}\,d\phi, \end{aligned}$$

with

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

We determine the function $\Lambda(b_t)$ as a fixed point through value function iteration. When both agents invest at a given b_t, they take $\Lambda(b_t)$ as given and optimize their holding θ_t^i 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, $\theta^B(b_t) + \theta^A(b_t) = 1$.

$$
\begin{aligned}
\mathbb{E}_t(p_{t+1} + d_{t+1} - p_t) &= \mathbb{E}_t(d_{t+1}) + \mathbb{E}_t\left[\frac{1}{r}m_{t+1} - \Lambda(b_{t+1})\right] - \left[\frac{1}{r}m_t - \Lambda(b_t)\right] \\
&\approx \mathbb{E}_t(d_{t+1}) + \frac{1}{r}\mathbb{E}_t[m_{t+1} - m_t] - \mathbb{E}_t^[b_{t+1} - b_t]\frac{d\Lambda}{db}.
\end{aligned} \tag{D.5}
$$

$$
\begin{aligned}
\mathbb{E}_t^{\text{obj}}[d_{t+1}] &= \mu, \\
\mathbb{E}_t^{\text{obj}}[m_{t+1}] &= \omega m_t + (1-\omega)\mu, \\
\mathbb{E}_t^{\text{obj}}[b_{t+1}] &= \omega b_t + \omega\mathbb{E}_t^{\text{obj}}\Big[(d_{t+1} - m_t)^2\Big] \\
&= \omega b_t + \mathbb{E}_t^{\text{obj}}\Big[(d_{t+1} - \mu + \mu - m_t)^2\Big] \\
&= \omega b_t + \omega\Big[\underbrace{\mathbb{E}_t^{\text{obj}}[(d_{t+1}-\mu)^2]}_{=\sigma^2} + 2(\mu - m_t)\underbrace{\mathbb{E}_t^{\text{obj}}[d_{t+1} - \mu]}_{=0} + (m_t - \mu)^2\Big] \\
&= \omega b_t + \omega\Big[\sigma^2 + (m_t - \mu)^2\Big].
\end{aligned}
$$

$$
\begin{aligned}
\mathbb{E}_t^{\text{subj}}[d_{t+1}] &= \mu_t^i, \\
\mathbb{E}_t^{\text{subj}}[m_{t+1}] &= \omega m_t + (1-\omega)\mu_t^i, \\
\mathbb{E}_t^{\text{subj}}[b_{t+1}] &= \omega b_t + \omega \mathbb{E}_t^{\text{subj}}\Big[(d_{t+1}-\mu_t^i)^2\Big].
\end{aligned}
$$

$$
\begin{aligned}
\mathbb{E}_t^{\text{subj}}\Big[(d_{t+1}-\mu_t^i)^2|\phi\Big] &= \frac{(n_t+1)}{n_t}\frac{1}{\phi}, \\
\implies \lim_{t\to\infty} \mathbb{E}_t^{\text{subj}}\Big[(d_{t+1}-\mu_t^i)^2|\phi\Big] &= (2-\omega)\frac{1}{\phi}.
\end{aligned}
$$

$$
\mathbb{E}_t^{\text{subj}}\Big[(d_{t+1}-\mu_t^i)^2\Big] = \frac{(2-\omega)(1-\omega)}{2\omega-1},
$$

$$
\mathbb{E}_t^{\text{subj}}[b_{t+1}] = \omega b_t + (1-\omega) b_t \frac{2-\omega}{2-\frac{1}{\omega}} \approx b_t
$$

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.[27]

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.2) shows that dividends 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}. 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.

[27]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.

Chen, J., Hong, H., and Stein, J. C. (2001). Forecasting crashes: trading volume, past returns, and conditional skewness in stock prices. *Journal of Financial Economics*, 37:345–381.

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*.

French, K. R., Schwert, W., and Stambaugh, R. (1987). Expected stock returns and volatility. *Journal of Financial Economics*, 19:3–29.

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.

Robins, R. P. and Smith, G. P. (2023). A new look at expected stock returns and volatility. *Critical Finance Review*, 12(1-4):225–270.

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.

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

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.