DISCUSSION PAPER SERIES

DP16441

Dynamics of Asset Demands with Confidence Heterogeneity

Adrian Buss, Raman Uppal and Grigory Vilkov

FINANCIAL ECONOMICS

CEPR

ISSN 0265-8003

Dynamics of Asset Demands with Confidence Heterogeneity

Adrian Buss, Raman Uppal and Grigory Vilkov

Discussion Paper DP16441
Published 09 August 2021
Submitted 06 August 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

Dynamics of Asset Demands with Confidence Heterogeneity

Abstract

To understand the dynamics of investor asset demands, we develop a multiperiod general-equilibrium model driven by a single latent variable, differences in beliefs, resulting from heterogeneity in investors' confidence regarding the return dynamics of assets. Consistent with the data, investors' asset holdings are concentrated and display large and persistent heterogeneity in asset demands across investors. Moreover, demand curves are steeper than with homogeneous beliefs. The time-series and cross-sectional variation in assets' realized and expected returns, as well as their volatilities, are driven by the mean and dispersion of latent demand.

JEL Classification: D53, G11, G12

Keywords: Institutional asset demand, asset-demand elasticity, investors' expectations, trend chasing, predictability

Adrian Buss - adrian.buss@insead.edu
INSEAD and CEPR

Raman Uppal - ruppal@mac.com
Edhec Business School and CEPR

Grigory Vilkov - vilkov@vilkov.net
Frankfurt School of Finance and Management

Dynamics of Asset Demands with Confidence Heterogeneity*

Adrian Buss Raman Uppal Grigory Vilkov

July 31, 2021

Abstract

To understand the dynamics of investor asset demands, we develop a multiperiod general-equilibrium model driven by a single latent variable, differences in beliefs, resulting from heterogeneity in investors' confidence regarding the return dynamics of assets. Consistent with the data, investors' asset holdings are concentrated and display large and persistent heterogeneity in asset demands across investors. Moreover, demand curves are steeper than with homogeneous beliefs. The time-series and cross-sectional variation in assets' realized and expected returns, as well as their volatilities, are driven by the mean and dispersion of latent demand.

JEL: D53 - Financial Markets, G11 - Portfolio Choice, G12 - Asset Pricing.
Keywords: Institutional asset demand, asset-demand elasticity, investors' expectations, trend chasing, predictability.

*Adrian Buss is affiliated with INSEAD and CEPR, Boulevard de Constance, 77305 Fontainebleau, France, email: adrian.buss@insead.edu; Raman Uppal is affiliated with Edhec Business School and CEPR, 10 Fleet Place, London, U.K. EC4M 7RB, email: raman.uppal@edhec.edu; and Grigory Vilkov is affiliated with Frankfurt School of Finance & Management, Adickesallee 32-34, 60322 Frankfurt, Germany, email: vilkov@vilkov.net.
This paper was previously circulated under the title "Investor Confidence and Portfolio Dynamics." We are grateful for helpful comments from Harjoat Bhamra, Oliver Boguth, Laurent Calvet, Bruce Carlin, Georgy Chabakauri, João Cocco, Victor DeMiguel, Bernard Dumas, Lorenzo Garlappi, Francisco Gomes, Bernard Herskovic, Christian Heyerdahl-Larsen, Stavros Panageas, Anna Pavlova, Kim Peijnenburg, Joel Peress, Valery Polkovnichenko, Lukas Roth, Nikolai Roussanov, Dimitri Vayanos, Gyuri Venter, Pietro Veronesi, Lu Zhang, and Irina Zviadadze, and seminar participants at Bank of Lithuania, EDHEC Business School, Einaudi Institute for Economics and Finance, EPFL, Frankfurt School of Finance & Management, Goethe University Frankfurt, HEC Lausanne, HEC Paris, Temple University, University of Birmingham, Vienna Graduate School of Finance, and at the following conferences: Adam Smith Workshop, American Finance Association Meetings, 1st Asset Pricing Conference at Collegio Carlo Alberto, 3rd Annual CEPR Spring Symposium, China International Conference in Finance, European Finance Association Meetings, 2nd World Symposium on Investment Research, Financial Intermediation Research Society (FIRS) Conference, Seventh Biennial McGill Global Asset Management Conference, and UBC Summer Finance Conference.

Traditional asset-pricing models imply asset demand with little heterogeneity across investors. However, Koijen and Yogo (2019) show empirically that investors' asset demand is characterized by large heterogeneity across different institutional investors. Consequently, the aggregate demand elasticity varies across assets that are held by different investors. They also find that a substantial part of this heterogeneity cannot be explained by assets' observable characteristics but is driven by latent variables. This latent demand, which is highly volatile, strongly influences assets' realized and expected returns as well as their volatilities.

Our main contribution is to develop a novel but simple theoretical framework of a multi-period general-equilibrium economy that can jointly rationalize the *static* and *dynamic* properties of asset demands. Our model is driven by a single latent variable: differences in beliefs arising from heterogeneity in investors' confidence about expected asset returns. This model generates large and persistent heterogeneity in investors' demand, both in the cross-section of assets and over time, with micro and macro demand elasticities that are lower than in the standard models with full confidence.

To study heterogeneity in asset demands across investors and over time, we develop a model that features two classes of risky assets (along with a risk-free asset) and two groups of investors. The two classes of assets differ in terms of the confidence investors have about their expected returns. The first class of assets represents assets well understood by all investors, for instance, large-cap stocks. The second class of assets represents assets whose fundamentals are more difficult to assess, such as small-cap stocks, stocks from new sectors such as technology and biotechnology, or a less-traditional asset class, such as private equity and venture capital. The two *groups* of investors have identical Epstein and Zin (1989) preferences but differ in terms of the confidence they have about each asset class. Both groups of investors are assumed to be fully confident about the first asset class (large-cap stocks). But, for the second class of assets (small-cap stocks), one group of investors representing, for instance, large institutions and "smart-money," is more confident about its estimate of expected returns on this asset than the other group of investors representing, for instance, smaller institutions or households. Because regulations restrict many institutional investors from taking short positions, we restrict all investors to have non-negative asset demands, which implies that financial markets are incomplete.

Many empirically observed patterns in investors' static and dynamic asset demands arise naturally in our model, which we describe below along with the economic mechanisms driving them.

First, heterogeneity in investors' confidence regarding the return dynamics of assets, coupled with Bayesian learning, leads to pronounced *heterogeneity* in asset demands across investors. Less-confident investors, typically smaller financial institutions, hold large-cap stocks, while more-confident investors, typically larger financial institutions, invest a greater share of wealth in the riskier asset. This heterogeneity in asset demands is present even when investors have the same initial beliefs, as long as they have heterogeneous confidence. This is because differences in confidence about the less-familiar asset's return, coupled with Bayesian updating, gives rise to precautionary-savings (Kimball, 1990) and intertemporal-hedging demands (Merton, 1971) in anticipation of future changes in beliefs.[1]

Second, despite Bayesian learning, there is strong *persistence* in the heterogeneity in asset demands across investors. This is because, under Bayesian updating, while confidence increases with each new cash-flow observation, the increase in confidence becomes smaller each successive period, and hence, changes in beliefs also become smaller with each additional observation. Thus, the precautionary-savings and intertemporal hedging demands decline only slowly, which explains the persistence in the heterogeneity of asset demands and the importance of the latent variable even in the long run.

Third, the *demand elasticity* of the less-familiar asset with respect to its own price is lower in the model with confidence heterogeneity, compared to the standard model where all investors have full confidence. The macro elasticity, that is the change in the demand for the bond with respect to the price of the risky assets, is also lower in the presence of confidence heterogeneity. Latent demand (that is, demand driven by differences in beliefs) plays an important role in reducing these micro and macro demand elasticities.

[1] Specifically, the higher perceived consumption volatility, which results from the higher perceived cash-flow volatility of the less-familiar risky asset, creates a precautionary motive to hold more wealth in the safe asset. The underinvestment in the less-familiar asset arises because less-confident investors wish to hedge changes in their perceived investment-opportunity set. In particular, if the dividend-growth rate of the less-familiar asset decreases, their future utility falls. To hedge this decline, they hold a portfolio that underweights the less-familiar asset (i.e., they have a negative intertemporal hedging demand).

Fourth, both classes of investors change their portfolios conditional on observing shocks to dividends, but in different ways. Less-confident investors display *trend-chasing* behavior; that is, they increase their holdings of the less-familiar asset following positive cash-flow news and vice versa. On the other hand, more-confident investors take on riskier positions when expected market returns are higher; that is, they exhibit superior *market-timing* abilities.

Fifth, the static and dynamic differences in asset demands across the two groups of investors lead to *heterogeneity in portfolio returns.* Less-confident investors earn lower investment returns than more-confident ones—in absolute terms and also on a risk-adjusted basis. In particular, their larger allocation to the safe asset leads to lower absolute expected portfolio returns and also to lower portfolio volatility. Moreover, their portfolio underdiversification—resulting from an underinvestment in the asset about which they are less familiar—reduces their risk-adjusted investment returns (i.e., the portfolio Sharpe ratio). Heterogeneity in absolute investment returns is further boosted by the market-timing ability of the more-confident investors.

Finally, our model highlights the *importance of latent demand for asset prices and returns.* In equilibrium, the less-familiar asset's return volatility, risk premium, and Sharpe ratio are endogenously higher than those of the more-familiar asset (despite the fundamentals for the two assets being the same) and decline only slowly over time, as investors' confidence about its dividend-growth rate increases. The high return volatility stems from an amplification of the cash-flows shocks,[2] which, together with the high volatility of the stochastic discount factor, also leads to a high risk premium. Intuitively, the higher Sharpe ratio for the less-familiar asset is required to induce the more-confident investors to hold a larger share of this asset—to compensate for the low demand from less-confident investors.

Importantly, our results are not a consequence of behavioral biases or investment mistakes nor do not they rely on the possibility of extreme beliefs or fat-tailed (belief) distributions. Instead, they are a consequence of small revisions in beliefs having a large impact on the long-run subjective consumption distribution. The only critical assumption underlying our findings is that investors prefer early resolution of uncertainty. Indeed, we explicitly show that models with log utility (i.e., in the absence of intertemporal hedging demands) and CRRA utility (i.e.,

[2]With a preference for early resolution of uncertainty, less-confident investors allocate a larger (smaller) fraction of their risky portfolio to this asset following positive (negative) cash-flow news about the less-familiar asset. This, in turn, increases (decreases) its price-dividend ratio exactly when its dividends are high (low).

with investors being neutral about the timing of the resolution of uncertainty) fail to deliver the portfolio and investment-return dynamics observed empirically. We also demonstrate that our results are distinctly different from those of many alternative models relying either on heterogeneous preferences or constant belief heterogeneity.

Our model is motivated by the broad features of asset demands and returns documented in Koijen and Yogo (2019), whose estimates are based on a single-period model in which investors have logarithmic utility function, heterogeneous beliefs, and face a shortsale constraint. Just like Koijen and Yogo (2019), investors in our model have heterogeneous beliefs (because of differences in confidence) and face a short-sale constraint. However, to study the *dynamics* of assets demands and returns, we consider a multiperiod setting (with Bayesian updating of beliefs). We also assume that investors have Epstein and Zin (1989) preferences, which allows us to study the importance of intertemporal hedging demands that are absent under log utility. Therefore, the results that we obtain for dynamically optimal portfolios nicely complement the analysis of myopic portfolios in Koijen and Yogo (2019).

Our work is also related to Collin-Dufresne, Johannes, and Lochstoer (2016a) and Ehling, Graniero, and Heyerdahl-Larsen (2018), who study models with a single risky asset with investors who have heterogeneous beliefs resulting from investors' experience and focus on the link between investor demographics (generations) and asset returns. The key difference with these papers is that in our framework learning is Bayesian, and our focus is on the evolution of *heterogeneity* in asset demands for multiple risky assets across investors, rather than the asset-pricing implications of biased learning.[3]

There is also an extensive literature that studies the implications of time-invariant heterogeneous beliefs for asset pricing; see, for instance, Scheinkman and Xiong (2003), Panageas (2005), Gallmeyer and Hollifield (2008), Xiong and Yan (2010), Prieto (2013), Chabakauri (2013, 2015), Baker, Hollifield, and Osambela (2016), and Borovička (2020). Panageas (2019) provides an excellent review of the literature on belief and preference heterogeneity.

[3]Bayesian learning also features in the representative-agent single-risky-asset models of Collin-Dufresne, Johannes, and Lochstoer (2016b) and Johannes, Lochstoer, and Mou (2016). In contrast, our model has multiple groups of investors (thereby, allowing for heterogeneity in confidence) and multiple risky assets. Empirically, Johannes et al. (2016) demonstrate that beliefs about long-run dynamics are volatile and drift considerably over time—consistent with our model.

The rest of the paper is organized as follows. Section 1 introduces our economic framework. We show how heterogeneity in investors' confidence affects the dynamics of asset demands in Section 2, the dynamics of asset returns in Section 3, and investors' portfolio returns in Section 4. Section 5 contrasts our predictions with those from alternative model specifications. Section 6 concludes. Technical details are relegated to the appendices.

1 Economic Framework

This section introduces our economic framework, which is designed to capture a key salient feature of financial markets—differences in investor confidence. Below, we describe the details of the model, investors' optimization problems, and the definition of equilibrium.

1.1 The Model

We study a general-equilibrium model set in discrete time, with time interval Δt and a finite horizon T.[4]

Financial assets: There are three financial assets in the economy. The first asset is a risk-free single-period discount bond in zero net supply, indexed by $n = 0$. In addition, there are two risky assets, indexed by $n \in \{1, 2\}$, each in unit supply, and modeled as a claim to a Lucas (1978) tree. Specifically, we assume that each asset's log dividend growth $\Delta d_{n,t+1} \equiv \ln\big[D_{n,t+1}/D_{n,t}\big]$ is described by an IID-Normal model with expected dividend-growth rate μ_n and dividend-growth volatility σ_n:

$$\Delta d_{n,t+1} = \mu_n + \sigma_n \, \varepsilon_{n,t+1}, \quad n \in \{1, 2\}, \tag{1}$$

where $\varepsilon_{n,t+1} \sim \mathcal{N}(0, 1)$, and $\varepsilon_{1,t+1}$ and $\varepsilon_{2,t+1}$ are assumed to be uncorrelated.[5] We interpret the first risky asset as a portfolio of large-cap stocks that are well understood by all investors. The second risky asset represents a portfolio of smaller stocks that are less well understood (for

[4]The choice of a finite horizon is dictated by the numerical solution technique. However, in our numerical illustration, we choose a very long horizon ($T = 1000$) that renders the impact of the finite horizon—even quantitatively—negligible (in that further extending the horizon has no impact on the results).

[5]This assumption simplifies Bayesian' updating because it implies that investors will only use news about the cash flow of the less-familiar asset in their learning. However, it is not crucial to our results, which are—qualitatively—unchanged if one allows for correlated dividends (cf. Appendix D). Note also that even though dividends are uncorrelated, asset returns, which depend on equilibrium prices, will be endogenously correlated.

example, because of lower analyst coverage or because they are stocks in a new sector, such as technology or biotechnology).

Investors: The economy is populated by two groups of investors, indexed by $k \in \{1, 2\}$. Investors have Epstein and Zin (1989) and Weil (1990) preferences over consumption of the single consumption good, $C_{k,t}$. Specifically, lifetime utility $V_{k,t}$ is defined recursively as

$$V_{k,t} = \left[(1-\beta)\, C_{k,t}^{1-\frac{1}{\psi}} + \beta\, E_t^k \left[V_{k,t+1}^{1-\gamma} \right]^{\frac{1}{\phi}} \right]^{\frac{\phi}{1-\gamma}}, \tag{2}$$

where E_t^k denotes the time-t conditional expectation under investor k's subjective probability measure, $\beta > 0$ is the rate of time preference, $\gamma > 0$ is the coefficient of relative risk aversion, $\psi > 0$ is the elasticity of intertemporal substitution (EIS), and $\phi = \frac{1-\gamma}{1-1/\psi}$.

The two groups of investors differ only in their level of confidence about the cash-flow dynamics of the less-familiar asset.[6] Intuitively, such heterogeneity could arise from differences in experience, skill, or access to analyst coverage. In particular, we assume that the more confident investors ($k = 1$) represent the larger financial institutions whose latent demand is relatively more stable (Koijen and Yogo, 2019). In contrast, the less confident investors ($k = 2$) represent smaller financial institutions and households.

We model investors' beliefs using parameter uncertainty combined with Bayesian updating (based on cash-flow news), which yields rich dynamics in investors' subjective beliefs regarding the distribution of future cash flows of the less-familiar stocks.[7] In particular, investor k starts at date $t = 1$ with a conjugate prior for the expected dividend-growth rate of the second asset, $\mu_2 \sim \mathcal{N}(\hat{\mu}_{k,1}, A_{k,1}\, \sigma_2^2)$. This prior, combined with the dividend dynamics in (1), implies a time-t posterior density function $p(\mu_2 | \Delta d_{2,1}, \ldots, \Delta d_{2,t}) = \mathcal{N}(\hat{\mu}_{k,t}, A_{k,t}\, \sigma_2^2)$, with the dynamics of $\hat{\mu}_{k,t}$ and $A_{k,t}$ given by

$$\hat{\mu}_{k,t} = \hat{\mu}_{k,t-1} + \left(\Delta d_{2,t} - \hat{\mu}_{k,t-1} \right) \frac{A_{k,t-1}}{1 + A_{k,t-1}}, \tag{3}$$

[6] It would be straightforward to also allow for heterogeneity in investors' preferences. We assume identical preferences to highlight the effects arising from differences in confidence alone. The case of heterogeneity in preferences alone is studied in Section 5.2.

[7] Investors "agree to disagree," that is, they do not revise their beliefs based on asset prices. Morris (1995) explains why it is reasonable to assume that investors have different priors and that this is fully consistent with rationality. There exists a vast literature that uses this formulation; see, for example, Dumas, Kurshev, and Uppal (2009), and the papers cited therein.

$$A_{k,t} = \frac{1}{1/A_{k,t-1} + 1}. \tag{4}$$

Consequently, even though the dividend dynamics of the less-familiar asset are driven by an IID model with constant parameters, from the investors' perspective the expected dividend-growth rate, $\hat{\mu}_{k,t}$, and its volatility, $\sqrt{1 + A_{k,t}}\,\sigma_2$, are time varying, with this time-variation driven by $A_{k,t}$. In particular, any difference in investors' initial precision, $A_{k,1}$, leads to (long-term) differences in investors' confidence (and beliefs) and volatility of latent demand.

For expositional ease, we assume that there is no uncertainty about the parameters of the dividend process for the first asset (i.e., the portfolio of large-cap stocks), and investors know that its correlation with the dividends of the less-familiar asset is zero.

There are a few important points to highlight about the dynamics of investors' beliefs in Equations (3) and (4) that will be useful for understanding the dynamics of asset demands. First, the perceived dividend-growth rate, $\hat{\mu}_{k,t}$, is a martingale, so revisions in subjective beliefs constitute permanent shocks to investors' continuation utility $V_{k,t}$. Second, the volatility of the expected dividend-growth rate, $\sqrt{1 + A_{k,t}}\,\sigma_2$, is decreasing over time, but at a decreasing rate. Specifically, using the recursive definition in (4), it is straightforward to express the date-t value of $A_{k,t}$ in terms of $A_{k,1}$, its date-1 value: $A_{k,t} = \frac{A_{k,1}}{1+t\,A_{k,1}}$. This implies that the change

$$A_{k,t+1} - A_{k,t} = \frac{-A_{k,1}^2}{\big(1 + t\,A_{k,1}\big)\big(1 + (t+1)\,A_{k,1}\big)} \tag{5}$$

is negative and, hence, $A_{k,t}$ is decreasing over time. Third, the effect of cash-flow news on the change in perceived beliefs, $\hat{\mu}_{k,t+1} - \hat{\mu}_{k,t}$, is decreasing over time, but at a decreasing rate. To see this, note that the recursive definition in (4) implies that the (last) term appearing in (3) that determines by how much beliefs change after each cash-flow news, simplifies to $\frac{A_{k,t-1}}{1+A_{k,t-1}} = \frac{A_{k,1}}{1+t\,A_{k,1}}$. The change in this expression is negative

$$\frac{A_{k,t+1}}{1 + A_{k,t+1}} - \frac{A_{k,t}}{1 + A_{k,t}} = \frac{-A_{k,1}^2}{\big(1 + (t+1)\,A_{k,1}\big)\big(1 + (t+2)\,A_{k,1}\big)}, \tag{6}$$

which confirms that the perceived dividend-growth rate diminishes over time.

Figure 1 illustrates these patterns in the perceived dividend-growth rate of the less-familiar asset, $\hat{\mu}_{k,t}$, for a simulated path of the economy. Two cases are illustrated: in the "Low

Figure 1: Perceived Dividend-Growth Rate of the Less-familiar Asset

The figure illustrates the interplay between investors' confidence and Bayesian updating in our model. It depicts the perceived (annualized) dividend-growth rate of the less-familiar asset for a simulated path of the economy. Two cases are illustrated: In the "Low confidence" case, prior precision is given by $A_{k,1} = 20$, while in the "High confidence" case, prior precision is given by $A_{k,1} = 40$. The beliefs in both cases are initially unbiased, i.e., the perceived growth rate at date 1, $\hat{\mu}_{k,1}$, is equal to the true growth rate $\mu_2 = 1.8\%$ p.a. The time interval Δt is set equal to one quarter and volatility σ_2 to 4.80% p.a. The shaded gray areas indicate periods of higher-than-expected (realized) dividend growth for the less-familiar asset, and the dashed horizontal line shows the true expected growth rate.

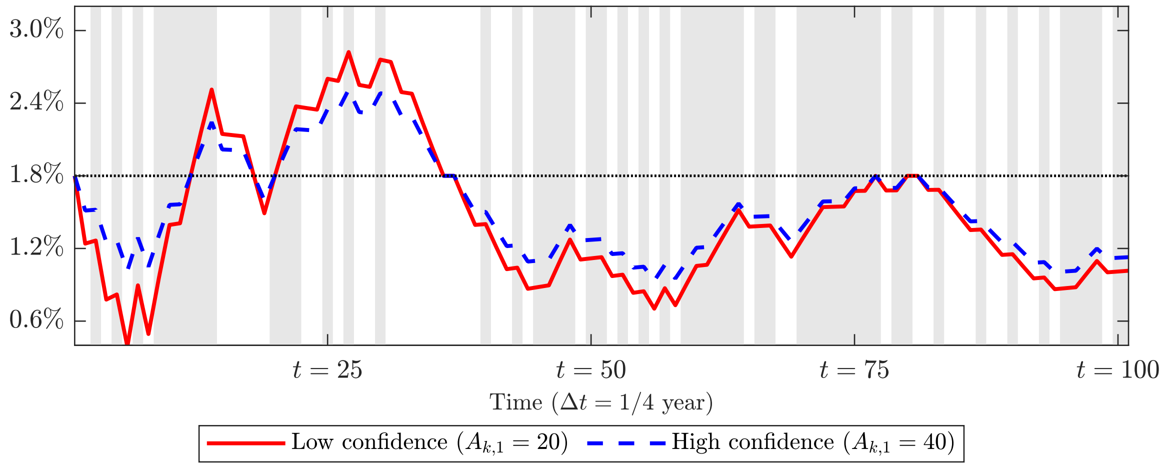

confidence" case, prior precision is given by $A_{k,1} = 20$, while in the "High confidence" case, it is given by $A_{k,1} = 40$. The beliefs in both cases are initially unbiased, that is, the (annualized) perceived growth rate at date 1, $\hat{\mu}_{k,1}$, is equal to the true growth rate $\mu_2 = 1.8\%$ p.a. Intuitively, when realized dividend growth is higher than expected, highlighted in shaded gray, investors revise their beliefs about μ_2 upward and vice versa for lower-than-expected growth. As a result, investors are sometimes optimistic and, at other times, pessimistic regarding the future dividend growth of the less-familiar asset; that is, they overstate (understate) its dividend-growth rate (compared to the true growth rate μ_2 highlighted by the dotted horizontal line). Notably, in the case where investors are less confident about the dividend dynamics of the less-familiar asset ($A_{k,1} = 20$), the revisions in expectations in response to dividend news are greater. Moreover, with each new observation, investors' estimates of the expected dividend-growth rate become more precise, and, thus, revisions in their beliefs decrease in size.

1.2 Investors' Optimization Problem and Equilibrium

The objective of investors in group k is to maximize their expected lifetime utility (2), by choosing consumption, $C_{k,t}$, and holdings in financial assets, $\theta_{k,n,t}, n \in \{0,1,2\}$, subject to the budget equation

$$C_{k,t} + \theta_{k,0,t}\, S_{0,t} + \sum_{n=1}^{2} \left(\theta_{k,n,t} - \theta_{k,n,t-1}\right)\, S_{n,t} \;\leq\; \theta_{k,0,t-1} + \sum_{n=1}^{2} \theta_{k,n,t-1}\, D_{n,t}, \tag{7}$$

where $S_{n,t}$ denotes the price of asset n. The left-hand side of the budget equation describes the use of funds for consumption, the purchase or sale of the (newly issued) short-term discount bond, and changes in the demand for the risky assets, while the right-hand side reflects the source of funds, stemming from the unit payoff of the (maturing) short-term bond as well as the dividends from the holdings of the risky assets. In addition, we impose a short-sale constraint for the second risky asset: $\theta_{k,2,t} \geq 0, k \in \{1,2\}$.[8]

The first-order conditions for consumption and demand for the risk-free and the two risky assets, which are derived in Appendix A, imply:

$$\begin{aligned} M_{k,t+1} &= \beta\, \xi_{k,t}\, \exp\Big(\left(-1/\psi\right) \Delta c_{k,t+1} - \left(\gamma - 1/\psi\right) v_{k,t+1}\Big), \\ S_{0,t} &= E_t^k \left[M_{k,t+1}\right], \\ S_{1,t} &= E_t^k \left[M_{k,t+1} \left(S_{1,t+1} + D_{1,t+1}\right)\right], \\ S_{2,t} &= \frac{1}{1 - \Lambda_{k,t}}\, E_t^k \left[M_{k,t+1} \left(S_{2,t+1} + D_{2,t+1}\right)\right], \end{aligned} \tag{8}$$

where $M_{k,t+1}$ is investor k's stochastic discount factor, $\Delta c_{k,t+1}$ is log consumption growth, $v_{k,t+1} \equiv \log(V_{k,t})$ is log-continuation utility, $\Lambda_{k,t}$ is the Lagrange multiplier associated with the short-sale constraint, and $\xi_{k,t} \equiv E_t^k \left[\exp\left((1-\gamma)\, v_{k,t+1}\right)\right]^{(\gamma - 1/\psi)/(1-\gamma)}$.

Equation (8) highlights that when relative risk aversion is not equal to the reciprocal of EIS ($\gamma \neq 1/\psi$), then shocks to the future log-continuation utility, $v_{k,t+1}$, are a source of priced risk—in addition to shocks to the one-period ahead log consumption growth, $\Delta c_{k,t+1}$. In particular,

[8] We constrain short sales because, in practice, it is not easy to short a risky asset, especially one with which investors are not very familiar; moreover, many financial institutions are prohibited from holding short positions. Qualitatively, the constraint does not affect any of our results; quantitatively, the results are slightly stronger if we allow for short sales (see Section 5.3). For the first asset, the constraint is never binding (because there is no disagreement) and, hence, is omitted.

through its impact on investors' continuation utility, variation in the perceived dividend-growth rate of the less-familiar asset ($\hat{\mu}_{k,t}$) becomes a priced risk factor.

Equilibrium in the economy is defined by consumption policies $\{C_{k,t}\}$, asset-demand decisions $\{\theta_{k,n,t}\}$, and price processes for the financial assets $\{S_{n,t}\}$, with $k \in \{1,2\}, n \in \{0,1,2\}$, such that: (a) $C_{k,t}$ and $\theta_{k,n,t}$ maximize investor k's expected lifetime utility (2) subject to the budget equation (7) and the short-sale constraint, $\theta_{k,2,t} \geq 0$; (b) aggregate demand equals aggregate supply:

$$\sum_{k=1}^{2} \theta_{k,0,t} = 0, \quad \text{and} \qquad \sum_{k=1}^{2} \theta_{k,n,t} = 1, \quad n \in \{1,2\}. \tag{9}$$

The state variables of the economy are: the consumption share of the second group of investors, $\omega_{2,t} \in (0,1)$; the dividend share of the first risky security $\delta_{1,t} \in (0,1)$, whose dynamics follow from the joint dividend dynamics in (1); the expected dividend-growth rate of the less-familiar asset as perceived by the two groups of investors, $\hat{\mu}_{k,t}$, $k \in \{1,2\}$, with its dynamics specified in (3); and the (deterministic) posterior variances of the beliefs of the two groups of investors, $A_{k,t}\,\sigma_2^2$, $k \in \{1,2\}$, with the dynamics specified in (4).

1.3 Parameter Values and Solution Method

To illustrate our key results, we solve the model for the set of parameter values listed in Table 1. Importantly, however, our results depend only on heterogeneity in confidence and the assumption that investors prefer early resolution of uncertainty, that is, $\psi > 1/\gamma$.[9] In Section 5 and Appendix D, we discuss the quantitative impact of variations in parameter values and explicitly show that our results are not sensitive to the choice of parameter values.

We set the trading frequency, Δt, to one quarter. The number of periods, T, is set to $1,000$—minimizing any effects from having a finite horizon. The initial share of the second (less-familiar) asset's dividends, $\delta_{2,1}$, is set to a conservative value of 0.20 to highlight the "small-cap" nature of the stocks. We specify identical processes for the two assets' dividends, calibrated to match jointly the historical mean and volatility of aggregate U.S. consumption

[9] This is a standard parametric assumption in the macroeconomics and asset-pricing literature that ensures that positive dividend shocks for an asset increase its price-dividend ratio; Bansal and Yaron (2004) and Ai, Bansal, Guo, and Yaron (2019) provide supporting empirical evidence.

Table 1: Model Parameters

The table reports the baseline parameter values used for our numerical illustrations.

Variable	Description	Baseline
Δt	Trading frequency	1/4 year
T	Total number of trading dates (quarters)	1,000
β	Rate of time preference (per quarter)	0.994
γ	Relative risk aversion	10
ψ	Elasticity of intertemporal substitution	1.0
$w_{2,1}$	Initial wealth share of the less-confident investors	2/3
μ_n	Expected dividend growth (per quarter)	0.45%
σ_n	Dividend growth volatility (per quarter)	2.40%
ρ	Correlation between dividend growth rates	0.0
$\delta_{2,1}$	Less-familiar asset's share of total initial dividends	0.20
λ	Leverage factor	2.5
$\hat{\mu}_{2,1}$	Initial mean of less-confident investors' prior distribution	0.45%
$A_{2,1}$	Initial precision of less-confident investors' prior distribution	20
$\underline{\mu}_2, \bar{\mu}_2$	Truncation boundaries for beliefs of less-confident investors	$[-0.55\%, 1.45\%]$

growth (1.80% and 3.70% p.a., respectively).[10] Because equity claims are generally levered, we report asset and investment return moments assuming a leverage factor of $\lambda = 2.5$.

We use a (quarterly) rate of time-preference $\beta = 0.994$, a coefficient of relative risk aversion $\gamma = 10$, and an EIS $\psi = 1$—usual choices in the literature. For ease of exposition, we assume that the first group of investors knows the dividend-growth rate of the less-familiar asset (i.e., has infinite precision).[11,12] In contrast, the second group is uncertain about its growth rate but learns from dividend realizations. Accordingly, we refer to the first group ($k = 1$) of investors, representing larger financial institutions, as "more-confident investors" and the second group ($k = 2$), representing smaller financial institutions and households, as "less-confident investors." We assume that the first group of investors is initially endowed with 1/3 of the total wealth

[10]Specifying identical processes for the dividend dynamics of the two assets guarantees a stable dividend-share distribution in the initial years, thereby eliminating effects arising mechanically from time variation in the dividend-share distribution. Note, however, that in the limit ($t \to \infty$), a bimodal distribution with dividend shares of zero and one arises—as is standard for such models; see, e.g., Cochrane, Longstaff, and Santa-Clara (2008). We have confirmed that our results remain unchanged for a stationary dividend-share distribution.

[11]Consequently, in this case, we have four state variables: the consumption share of the less-confident investors, $\omega_{2,t}$; the dividend share of the first risky asset, $\delta_{1,t}$; the less-familiar asset's dividend-growth rate as perceived by less-confident investors, $\hat{\mu}_{2,t}$; and the posterior variance of their beliefs, $A_{2,t}\,\sigma_2^2$.

[12]The assumption that one group has infinite precision is not crucial for our results; we only require that the beliefs of one group of investors are less precise than those of the other. The model could easily be extended to incorporate generalizations such as parameter uncertainty for both investors, parameter uncertainty for both risky assets, and uncertainty about the assets' dividend-growth volatilities.

(Koijen and Yogo (2019, Table 4) report that the thirty largest institutions hold about 30% of the market), while the less-confident investors are endowed with 2/3 of the total wealth.[13]

Specifically, for the second group of agents we set their initial prior to be $A_{2,1} = 20$ (equivalent to 20 quarters of data) and $\hat{\mu}_{2,1} = \mu_2$ (i.e., beliefs are initially unbiased). To make sure that our results are not driven by extreme (and potentially unreasonable) levels of the perceived dividend-growth rate, we use fairly tight truncation bounds of $\underline{\mu}_2 = -0.55\%$ and $\bar{\mu}_2 = 1.45\%$ (i.e., deviations of plus/minus 1% from the true mean).[14]

Identifying the equilibrium is a non-trivial task. In particular, we extend the numerical solution approach in Dumas and Lyasoff (2012) along several dimensions; for instance, we incorporate parameter uncertainty with Bayesian learning, multiple risky assets, and Epstein-Zin-Weil preferences.[15] The details of our solution approach are provided in Appendix B.

We illustrate our results via plots, in which we focus on the first 100 quarters (25 years) of the economy and report averages across $100,000$ simulated paths of the economy. To illustrate how the *distribution* of the key quantities is affected by confidence heterogeneity, we frequently also plot the first and third quartile of the simulated paths (in light dash-dotted lines). In addition to showing the case of heterogeneous confidence, the plots include the case where both groups of investors have full confidence, which serves as a benchmark.

2 Dynamics of Asset Demands

In this section, we explain how heterogeneity in investor confidence gives rise to distinct asset-demand patterns as observed in the data, including persistent heterogeneity in asset demands, a reduction in the price sensitivity of asset demands, and trend-chasing behavior.

[13] In Section 5.3, we demonstrate how our findings vary quantitatively with the initial wealth distribution.

[14] Formally, we use a truncated Normal prior for μ_2, which results in a truncated Normal posterior with the same truncation bounds. Conveniently, the updating equations for the hyperparameters, $\hat{\mu}_{k,t}$ and $A_{k,t}$, remain the same—although $\hat{\mu}_{k,t}$, in general, no longer corresponds to the subjective conditional mean of the less-familiar asset's dividend growth (see the online appendix of Collin-Dufresne et al. 2016b). Note also that, for EIS greater (smaller) than one, truncation is generally required to ensure the existence of equilibrium as a positive (negative) probability for an arbitrarily high $\hat{\mu}_{k,t}$ leads to a violation of the transversality condition. In Section 5.3, we demonstrate how our findings vary quantitatively with the truncation bounds.

[15] For instance, for our baseline setting, we solve more than one million (small) equation systems, each involving the interpolation of future prices, future holdings, and future value functions over the grid of state variables, which takes about 24 hours on a 24-core workstation.

Figure 2: Asset-Demand Dynamics

The figure illustrates the portfolio shares of less-confident investors (left column) and more-confident investors (right column) over time. Panels A and B plot the average proportion of wealth invested in the risk-free bond, Panels C and D in the less-familiar asset representing small caps, and Panels E and F in the more-familiar asset representing large caps. Averages are calculated across 100,000 simulation paths. The light dash-dotted lines plot the first and third quartile of the distribution of the simulated paths. All graphs are based on the parameter values described in Table 1.

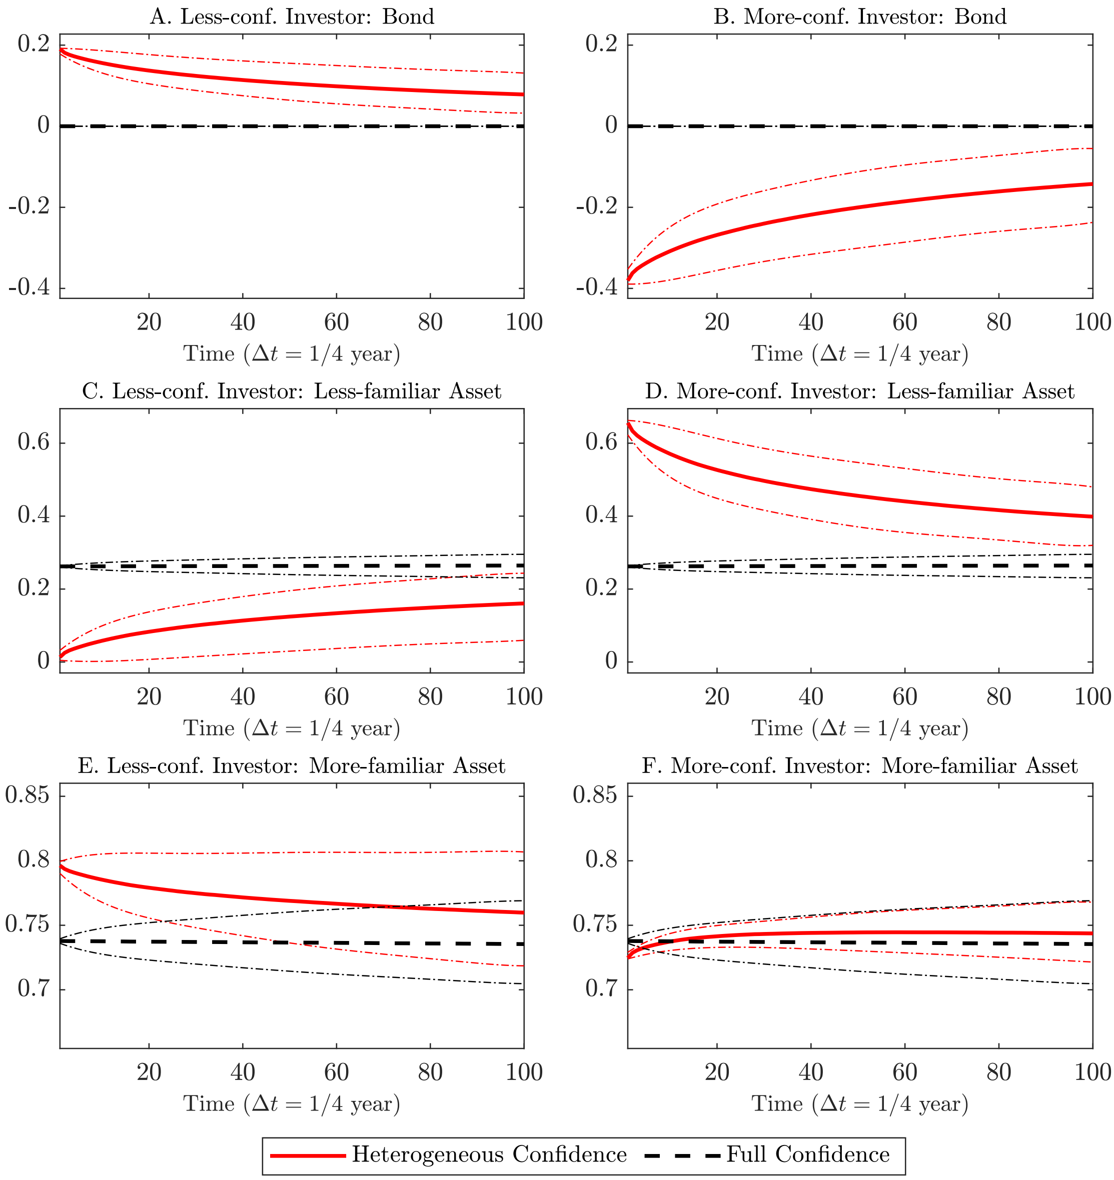

2.1 Heterogeneity in Asset-Demand Dynamics

Figure 2 plots the proportion of wealth each investor group allocates to the risk-free asset and the two risky assets—averaged across all simulation paths. First, we observe from Panels A and B that the model can generate a large and persistent heterogeneity in investors' asset demands (as in the data). Less-confident investors persistently allocate a larger fraction of their wealth to the risk-free asset. Indeed, in equilibrium, they have, on average, a long position in the safe asset. This is a consequence of their higher perceived consumption-growth volatility (resulting from the higher perceived dividend-growth volatility of the less-familiar asset), which gives rise to a precautionary-savings demand for the risk-free asset (as long as relative risk aversion $\gamma > 1$). In contrast, more-confident investors have, because of market-clearing, a short position in the safe asset.

Second, the risky-asset components of investors' portfolios also differ markedly. Specifically, less-confident investors persistently underweight the less-familiar small-cap stocks (Panel C) and overweight the more-familiar large cap stocks (Panel E)—compared to more-confident investors and also the full-confidence market weights. Market clearing implies that more-confident investors are overinvested in the less-familiar asset (Panel D)—relative to less-confident investors and relative also to the setting where all investors are fully confident. The more-confident investors' short position in the bond (Panel B), however, allows them to maintain their investment in the more familiar large-cap stocks at about the same level as the full-confidence case (Panel F). Thus, confidence heterogeneity has only a smaller impact on the diversification of their portfolios. As a result, consistent with empirical findings (Koijen and Yogo, 2019), larger financial institutions are better diversified, while smaller financial institutions and households are concentrated in more-familiar stocks.

The key economic force explaining this heterogeneity in investors' risky-asset demands is the *intertemporal-hedging demand*, which has been highlighted in the work of Merton (1971) on dynamic portfolio choice and of Campbell and Viceira (2002) on strategic asset allocation. Specifically, less-confident investors have a strong *negative* intertemporal-hedging demand for the less-familiar asset, which pushes down their overall demand for the asset. We show this explicitly in Figure 3 by decomposing investors' holdings in the less-familiar asset into two components: (1) a myopic component and (2) an intertemporal-hedging component. Observe

Figure 3: Less-familiar Asset: Portfolio-Share Decomposition

The figure plots the (average) portfolio share allocated to the less-familiar asset as well as its decomposition into the myopic and intertemporal-hedging components—for less-confident investors (Panel A) and more-confident investors (Panel B). Averages are calculated across $100,000$ simulation paths. All graphs are based on the parameter values described in Table 1.

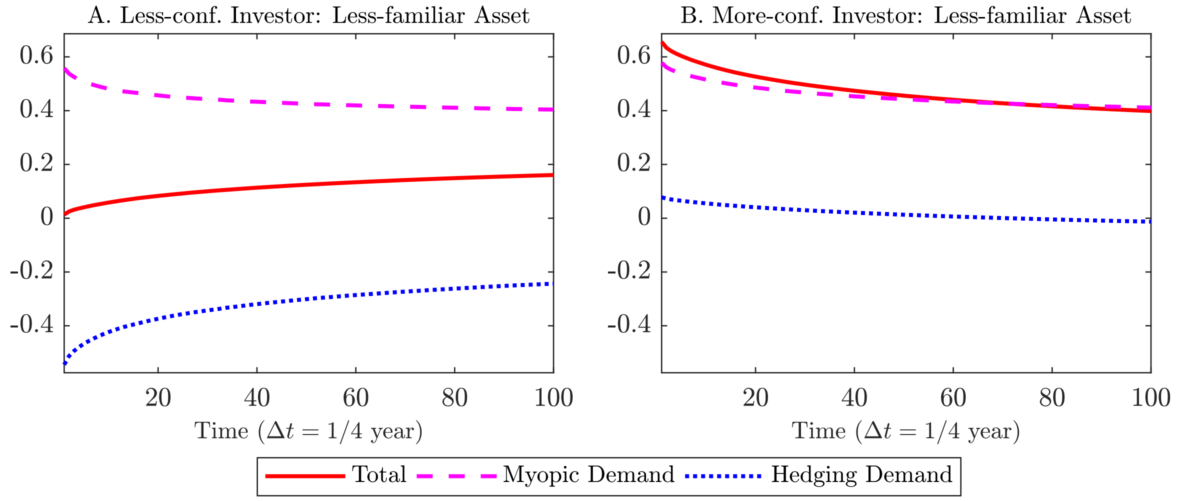

that the less-confident investors' myopic demand for the less-familiar asset is large and positive (Panel A) and very similar to the myopic demand of the more-confident investors (Panel B).[16] However, there is a striking difference in the intertemporal-hedging demands of the two groups of investors. While for less-confident investors the intertemporal-hedging demand is strongly negative, for more-confident investors it is quite small, arising mostly from (indirect) changes in their expected utility induced by changes in the beliefs of *less*-confident investors.[17]

To understand the intuition for the less-confident investors' negative intertemporal-hedging demand for the less-familiar asset, note that their continuation utility is highly sensitive to variations in their beliefs. Specifically, even though changes in beliefs are rather small on a per-period basis, their impact on the subjective continuation-utility is large because changes in beliefs permanently affect investors' (long-term) consumption dynamics. Formally, one can decompose shocks to each investor's subjective log-continuation utility (normalized by total

[16]The tiny difference in the myopic demands of the less- and more-confident investors arises from differences in the perceived dividend-growth volatility, with less-confident investors perceiving the less-familiar asset to be slightly riskier.

[17]More-confident investors hedge also changes in the other stochastic state variables, but this effect is even smaller (as can be seen from the sensitivities of their value function to the state variables, illustrated in Panel B of Figure 4 below).

Figure 4: Sensitivity of Continuation Utility to Changes in State Variables

Panels A and B, respectively, plot the average sensitivity of the log-continuation utility of less- and more-confident investors with respect to the three stochastic state variables. Averages are computed across 100,000 simulation paths. The graph is based on the parameter values described in Table 1.

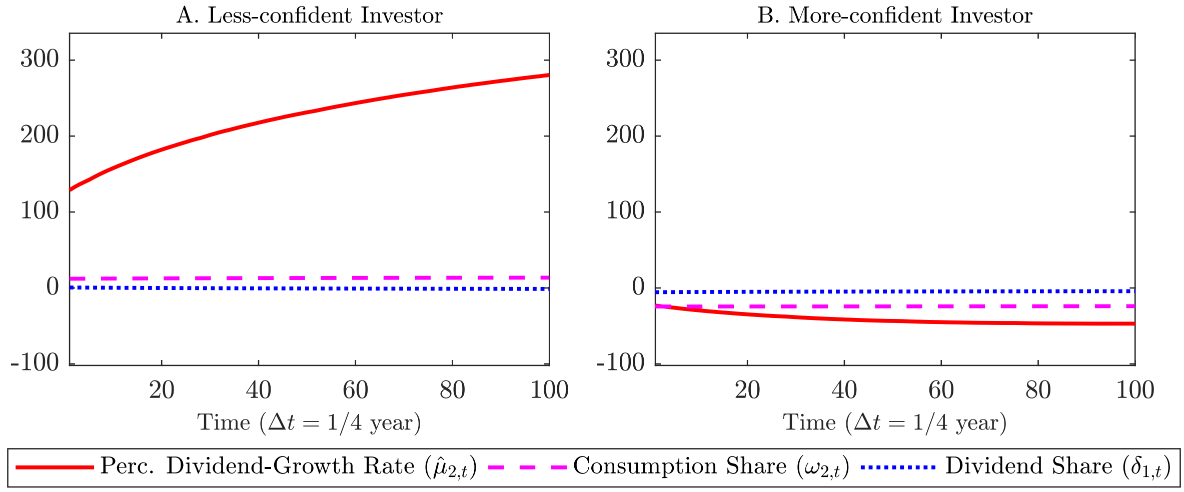

output), $\hat{v}_{k,t+1}$, into the sensitivity of their utility with respect to the stochastic state variables and the shocks to these state variables:

$$\hat{v}_{k,t+1} - E_t^k[\hat{v}_{k,t+1}] = \tag{10}$$

$$\left[\frac{\partial \hat{v}_{k,t}}{\partial \hat{\mu}_{2,t}} \left(\hat{\mu}_{2,t+1} - E_t^k[\hat{\mu}_{2,t+1}]\right) + \frac{\partial \hat{v}_{k,t}}{\partial \omega_{2,t}} \left(\omega_{2,t+1} - E_t^k[\omega_{2,t+1}]\right) + \frac{\partial \hat{v}_{k,t}}{\partial \delta_{1,t}} \left(\delta_{1,t+1} - E_t^k[\delta_{1,t+1}]\right)\right],$$

where $\hat{\mu}_{2,t}$, $\omega_{2,t}$, and $\delta_{1,t+1}$ denote, respectively, the less-familiar asset's dividend-growth rate as perceived by the less-confident investors, less-confident investors' share of aggregate consumption, and the more-familiar asset's share of total dividends. Panel A of Figure 4 shows that the sensitivity of the less-confident investors' subjective continuation utility with respect to the perceived dividend-growth rate ($\hat{\mu}_{2,k}$) is several orders of magnitude greater than the sensitivities to the other two stochastic state variables. Thus, it is the perceived dividend-growth rate that is the critical force driving their intertemporal-hedging demand. As expected, the sensitivity of more-confident investors' log-continuation utility to changes in beliefs of the less-confident investors is much smaller and of opposite sign (Panel B of Figure 4).

In particular, because the sensitivity of less-confident investors' subjective continuation utility with respect to the less-familiar asset's perceived dividend-growth rate is positive (and hence, their utility decreases when the perceived return of the less-familiar asset declines), they have an incentive to set up a portfolio that will perform well when the perceived dividend-growth rate is low. If asset prices rise (fall) following positive (negative) cash-flow news (as is the case with early resolution of uncertainty), then this can be achieved through a negative hedging position in the less-familiar asset, leading to a positive return when investors' utility is low.

As is apparent from Figure 2, asset-demand heterogeneity persists for long periods and weakens only gradually over time. These gradual changes are driven by less-confident investors' learning. Specifically, as their confidence increases, the perceived dividend-growth rate settles down. Thus, their precautionary-savings motive weakens and so the share of wealth allocated to the bond declines. Moreover, the revisions in less-confident investors' beliefs, and thus, the fluctuations in their continuation utility, become less pronounced, lowering the magnitude of the intertemporal-hedging component (Panel A of Figure 3). The reason why these changes occur slowly is that with each new observation of dividends, the increase in the precision of less-confident investors' beliefs becomes smaller, as explained in Equations (5) and (6).

To understand the role of sentiment in the asset demands of agents, we perform a simple simulation-based experiment, resembling a typical cross-sectional empirical test. First, we simulate a large number of model states, all at a fixed point in time, but with different histories, and record for each of these states (nodes in the tree) the portfolio holdings of each investor and the value of the latent variable (beliefs of the less-confident agent). Second, for both investors we compute the weight of each asset in the risky component of their portfolios, and regress it on the observed value weight of one of the assets in the economy, which is partly endogenous due to pricing, but is largely exogenous because of its dependence on the future dividends stream. Using the market weight of an asset as the only variable to explain the weight of the less-familiar asset delivers an R^2 of 32% for the portfolio of the investor with more confidence and 78% for the investor with less confidence. Adding the latent state variable (sentiment) to the right-hand side of the regression, however, boosts the R^2 values to 83% and 94%, respectively. Thus, the simulation experiment, while being a stylized exercise compared to the typical empirical test, shows that in a model with confidence heterogeneity, the latent variable explains asset demands, with this variable being more important for the investors who hold the less-familiar asset.

In summary, our model can simultaneously explain the large and persistent heterogeneity in investors' allocations to safe and risky assets documented in the empirical literature on the asset demands of institutional investors (Koijen and Yogo, 2019) and endowment funds (Lerner, Schoar, and Wang, 2008, Brown, Garlappi, and Tiu, 2010, Goetzmann and Oster, 2013). Moreover, the asset-demand patterns generated by our model are a consequence of a new channel that has not been studied before—limited confidence about the less-familiar asset's return, which—coupled with Bayesian updating—gives rise to precautionary-savings and intertemporal-hedging demands.

2.2 Price Sensitivity of Demand

For both investors, confidence heterogeneity has a considerable impact on the price sensitivity demand; that is, the ratio of the change in asset demand relative to the change in its price ($\partial Q/\partial P$), which is the inverse of the demand curve's slope. These sensitivities for the demands of the two assets by the two groups of investors are illustrated in Figure 5.

In particular, Panels A and B report the price sensitivity of investors' demand for the less-familiar asset; that is, how the demand for the less-familiar asset changes as its price goes up. In the presence of confidence heterogeneity, the demand of both investors for the less-familiar asset is *less* sensitive to changes in its price. In the case of the less-confident investor, this is again a consequence of the intertemporal hedging demand. Specifically, while the myopic component of the investor's demand remains quite sensitive to changes in its price, the hedging component (which is designed to hedge against adverse future outcomes rather than to benefit from a high expected return) is considerably less sensitive, thereby rendering the total demand less sensitive to changes in the less-familiar asset's price. Over time, as the less-confident investor becomes more confident, the importance of the hedging component declines, and thus, the demand elasticity goes up (in absolute terms). Interestingly, the elasticity of the more-confident investor is also lower than in the case of full confidence (although the effect is quantitatively smaller than that for the less-confident investor) and is driven by the higher return variance in the presence of confidence heterogeneity.[18] In unreported results, we find that the demand of the two investors for the *more*-familiar asset is also less sensitive to changes in the price of the *less*-familiar asset; that is, the cross-elasticities are lower with confidence heterogeneity.

[18]Note, even in the case of full confidence, the price elasticity of the two investors is not the same. This is a consequence of the different wealth (consumption) shares of the two investors.

Figure 5: Price Sensitivity of Investors' Demand

Panels A and B, respectively, plot the average price sensitivity of investors' demand for the less-familiar asset; that is, the ratio of the change in an investor's demand for the less-familiar asset and the change in its price. Panels C and D, respectively, plot the average price sensitivity of investors' demand for the more-familiar asset; that is, the ratio of the change in an investor's demand for the more-familiar asset and the change in its price. Panels E and F, respectively, plot the average ratio of a change in an investor's demand for the bond and the price of the "market portfolio" (i.e., the combination of both risky assets). Averages are computed across 100, 000 simulation paths. The graph is based on the parameter values described in Table 1.

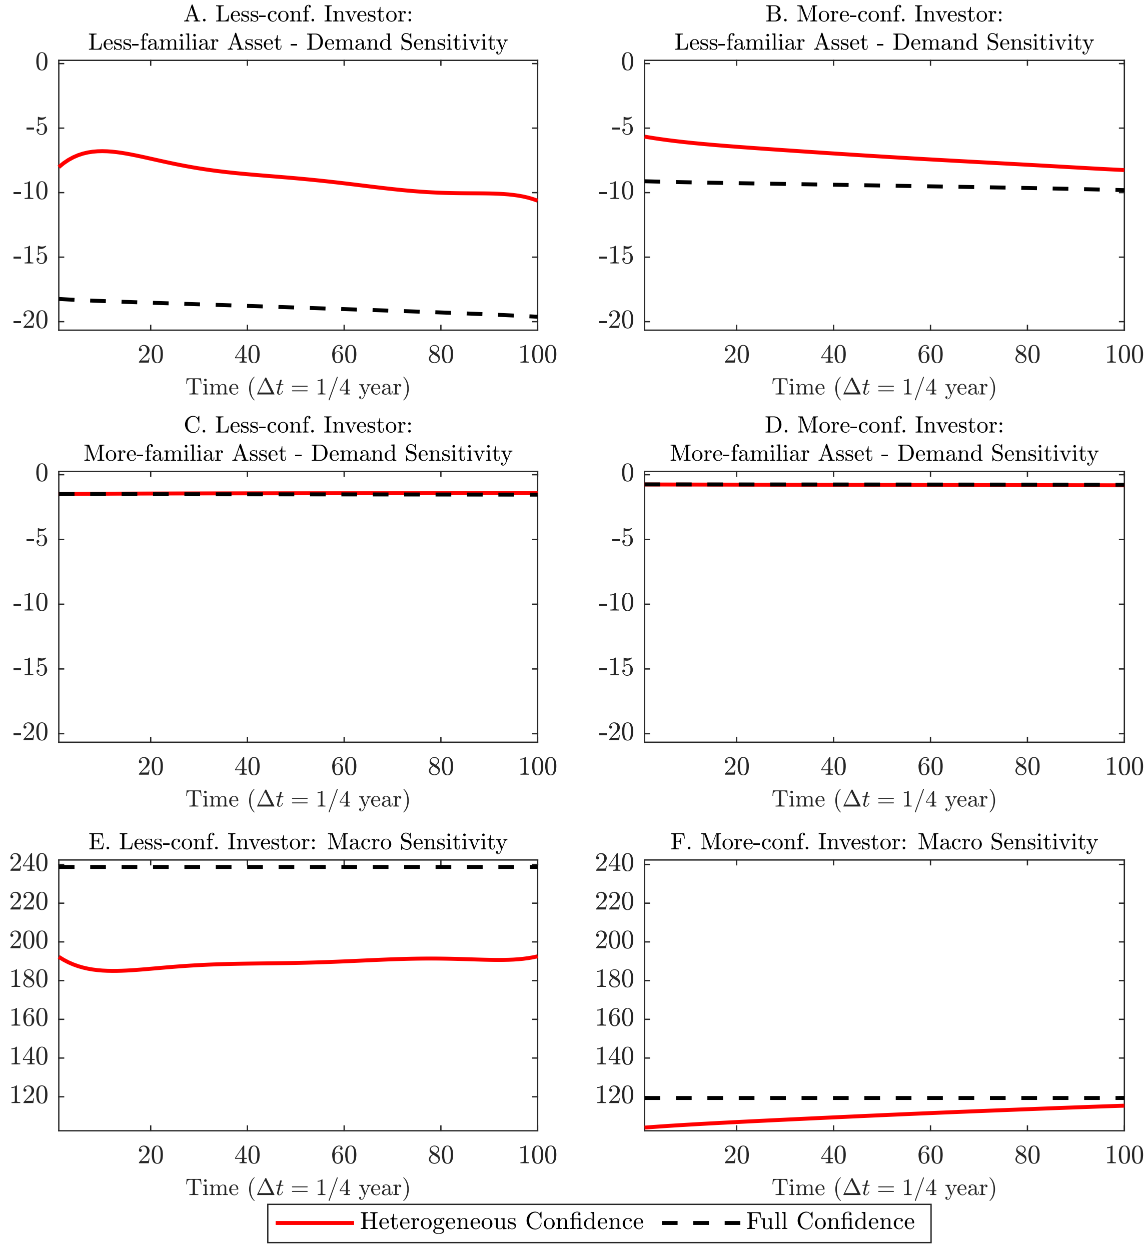

Panels C and D document that the price sensitivity of the more-familiar asset is considerably smaller than that of the less-familiar asset—due to the absence of differences in confidence for the dividend dynamics of the asset. Moreover, because the intertemporal hedging demand for this asset is small, confidence heterogeneity has only a negligible impact on the elasticity of investors' demand.

Finally, Panels E and F report "macro elasticities," that is, how the demand of the investors' for the bond changes as the price of the "market portfolio" (i.e., the sum of the two risky assets) goes up.[19] In the presence of confidence heterogeneity, the macro elasticities are again lower than the case with full confidence. This is again driven by the intertemporal hedging demand which implies smaller changes in the demand for the risky assets, and hence, for the bond, as the prices of the risky assets change.

To sum up, in the presence of confidence heterogeneity, both micro and macro demand elasticities are lower than in the case of full confidence. Hence, confidence heterogeneity can help to explain, at least partly, the low price elasticity of investors' demand documented empirically by Gabaix and Koijen (2021).[20]

2.3 Trend-Chasing

The revisions in less-confident investors' beliefs upon the arrival of cash-flow news also create "trend-chasing" in their demand for the less-familiar asset. For example, following positive cash-flow news for the less-familiar asset, they revise their subjective expectations regarding the dividend-growth rate of the less-familiar asset upward, and accordingly, increase their demand. Conversely, following negative cash-flow news, their demand declines. Hence, there is a positive correlation between past returns and less-confident investors' current demand, as illustrated in Panel A of Figure 6. This trend chasing is consistent with the empirical evidence for inexperienced fund managers (Greenwood and Nagel, 2009) and smaller endowment funds (Goetzmann and Oster, 2013).

Revisions in investors' beliefs can also explain why less-confident investors occasionally allocate substantial fractions of their wealth to assets that they are, in fact, not very confident about,

[19] Note, the higher absolute value of the macro elasticities is simply a result of the lower bond price.

[20] For an early paper with an excellent discussion of the slope of asset-demand curves, see Petajisto (2009).

Figure 6: Trend Chasing in the Demand for the Less-familiar Asset

Panel A plots the average change in the holdings of less-confident investors (in excess of the expected change) *conditional* on positive or negative cash-flow news for the less-familiar asset. Panel B depicts the proportion of wealth less-confident investors allocate to the less-familiar asset *conditional* on positive (negative) cash-flow news for the less-familiar asset in the first eight periods; namely, for the 10% of paths with the most positive (negative) cash-flow news. "Average" refers to the unconditional average across all paths. All quantities are based on 100,000 simulated paths of the economy and the parameter values described in Table 1.

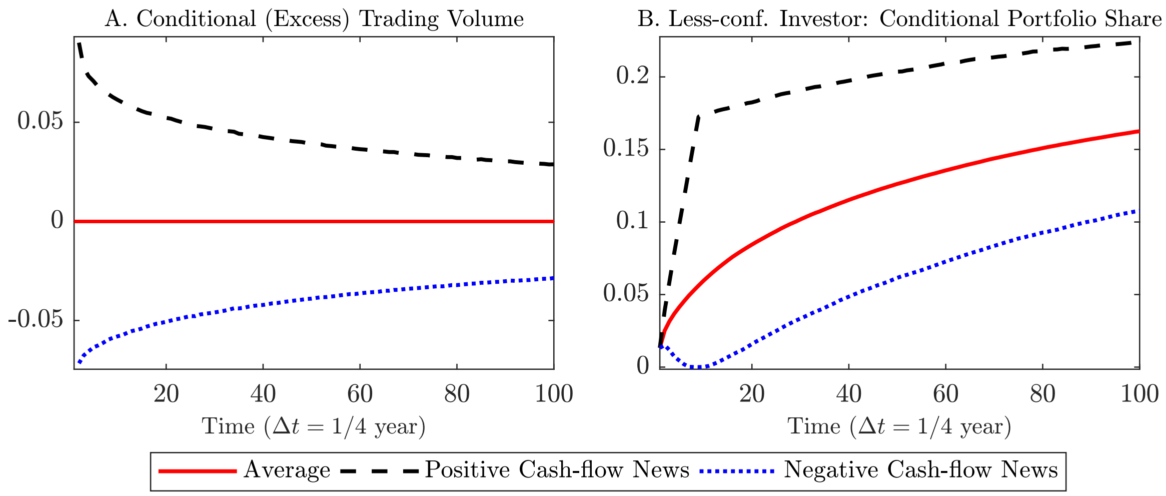

as was the case, for example, for dot-com stocks (Greenwood and Nagel, 2009). Intuitively, because less-confident investors' beliefs are not very precise in the early periods, shocks have a substantial impact on the perceived dividend-growth rate. Accordingly, less-confident investors can easily become overly optimistic about an asset, that is, overstate its dividend-growth rate, and consequently, also its return. As a result, they quickly increase their portfolio share in such assets. The average holdings in subsequent years will then increase even further, driven by the decline in the intertemporal-hedging component. We illustrate this in Panel B of Figure 6, which plots the less-familiar asset's share in less-confident investors' portfolios *conditional* on cash-flows news in the first few periods. In particular, the figure shows the portfolio share for the 10% of paths with the most positive cash-flow news and the 10% of paths with the most negative cash-flow news in the first two years (eight quarters).

3 Dynamics of Asset Returns

There is a large empirical literature that links investors' characteristics to their demand for assets (see Koijen and Yogo (2019) and the references therein), and the resulting consequences for asset returns. In the preceding section, we discussed how confidence heterogeneity drives asset demands. In this section, we investigate the implications of differences in investor confidence for the short- and long-run dynamics of equilibrium asset returns. In the next section, we will study the consequences of these asset returns for the heterogeneity in investment returns across the different classes of investors.[21]

3.1 Stochastic Discount Factor

To understand the forces driving asset returns in equilibrium, it is instructive to start by examining the impact of limited confidence on investors' stochastic discount factors (SDFs). Figure 7 illustrates the dynamics of the volatility of investors' SDFs. Notably, with heterogeneity in confidence, the volatilities of the SDFs of *both* groups of investors are higher than in the case of full confidence. Moreover, they are both generally declining over time.

Interestingly, the higher SDF volatilities for the two groups of investors are driven by different economic forces. In particular, note that shocks to the log SDF in Equation (8) can be written as

$$m_{k,t+1} - E_t^k[m_{k,t+1}] = -(1/\psi)\left(\Delta c_{k,t+1} - E_t^k[\Delta c_{k,t+1}]\right) - (\gamma - 1/\psi)\left(v_{k,t+1} - E_t^k[v_{k,t+1}]\right). \quad (11)$$

Hence, if relative risk aversion differs from the reciprocal of EIS, the SDF is driven not only by shocks to one-period-ahead consumption growth $\Delta c_{k,t+1}$ (as in standard CRRA-utility models) but also by shocks to the (forward-looking) subjective log-continuation utility $v_{k,t+1}$. As shown in Equation (10) above, the second term on the right-hand side of (11) can be further decomposed into (i) the sensitivity of investors' subjective log-continuation utility with respect to the stochastic state variables and (ii) shocks to these state variables.

For less-confident investors, the increase in their SDF's volatility in the first few periods (Panel A of Figure 7) comes from the continuation-utility component of the SDF. In particular,

[21] We do not include a discussion of asset *prices* in addition to our analysis of asset *returns*, because dividend-growth rates in the model are exogenous and IID, and therefore, price-dividend ratios are essentially the inverse of (long-term) expected returns.

Figure 7: Stochastic Discount Factors

Panels A and B show the average conditional volatility of the stochastic discount factor for the less- and more-confident investors, respectively. Averages are computed across 100,000 simulation paths. The light dash-dotted lines plot the first and third quartile of the distribution of the simulated paths. All graphs are based on the parameter values described in Table 1.

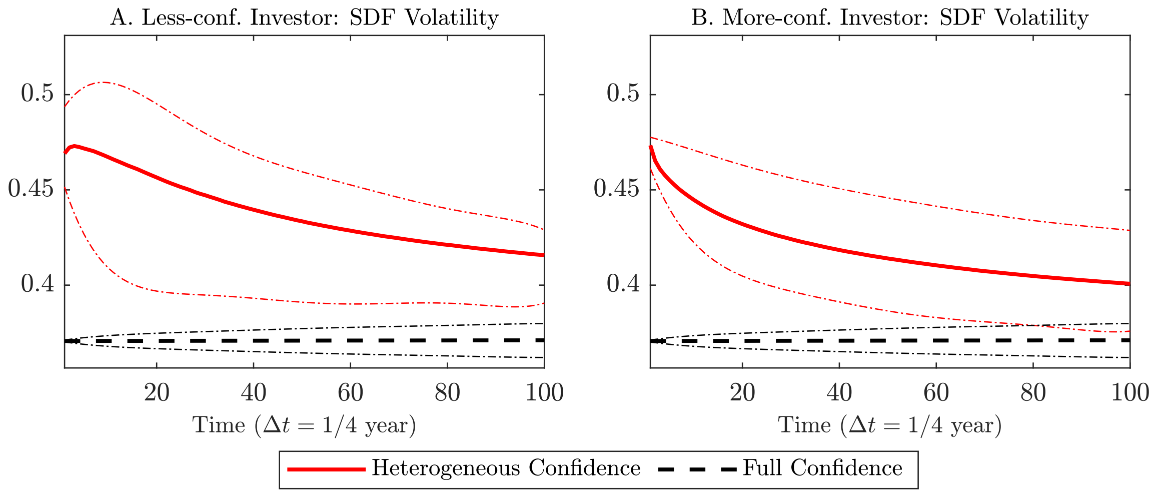

the high sensitivity of their continuation utility to the perceived dividend-growth rate of the less-familiar asset (cf. Panel A of Figure 4), coupled with the high initial volatility of revisions in their beliefs, increases the SDF's volatility. Lower consumption-growth volatility (compared to the case of full confidence), which stems from their precautionary demand for the risk-free asset, only partially offsets this increase. Over time, both the sensitivity of the continuation-utility and the volatility of revisions in the perceived dividend-growth rate decline, explaining the gradual reduction in the volatility of the less-confident investors' SDF.

In contrast, the increase in the volatility of more-confident investors' SDF (Panel B of Figure 7) results from higher consumption-growth volatility, which in turn results from the higher return volatility of their portfolios. In particular, because more-confident investors, in equilibrium, go short the risk-free asset and overweight the less-familiar asset, their portfolio returns are more volatile (than in the case of full confidence). Over time, as less-confident investors' confidence increases, their holding of the less-familiar asset increases and that of the safe asset decreases. Market clearing then implies a decrease in more-confident investors' holding of the less-familiar asset and their short position in the bond, which explains the decline in the SDF volatility.

Figure 8: Dynamics of Asset Returns

The figure illustrates the dynamics of the return moments of the less- and more-familiar risky assets. Panels A and B show the average conditional return volatilities, Panels C and D the average conditional risk premia, and Panels E and F the average conditional Sharpe ratios. Averages are calculated across 100,000 simulation paths and are computed under the objective beliefs. The light dash-dotted lines plot the first and third quartile of the distribution of the simulated paths. All graphs are based on the parameter values described in Table 1.

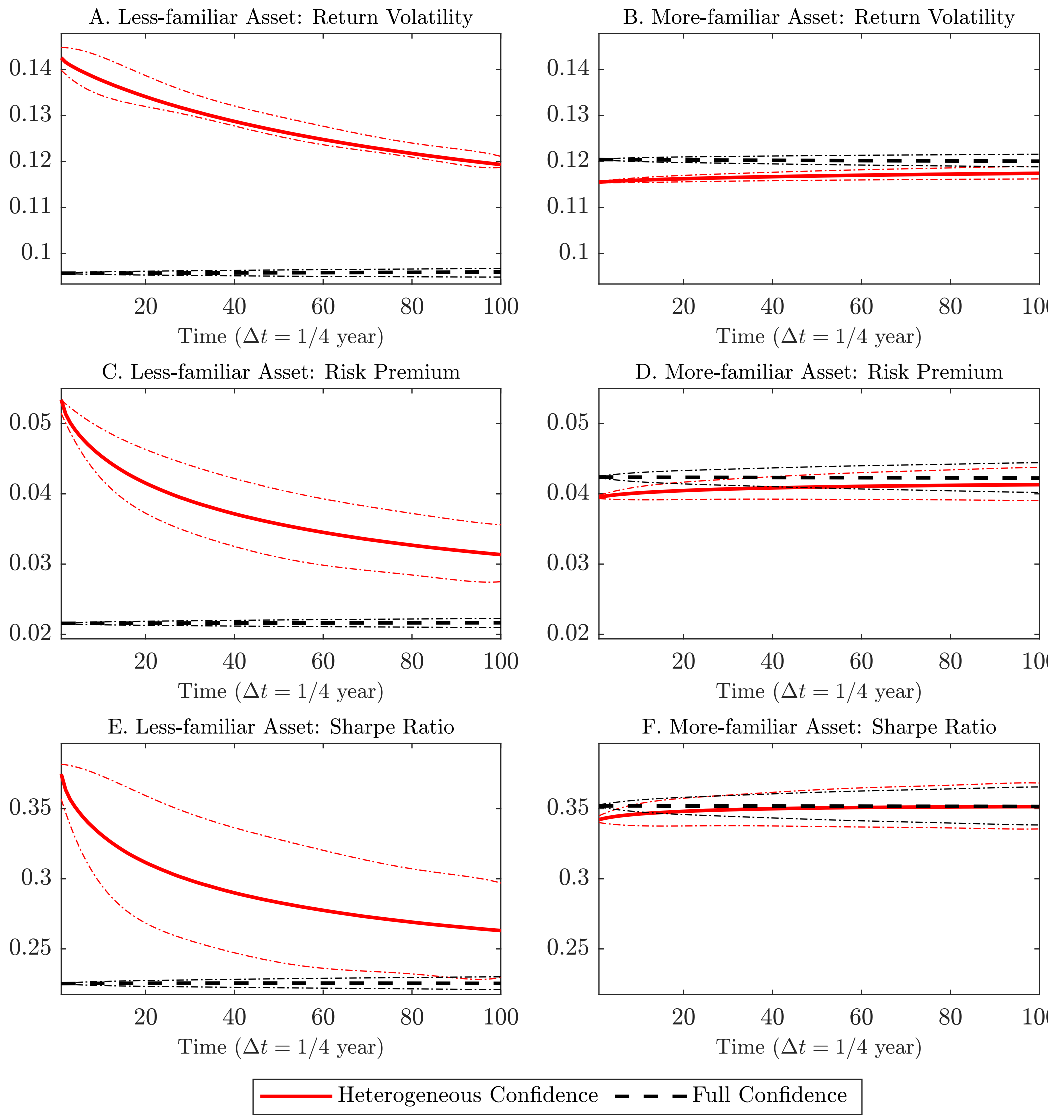

3.2 Asset Returns

Figure 8 illustrates the dynamics of the return moments of the two risky assets. In the presence of heterogeneity in confidence, the return volatility of the less-familiar (small-cap) asset is substantially higher than in the case of full confidence (Panel A). To understand the economic intuition underlying this higher volatility, recall that positive cash-flow news leads to an upward revision in the perceived dividend-growth rate (cf. Equation (3)) and vice versa. Consequently, with a preference for early resolution of uncertainty, the price-dividend ratio of the less-familiar asset increases (declines) exactly when its dividends are high (low), thereby *amplifying* the variations in dividends and creating "excess volatility."[22] Over time, as the precision of less-confident investors' beliefs increases, the magnitude of changes in their beliefs declines, and so does the return volatility. Note, however, that this occurs very gradually over time.

The higher return volatility, in combination with the higher volatility of investors' SDFs, also explains the increase in the risk premium of the less-familiar (small cap) asset, relative to the case of full confidence (Panel C). Moreover, because the less-familiar asset's risk premium increases more than its return volatility, its Sharpe ratio is considerably larger than in the full-confidence case (Panel E). Intuitively, the higher Sharpe ratio is required to induce the more-confident investors to hold a larger proportion of their wealth in the less-familiar asset, to compensate for the lower demand from less-confident investors. The Sharpe ratio also declines only slowly over time.

Importantly, consistent with data for small-cap stocks, the less-familiar asset's return volatility, risk premium, and Sharpe ratio (Panels A, C, and E) are all endogenously higher than those of the more-familiar large-cap stocks in the first few years (Panels B, D, and F)—even though their fundamentals are the same. These results are in sharp contrast to those from models with full confidence, in which the return volatility, risk premium, and Sharpe ratio of the less-familiar asset would be *lower* because of this asset's smaller share of aggregate dividends, and hence, smaller covariance with aggregate consumption (see, e.g., Cochrane et al. 2008).

We conclude our discussion of asset returns by describing, for completeness, the dynamics of the risk-free rate, which follow from the demand for the safe asset. Intuitively, the less-confident investors' precautionary-savings demand—resulting from their lower confidence—implies a higher bond price relative to the full-confidence case. Thus, the risk-free rate is lower

[22] Interestingly, in our model, excess volatility arises even if EIS ≤ 1 (as long as EIS exceeds the reciprocal of relative risk aversion) because of the substitution between the two *risky* assets. In contrast, in a single-risky-asset setting, such as Collin-Dufresne et al. (2016b), excess volatility can arise only if EIS > 1, because of the substitution between the risk-free and the single risky asset. Setting EIS > 1 in our model would also activate this channel, and hence, further amplify the variations in the less-familiar asset's price-dividend ratio and lead to a further increase the excess volatility (see also Appendix D).

Figure 9: Risk-free Rate

The figure illustrates the average risk-free rate over time. Averages are calculated across 100,000 simulation paths. The figure is based on the parameter values described in Table 1.

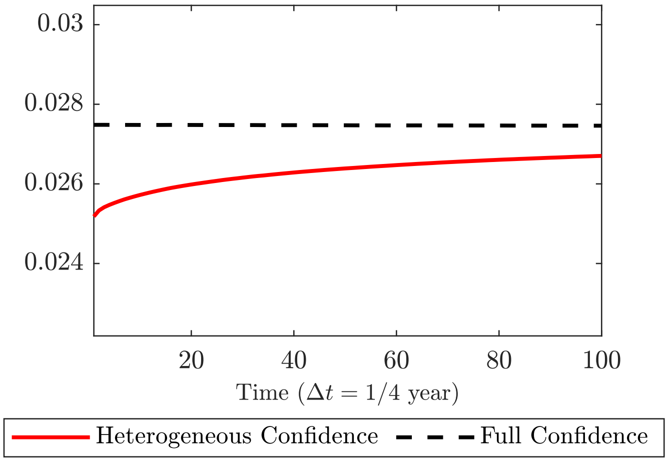

initially, as illustrated in Figure 9. Over time, as less-confident investors gain confidence, their precautionary-savings demand declines, and, hence, the risk-free rate slowly moves toward the level in the full-confidence case.

4 Investment Returns and Wealth Dynamics

We now describe the dynamics of investors' portfolio returns and how they affect the evolution of the wealth of the two groups of investors.

4.1 Heterogeneity in Investment Returns

Figure 10 illustrates how the expected investment returns and volatilities differ across the two groups of investors. Specifically, Panel A shows that less-confident investors earn—on average—substantially lower absolute portfolio returns than more-confident investors (for comparison, the common expected portfolio return in the full-confidence case is around 7.5%). This lower expected investment return is the result of two economic forces. First, because of less-confident investors' precautionary-savings demand for the risk-free asset, they hold a smaller position in risky assets, and thus, benefit less from risk premia. Second, compared to more-confident

Figure 10: Dynamics of Heterogeneity in Portfolio Returns

The figure illustrates the dynamics of the differences in investors' portfolio returns. In particular, Panels A, B, and C depict the *difference* in the average portfolio return, volatility, and Sharpe ratio of the less-confident and more-confident investors (Less-conf. − More-conf.). "Absent GE" refers to the setting in which we keep the assets' return moments fixed at their full-confidence level even as we allow some investors to be less than fully confident about the returns of the less-familiar asset. Portfolio returns are annualized, averaged across 100,000 simulation paths, and computed under the objective beliefs. All graphs are based on the parameter values described in Table 1.

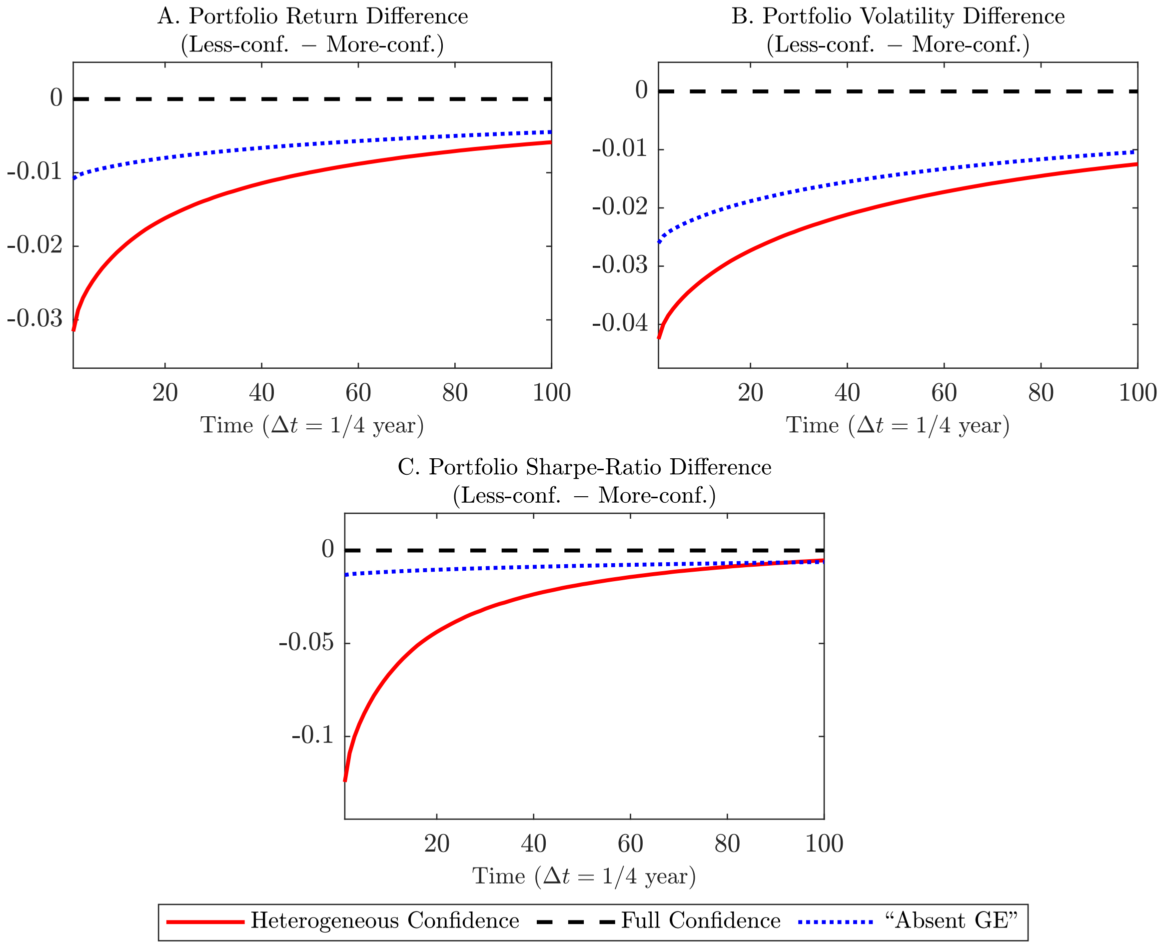

investors, they invest only a small fraction of their risky-asset portfolio in the less-familiar asset, thus benefiting less from the higher risk premium it earns in the early years.

Because of the less-confident investors' larger allocation to the safe asset, their investment returns are, on average, also considerably less volatile than those of more-confident investors, as illustrated in Panel B (for comparison, the common portfolio-return volatility in the full-confidence case is around 11.5%). This effect is further strengthened by less-confident investors' lower demand for the less-familiar asset, which initially is much more volatile than the more-familiar asset representing large-cap stocks.

Notably, however, even adjusted for portfolio risk, less-confident investors' investment returns are considerably lower than those of more-confident investors, as one can see from the difference in portfolio Sharpe ratios in Panel C (for comparison, the Sharpe ratio in the full-confidence case of the common portfolio is around 0.42). Hence, less-confident investors' lower absolute investment returns cannot be attributed to just their larger investment in the risk-free asset but, instead, are also a consequence of their portfolio underdiversification. That is, because of their limited confidence, less-confident investors' expected portfolio excess return declines disproportionately more than portfolio volatility.

The heterogeneity in investment returns between less- and more-confident investors shows substantial persistence; that is, the difference in investors' (risk-adjusted) investment returns diminishes only gradually over time. This is a consequence of the persistent heterogeneity in asset demands that we have discussed above. Indeed, only as less-confident investors slowly become more confident, thus reducing their allocation to the safe asset and improving their portfolio's diversification, do their raw and risk-adjusted investment returns improve (in absolute terms but, more importantly, also relative to those of more-confident investors).

It is important to highlight that the heterogeneity in investors' portfolio returns also arises in the *absence* of general-equilibrium asset-pricing effects; that is, in the setting in which we keep the assets' return moments fixed at their full-confidence level even as we specify some investors to be less than fully confident about the returns of the less-familiar asset (see "Absent GE" in Figure 10). However, as the figure illustrates, if one were to ignore the general-equilibrium implications for asset returns—that is, the high return volatility, risk premium, and Sharpe ratio of the less-familiar asset in its early years—one would severely underestimate the heterogeneity in portfolio returns.

In summary, the properties of investment returns in our model are consistent with the empirical literature (Lerner et al., 2008, Greenwood and Nagel, 2009, Brown et al., 2010, Goetzmann and Oster, 2013). In particular, there is substantial heterogeneity in investors' portfolio returns, with less-confident investors earning, on average, substantially lower (risk-adjusted) returns compared to more-confident investors. Like in the data, this is a consequence

Figure 11: Portfolio Volatility and Asset Returns

The figure illustrates the relation between investors' risk taking and asset returns. In particular, Panels A and B, respectively, plot the less- and more-confident investors' average portfolio volatility *conditional* on a high (low) market risk premium. "High (Low) Aggregate Risk Premium" refers to states in which the risk premium on the aggregate market portfolio is higher than its unconditional mean (on the same date). "Average" refers to the unconditional average across all paths. All quantities are based on 100,000 simulated paths of the economy and the parameter values described in Table 1.

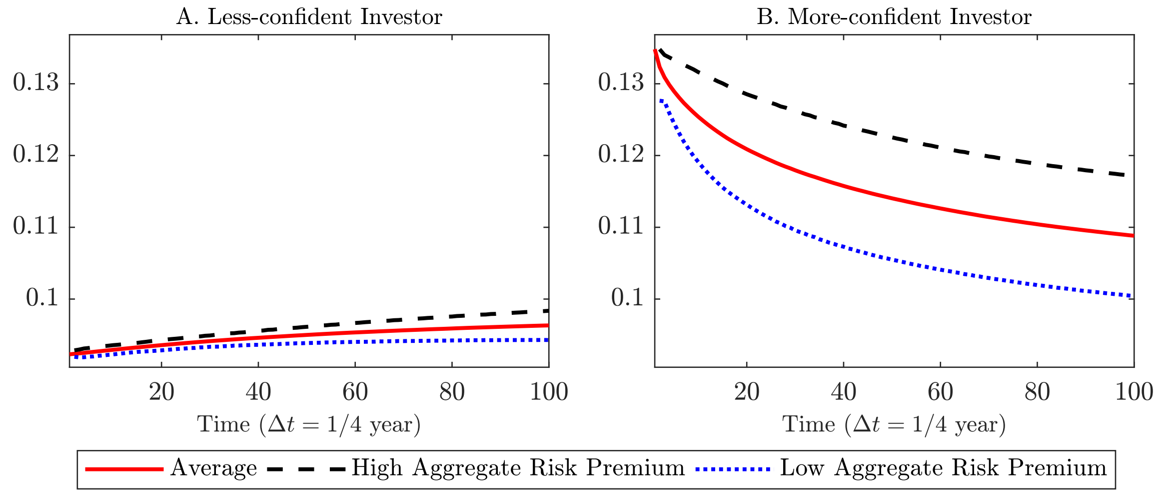

of a larger allocation to the safe asset, a smaller allocation to riskier assets that have higher expected returns, and relative underdiversification by the less-confident investors. Moreover, these differences in investment returns persist for long periods.

4.2 Asset Demands Conditional on Asset Returns

The dynamic nature of our model also allows us to study investors' risk-taking behavior *conditional* on asset returns. Figure 11 shows that the same pattern arises in our framework; that is, during periods with a high aggregate risk premium, more-confident investors choose a considerably more volatile portfolio (Panel B), compared to periods with a low aggregate risk premium. In contrast, the volatility of less-confident investors' portfolios does *not* display this behavior; that is, the portfolio volatility does not change much depending on expected returns (Panel A).

Again, this is a consequence of less-confident investors' learning, in particular, their asset demand in response to revisions in their beliefs. Specifically, following downward revisions in the perceived dividend-growth rate of the less-familiar asset, less-confident investors allocate more wealth to the safe asset and reduce their allocation to the less-familiar asset. While the higher allocation to the safe asset reduces portfolio risk, the loss in diversification creates an offsetting effect, leaving less-confident investors' portfolio volatility mostly unchanged. In response to the asset demands of less-confident investors, market clearing requires more-confident investors to reduce their allocation to the safe asset and increase their holding of the less-familiar asset, *unambiguously* driving up their portfolio volatility. At the same time, the risk premium on the aggregate market goes up. That is, the lower demand of the less-confident investors for the less-familiar asset leads to an increase in its risk premium (and expected return). Moreover, the stronger demand for the safe asset lowers the risk-free rate and leads to an increase in the risk premium of the more-familiar asset (whose expected return declines only marginally because of a slightly stronger demand).

Consequently, we observe more-confident investors taking on more risk exactly when the risk premium on the aggregate market is high, whereas less-confident investors' portfolio risk is mostly unchanged. This pattern in investors' risk taking highlights another channel through which more-confident investors obtain higher investment returns, contributing to the heterogeneity in investors' investment returns.

4.3 Wealth Dynamics

Heterogeneity in investors' portfolio returns plays a critical role in determining the dynamics of wealth. Specifically, in the presence of confidence heterogeneity, more-confident investors accumulate financial wealth at a considerably faster rate—because of their higher investment returns. Consequently, as illustrated in Panels A and B of Figure 12, the average share of aggregate financial wealth held by more-confident investors increases over time, whereas that of less-confident investors declines.[23] The divergence in wealth shares is most pronounced in the early years when investors' investment returns differ the most (cf. Figure 10). It slows down only in later years when less-confident investors start to reduce their allocation to the safe asset and invest more in the less-familiar risky asset.

[23] In the case of full confidence, both groups of investors are identical, and thus, their wealth shares remain equal to the initially endowed wealth shares of 2/3 and 1/3 for less- and more-confident investors, respectively.

Figure 12: Dynamics of Wealth Shares

Panels A and B plot the average share of aggregate wealth held by the less- and more-confident investors. For ease of comparison, we have normalized certainty-equivalent consumption by total output. "No Access Less-fam. Asset" refers to a setting in which less-confident investors have no access to the less-familiar asset. "$A_{2,1} = 40$" refers to the setting in which less-confident investors are endowed with higher initial confidence. Averages are calculated across 100,000 simulation paths. All graphs are based on parameter values listed in Table 1.

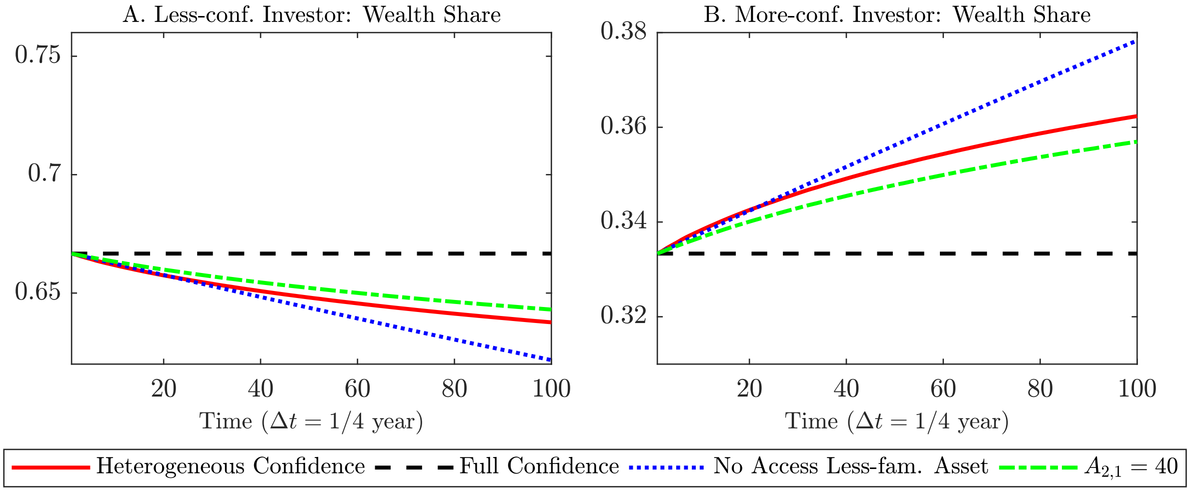

Figure 12 also shows the results for some additional cases. First, it depicts the case in which less-confident investors have no access to the less-familiar asset, whereas more-confident investors can invest in this asset. Comparing this setting to that with confidence heterogeneity highlights that giving less-confident investors access to new assets, slows down the rise in wealth inequality. Second, Figure 12 depicts the case in which less-confident investors have higher initial confidence about the dynamics of the less-familiar asset ($A_{2,1} = 40$ compared to $A_{2,1} = 20$ for our baseline case). Again, the consequence is a slower increase in wealth inequality.

5 Comparison with Alternative Models

To highlight the critical role of heterogeneity in investors' confidence about the *mean* dividend-growth rate, and to distinguish our mechanism from other potential explanations, we now compare the asset-demand and investment-return dynamics arising in our model to those under several alternative frameworks. In particular, we demonstrate that our model generates distinctly different asset-demands and investment-return dynamics than these alternative modeling frameworks (which generally do not match well the empirical evidence).

5.1 Alternative Belief Specifications

We start with an analysis of alternative belief specifications, the results of which are displayed in Figure 13. To highlight the differences to our baseline specification, we shut down confidence heterogeneity regarding the mean dividend-growth rate in these alternative specifications.

First, we consider the case in which, instead of lower confidence about the less-familiar asset's expected dividend-growth rate, the group still labeled as "less-confident investors" perceive the less-familiar asset to be riskier than it is. In particular, we assume that less-confident investors believe that the (quarterly) dividend-growth volatility of the less-familiar asset is 2.94% (equivalent to a 50% higher perceived cash-flow variance) and that they do not update their beliefs. The higher perceived risk creates a precautionary-savings demand, and hence, these investors allocate a slightly higher proportion of wealth to the risk-free asset. As Panel A of Figure 13 shows, the demand for the risk-free asset is minimal compared to that in the baseline model. Also, their demand for the less-familiar asset is lower than that of more-confident investors. However, as Panel B shows, in contrast to the baseline model, this underinvestment is very small because it is exclusively due to differences in the investors' myopic asset demand. Hence, a higher perceived risk can hardly explain the large heterogeneity in asset demands observed in the data.

Second, we consider the case in which less-confident investors are less confident about the *volatility* of the less-familiar asset (instead of the mean, as in our baseline model) but learn about it from realized dividend growth.[24] In this case, there is some heterogeneity in investors' asset demands; in particular, less-confident investors allocate more wealth to the safe asset and less wealth to the less-familiar asset initially. But, the magnitude of the asset-demand and investment-return heterogeneity is considerably smaller than in our baseline case because perceived dividend-growth volatility only governs variation in the second moments, and hence, investors' continuation utility is substantially less sensitive to these variations. Moreover, the little heterogeneity present in the early periods vanishes very quickly so that investors are very soon holding the same portfolios; that is, there is no persistence in asset-demand heterogeneity—in sharp contrast to the large persistence in our baseline case. The short-lived heterogeneity is a consequence of the fact that variance parameters are, in general, relatively easy to learn.

[24] The details of this setting are given in Appendix C.

Figure 13: Alternative Belief Specifications

The figure illustrates how asset-demand heterogeneity and investment returns evolve under alternative belief specifications. Panels A and B plot the average proportion of wealth invested in the risk-free bond and Panels C and D in the less-familiar asset—by less- and more-confident investors. Panels E and F depict the *differences* in the average portfolio volatility and the portfolio Sharpe ratios of less- and more-confident investors (Less-conf. − More-conf.). Portfolio returns are annualized and computed under the objective beliefs. "Higher Perc. Risk" refers to a setting in which less-confident investors perceive the dividend-growth volatility of the less-familiar asset to be higher; in particular, they perceive the (quarterly) volatility to be 2.94% (instead of 2.40%) in every period. "Conf. Heterog. Volatility" refers to a setting in which investors differ in their confidence regarding the less-familiar asset's dividend-growth *volatility* but learn about it over time. "Fixed Heterog. Beliefs" refers to a setting in which less-confident investors have time-invariant pessimistic beliefs (with $\hat{\mu}_{2,t} = 0.3\%, \forall t$), while the other class of investors has unbiased beliefs. In all three cases, investors do not need to learn about the mean dividend-growth rate. Averages are calculated across 100,000 simulation paths. Other than the values for the parameters described above, all parameters take the values described in Table 1.

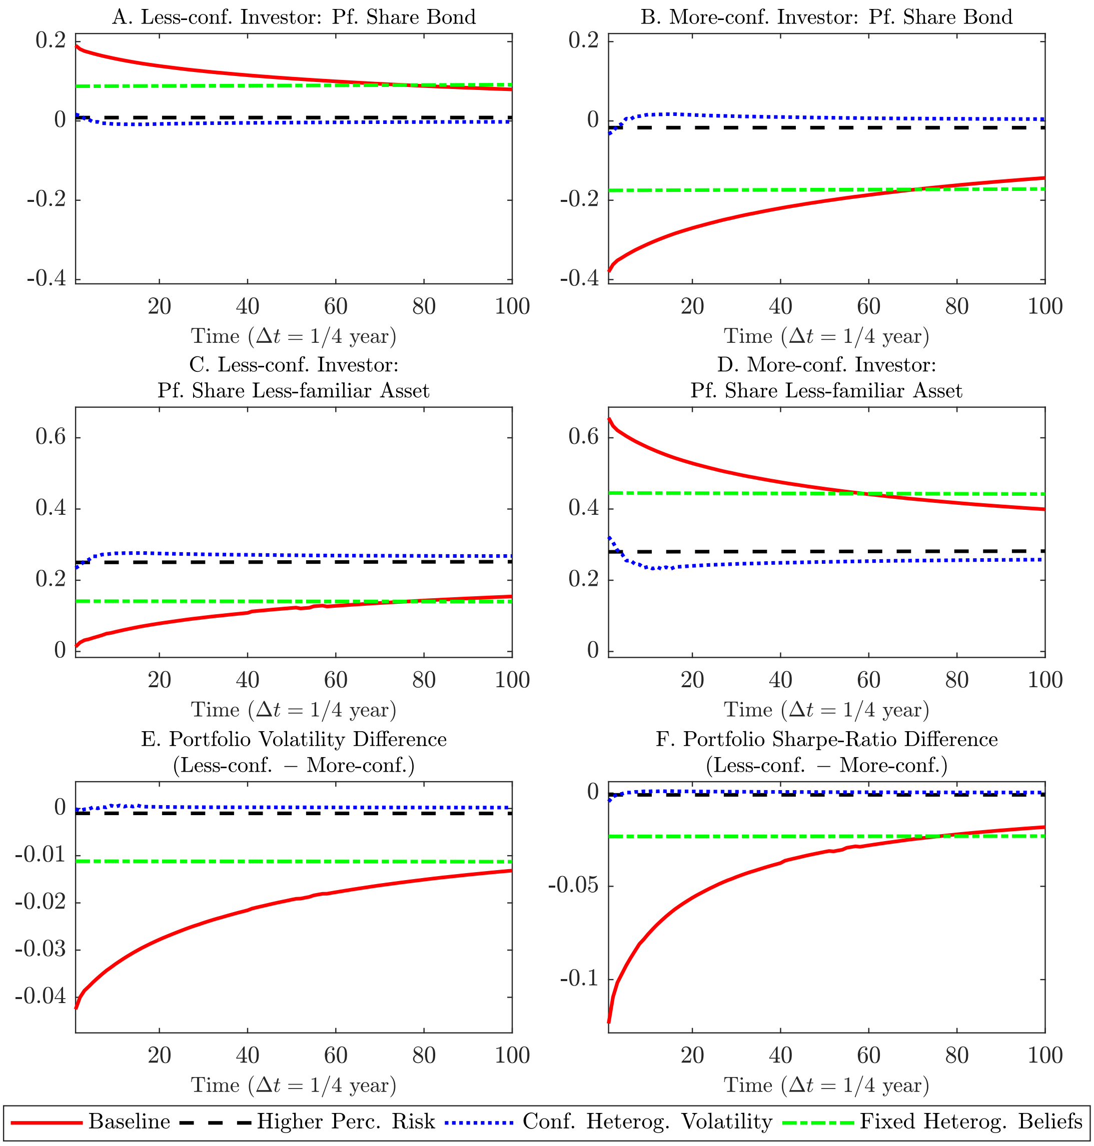

Third, we study the case in which "less-confident investors" have time-invariant pessimistic beliefs, while the other group of investors has unbiased beliefs. In this case, the pessimistic investors allocate more wealth to the risk-free asset (because they perceive consumption growth to be lower) and less wealth to the less-familiar asset (exclusively due to a low myopic demand). While these patterns in average portfolio allocations resemble most closely those in our baseline framework, time-invariant beliefs have, naturally, a difficult time rationalizing the *dynamics* of investors' portfolio allocations and returns observed empirically.

Notably, all three models also fail to generate the asset demands conditional on realized and expected returns that are observed empirically—trend-chasing by less-confident investors and increased risk taking when the market risk premium is high by more-confident investors.

5.2 Alternative Preference Specifications

We now turn to the case of alternative preference specifications. The results of these experiments are displayed in Figure 14.

First, we consider the case in which investors differ in their confidence regarding the mean dividend-growth rate of the less-familiar asset (exactly as in our baseline model) but assume that investors have *log* utility instead of Epstein-Zin-Weil recursive utility. In this case, there is practically no heterogeneity in investors' average asset demands and investment returns despite heterogeneous confidence. That is, allocations of the less- and more-confident investors to the safe asset are practically the same (Panels A and B). Moreover, as is well known, investors with log utility do not hedge changes in the state variables. Thus, there is no intertemporal hedging demand for the less-familiar asset, and therefore, on average, investors' allocations for the less-familiar asset are also similar (Panels C and D). As a result, their (risk-adjusted) average investment returns are also the same (Panel E). Thus, a model with log utility cannot generate the asset-demand *dynamics* observed in the data.

Next, we consider the case where investors disagree about the mean dividend-growth rate of the less-familiar asset but have *CRRA* utility (with the same risk aversion $\gamma = 10$ as in our baseline model). In this case, the asset demands in the early periods are comparable to those

Figure 14: Alternative Preference Specifications

The figure illustrates how asset-demand heterogeneity and investment returns evolve under alternative preference specifications. Panels A and B plot the average proportion of wealth invested in the risk-free bond and Panels C and D in the less-familiar asset—by less- and more-confident investors. Panels E and F depict the *differences* in the average portfolio volatility and the portfolio Sharpe ratios of less- and more-confident investors (Less-conf. − More-conf.). Portfolio returns are annualized and computed under the objective beliefs. "Log Utility" refers to a setting in which both investors have log utility and differ in their confidence regarding the mean dividend-growth rate of the less-familiar assets. "CRRA" refers to a setting in which both investors have CRRA utility ($\gamma = 10$) and differ in their confidence regarding the mean dividend-growth rate of the less-familiar assets. "Heterog. RA" refers to a setting in which "less-confident" investors have a risk aversion of 7 whereas the other investors have a risk-aversion of 13—in the absence of confidence heterogeneity. "Heterog. EIS" refers to a setting in which "less-confident" investors have an EIS of 1.5 whereas the other investors have an EIS of 0.5—in the absence of confidence heterogeneity. Averages are calculated across 100,000 simulation paths. Other than the values for the parameters described above, all parameters take the values described in Table 1.

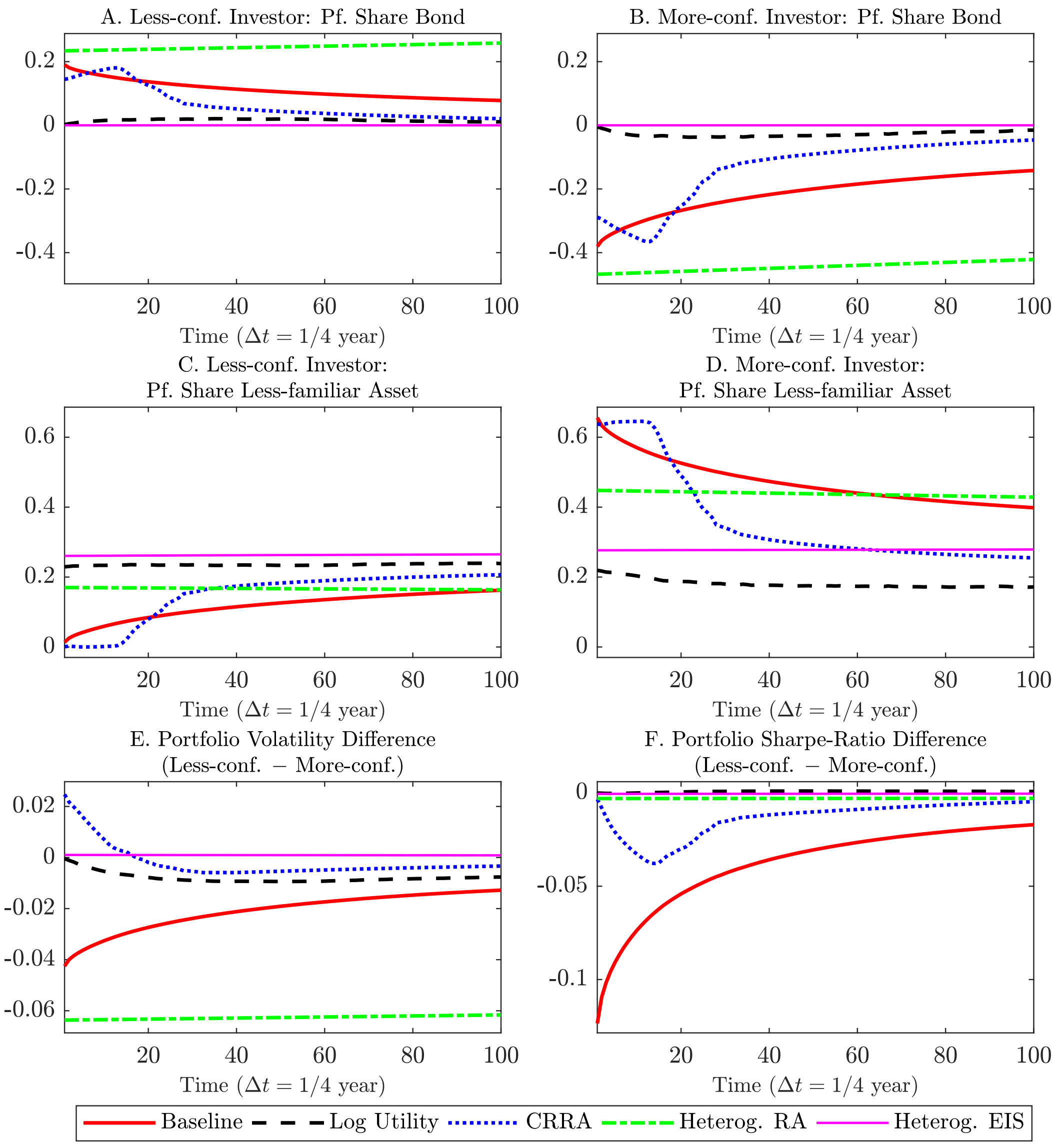

in our baseline model, but there is much less persistence.[25] Moreover, the patterns in the risk and returns of investors' portfolios are quite different from the baseline model and the data. In particular, less-confident investors' portfolio returns are initially *more* volatile—because of the more-familiar asset's higher return volatility (compared to that of the less-familiar asset). Moreover, as is well known, models of CRRA preferences, in general, have a hard time matching asset-pricing moments; indeed, the risk-free rate shoots up, and risk premia diminish.

Next, we study cases in which investors differ in their *preferences* but have the same level of (full) confidence. If the two groups of investors differ only in their relative risk aversion, then the more risk-averse investors (which, for simplicity, we continue to label as "less-confident investors" in the graphs) allocate a larger fraction of their wealth to the risk-free asset. This pattern is similar to the one present in our framework and is again driven by their stronger precautionary-savings demand. More risk-averse investors also reduce their investment in the less-familiar risky asset, but the higher risk aversion leads them to invest less in *both* risky assets. Consequently, there is no loss in diversification benefits. Therefore, while the more risk-averse investors earn lower absolute investment returns and face lower portfolio risk, in contrast to the empirical evidence, the *risk-adjusted* investment returns of the two groups of investors are the same, because the decline in portfolio volatility offsets precisely the decline in the portfolio excess return.

If the two groups of investors differ only in their EIS, then the investors with a stronger desire to smooth consumption over time (labeled "less-confident investors") allocate marginally more of their wealth to the safe asset and slightly reduce their allocation to both risky assets—relative to the investors with lower EIS. Overall, however, asset-demand heterogeneity is quite limited, and both groups of investors practically earn the same risk-adjusted investment returns (cf. Panel F). So, heterogeneity in EIS also fails to generate the heterogeneity in asset-demand dynamics observed empirically.

[25]With CRRA preference, positive (negative) cash flow news generally leads to a decline (increase) in the less-familiar asset's price-dividend ratio, offsetting the fluctuations in dividends. As a result, the less-familiar asset's return volatility is very low in the early periods which causes rather "extreme" holdings—due to a large precautionary-savings demand and a large hedging demand (despite fluctuations in expected utility being lower by a factor of almost 100 compared to the baseline model). The low return volatility—coupled with the short-sale constraint—also explains the small non-monotonicities in asset demands. Specifically, because investors' asset demands s are very sensitive to changes in the perceived dividend-growth rate, the short-sale constraint binds more frequently. If one were to relax the short-sale constraint a bit, these non-monotonicities would vanish.

If the two groups of investors differ in their rate of time preferences (omitted in the figure for legibility), less-patient investors (dubbed "less-confident investors") go slightly long the risk-free asset but *overweight* the less-familiar asset because more-patient investors prefer the more-familiar asset, which has a higher expected return in this setting. As a result, less-confident investors have lower portfolio volatility but a higher portfolio Sharpe ratio (because their portfolio is better diversified), contrary to the empirical evidence.

5.3 Other Model Specifications

Figure 15 reports the implications of variations in the key model parameters.

First, instead of prohibiting short sales as in the baseline model, if we allow for short sales it strengthens our findings; in particular, the heterogeneity in asset demands and investment returns increases. Intuitively, when short-sales are permitted, less-confident investors take more extreme positions in the less-familiar asset (Panel C) and, consequently, also in the bond (Panel A) and the more-familiar asset (not shown). As a result, the volatility and Sharpe ratio of their portfolio returns also differ more strongly from those of the more-confident investors (Panels E and F), strengthening the implications for wealth dynamics.

Second, widening the support of the dividend-growth rate as perceived by the less-confident investors, $\hat{\mu}_{2,t}$, (truncation bounds of $[-1.55\%, 2.45\%]$ p.a. instead of $[-0.55\%, 1.45\%]$) has—even quantitatively—a negligible impact on the results. That is, while there are some very small effects early in the sample (i.e., for dispersed prior beliefs), they quickly vanish as the investors become more confident.

Third, reducing the initial wealth share of the less-confident investors ($w_{2,1} = 1/2$ instead of $w_{2,1} = 2/3$) limits the heterogeneity in investors' asset demands and investment returns. Qualitatively, however, the results remain unchanged.

Finally, as expected, increasing the "size" of the less-familiar asset, i.e., increasing the set of stocks that are not well understood, by increasing this asset's initial dividend share ($\delta_{2,1} = 0.35$ instead of $\delta_{2,1} = 0.20$), strengthens the impact of confidence heterogeneity. That is, both the precautionary-savings demand and the intertemporal-hedging demand increase—widening the

Figure 15: Other Model Specifications

The figure illustrates how asset-demand heterogeneity and investment returns evolve under other model specifications. Panels A and B plot the average proportion of wealth invested in the risk-free bond and Panels C and D in the less-familiar asset—by less- and more-confident investors. Panels E and F depict the *differences* in the average portfolio volatility and the portfolio Sharpe ratios of less- and more-confident investors (Less-conf. − More-conf.). Portfolio returns are annualized and computed under the objective beliefs. "Short-sales" refers to a setting in which short sales are allowed. "Wider Support $\hat{\mu}_{k,t}$" refers to a setting with wider truncation bounds for $\hat{\mu}_{k,t}$; equal to $[-1.55\%, 2.45\%]$ p.a. "$w_{2,1} = 0.5$" refers to a setting in which less-confident investors are initially endowed with half of the wealth. $\delta_{2,1} = 0.35$ refers to a setting in which the initial dividend share of the less-familiar asset is 0.35. Averages are calculated across 100,000 simulation paths. Other than the values for the parameters described above, all parameters take the values described in Table 1.

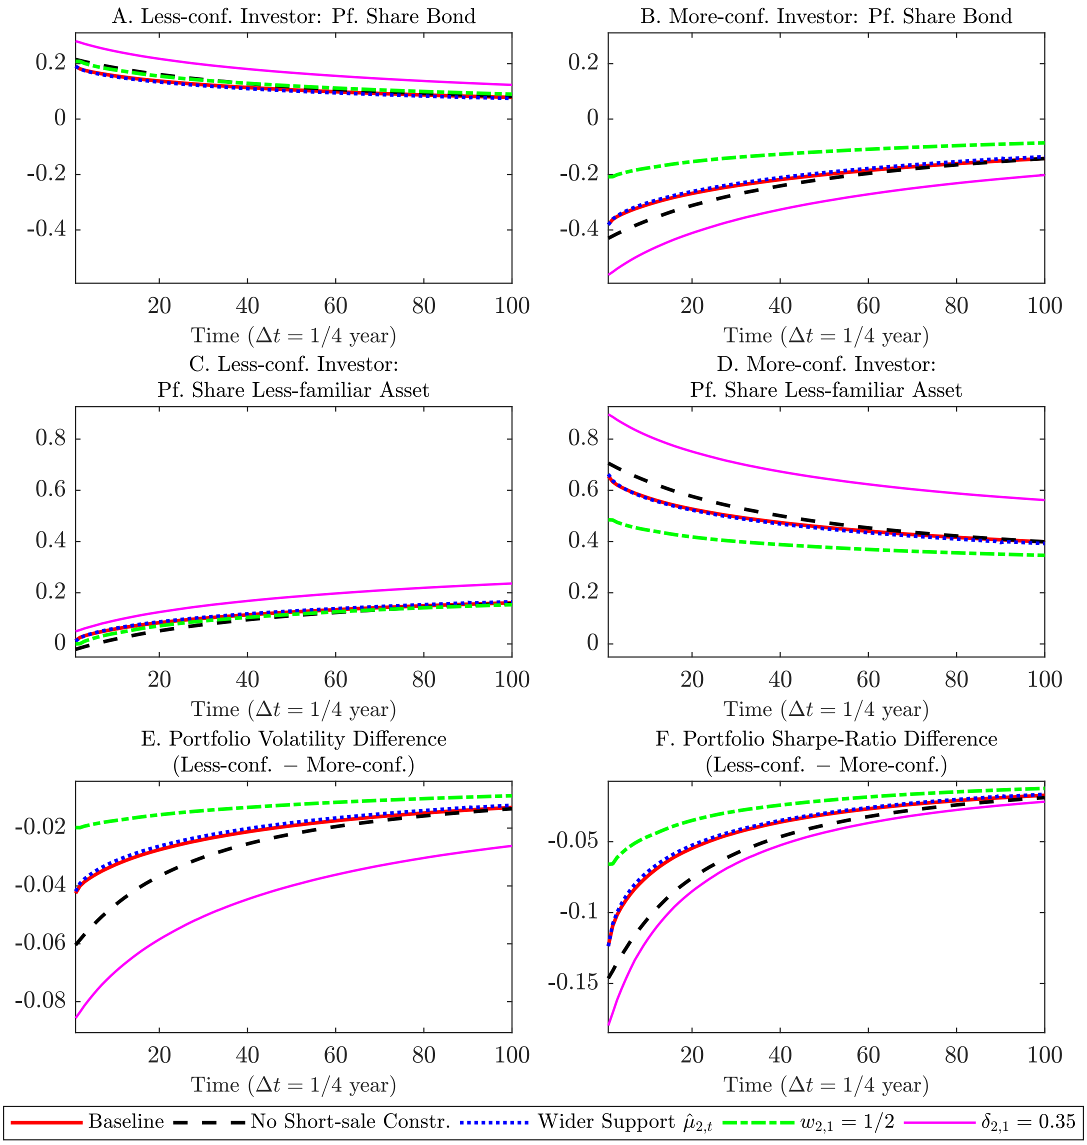

gap in investors' asset demands. As a result, the differences in investors' investment returns are also more pronounced.

6 Conclusion

Recent empirical evidence has documented many intriguing patterns in the asset demands of financial institutions and households that persist for long periods (Koijen and Yogo, 2019). In this paper, we develop a dynamic general-equilibrium framework that has multiple investors and multiple risky assets. The salient feature of the model is the presence of a latent variable—differences in beliefs, which are a consequence of heterogeneity in investors' confidence about expected asset returns.

The model we develop is consistent with several static and dynamic properties of asset demands observed empirically. In the model, investors' asset holdings are concentrated in a subset of available assets, with large and persistent heterogeneity in asset demands across investors. Less-confident (typically smaller) institutions invest largely in large-cap stocks, trend-chase, and earn lower absolute and risk-adjusted investment returns, while more-confident (typically larger) institutions hold a larger share of wealth in risky assets, time the market, and earn consistently higher investment returns. The model also generates elasticities for asset demands that are lower than those in the traditional model with full confidence. The time-series and cross-sectional variation in assets' realized and expected returns, as well as their volatilities, are driven by the latent variable.

The model of an endowment economy that we develop in this paper is a first step toward addressing Koijen's statement in Brunnermeier et al. (2021, page 2143) that: "A central question is how investors' demand responds to price changes and to changes in asset characteristics, and how investors substitute across various assets and asset classes." It would be straightforward to extend the model to allow for endogenous production so that one can study the implications of the latent variable also for real quantities. The general-equilibrium nature of the model also makes it suitable for answering questions about macroeconomic policies, such as quantitative easing, that are designed to work through their effect on the demand for financial assets.

A Optimality Conditions and Equilibrium

The objective of each investor k is to maximize her expected lifetime utility given in Equation (2), by choosing consumption, $C_{k,t}$, and the holdings in the available financial assets, $\theta_{n,k,t}, n \in \{0, \dots, 2\}$:

$$V_{k,t}\big(\{\theta_{k,n,t-1}\}\big) = \max_{C_{k,t},\{\theta_{k,n,t}\}} \left[(1-\beta)\, C_{k,t}^{1-\frac{1}{\psi}} + \beta E_t^k \left[V_{k,t+1}\big(\{\theta_{k,n,t}\}\big)^{1-\gamma} \right]^{\frac{1}{\phi}} \right]^{\frac{\phi}{1-\gamma}},$$

subject to the budget equation (7) and the short-sale constraint $\theta_{k,2,t} \geq 0$.[26]

Denoting the Lagrange multiplier associated with the budget equation by $\eta_{k,t}$ and that of the short-sale constraint by $\Lambda_{k,t}$, the Lagrangian can be written as

$$\begin{aligned} \mathcal{L}_{k,t} = \sup_{C_{k,t},\{\theta_{k,n,t}\}} \inf_{\eta_{k,t}} & \left[(1-\beta)\, C_{k,t}^{1-\frac{1}{\psi}} + \beta E_t^k \left[V_{k,t+1}^{1-\gamma} \right]^{\frac{1}{\phi}} \right]^{\frac{\phi}{1-\gamma}} \\ & + \eta_{k,t} \left(\theta_{k,0,t-1} + \sum_{n=1}^{2} \theta_{k,n,t-1}\, D_{n,t} - C_{k,t} - \theta_{k,0,t}\, S_{0,t} - \sum_{n=1}^{2} \Delta\theta_{k,n,t}\, S_{n,t} \right) + \Lambda_{k,t}\eta_{k,t}\theta_{k,2,t}, \end{aligned}$$

and the corresponding Karush-Kuhn-Tucker first-order conditions are given by

$$\begin{aligned} \frac{\partial \mathcal{L}_{k,t}}{\partial C_{k,t}} &= \frac{1}{1-\frac{1}{\psi}} \left\{ (1-\beta)\, C_{k,t}^{1-\frac{1}{\psi}} + \beta E_t^k \left[V_{k,t+1}^{1-\gamma} \right]^{\frac{1-\frac{1}{\psi}}{1-\gamma}} \right\}^{\frac{1}{1-\frac{1}{\psi}}-1} (1-\beta) \left(1 - \frac{1}{\psi}\right) C_{k,t}^{-\frac{1}{\psi}} - \eta_{k,t} \\ &= (1-\beta)\, C_{k,t}^{-\frac{1}{\psi}}\, V_{k,t}^{\frac{1}{\psi}} - \eta_{k,t} \equiv 0, \end{aligned} \tag{A1}$$

$$\frac{\partial \mathcal{L}_{k,t}}{\partial \eta_{k,t}} = \theta_{k,0,t-1} + \sum_{n=1}^{2} \theta_{k,n,t-1}\, D_{n,t} - C_{k,t} - \theta_{k,0,t}\, S_{0,t} - \sum_{n=1}^{2} \Delta\theta_{k,n,t}\, S_{n,t} \equiv 0, \quad \text{and} \tag{A2}$$

$$\begin{aligned} \frac{\partial \mathcal{L}_{k,t}}{\partial \theta_{k,n,t}} &= \frac{1}{1-\frac{1}{\psi}}\, V_{k,t}^{\frac{1}{\psi}}\, \beta\, \frac{1-\frac{1}{\psi}}{1-\gamma}\, E_t^k \left[V_{k,t+1}^{1-\gamma} \right]^{\frac{\gamma-\frac{1}{\psi}}{1-\gamma}} (1-\gamma)\; E_t^k \left[V_{k,t+1}^{-\gamma} \frac{\partial V_{k,t+1}}{\partial \theta_{k,n,t}} \right] - \eta_{k,t} S_{n,t} \\ &= \beta\, V_{k,t}^{\frac{1}{\psi}}\, E_t^k \left[V_{k,t+1}^{1-\gamma} \right]^{\frac{1-\frac{1}{\psi}}{1-\gamma}-1} E_t^k \left[V_{k,t+1}^{-\gamma} \frac{\partial V_{k,t+1}}{\partial \theta_{k,n,t}} \right] - \eta_{k,t}\, S_{n,t} \equiv 0, \quad n \in \{0,1\} \end{aligned} \tag{A3}$$

$$\frac{\partial \mathcal{L}_{k,t}}{\partial \theta_{k,2,t}} = \beta\, V_{k,t}^{\frac{1}{\psi}}\, E_t^k \left[V_{k,t+1}^{1-\gamma} \right]^{\frac{1-\frac{1}{\psi}}{1-\gamma}-1} E_t^k \left[V_{k,t+1}^{-\gamma} \frac{\partial V_{k,t+1}}{\partial \theta_{k,2,t}} \right] - \eta_{k,t}\, (1-\Lambda_{k,t}) S_{n,t} \equiv 0, \tag{A4}$$

$$\Lambda_{k,t}\theta_{k,2,t} \geq 0, \quad \theta_{k,2,t} \geq 0, \quad \Lambda_{k,t} \geq 0.$$

[26] For brevity, in the following derivations, we do not explicitly write the dependence of $V_{k,t}$ on the incoming (i.e., date $t-1$) asset holdings, $\{\theta_{k,n,t-1}\}$.

Using the Envelope Theorem, we can compute the derivatives of the value function $V_{k,t}$ with respect to $\theta_{k,n,t-1}$:

$$\frac{\partial V_{k,t}}{\partial \theta_{k,0,t-1}} = \frac{\partial \mathcal{L}_{k,t}}{\partial \theta_{k,0,t-1}} = \eta_{k,t}, \tag{A5}$$

$$\frac{\partial V_{k,t}}{\partial \theta_{k,n,t-1}} = \frac{\partial \mathcal{L}_{k,t}}{\partial \theta_{k,n,t-1}} = \eta_{k,t}\left(D_{n,t} + S_{n,t}\right), \quad n \in \{1,2\}. \tag{A6}$$

In summary, the optimality conditions for each investor k are given by the following set of equations. First, the budget equation from (A2):

$$C_{k,t} + \theta_{k,0,t}\, S_{0,t} + \sum_{n=1}^{2} \Delta\theta_{k,n,t}\, S_{n,t} \;=\; \theta_{k,0,t-1} + \sum_{n=1}^{2} \theta_{k,n,t-1}\, D_{n,t}, \tag{A7}$$

which equates the uses and sources of funds. Second, following from (A3) to (A6), the pricing equations which equate the price of an asset to the expected payoff from holding it:

$$\begin{aligned} S_{0,t} &= E_t^k\left[M_{k,t+1}\right], \\ S_{1,t} &= E_t^k\left[M_{k,t+1}\left(S_{1,t+1} + D_{1,t+1}\right)\right], \\ S_{2,t} &= \frac{1}{1-\Lambda_{k,t}}\, E_t^k\left[M_{k,t+1}\left(S_{1,t+1} + D_{1,t+1}\right)\right], \end{aligned}$$

where the stochastic discount factor $M_{k,t+1}$, given in Equation (8) on page 10, subsumes the Lagrange multiplier $\eta_{k,t}$ from Equation (A1). Finally, the complementary slackness and inequality conditions: $\Lambda_{k,t}\theta_{k,2,t} \geq 0$, $\theta_{k,2,t} \geq 0$, $\Lambda_{k,t} \geq 0$.

Equilibrium is then be characterized by the following equations: the budget equation (A7), the "kernel conditions" that equate the prices of the assets across investors:

$$E_t^1\left[M_{1,t+1}\right] = E_t^2\left[M_{2,t+1}\right], \tag{A8}$$

$$E_t^1\left[M_{1,t+1}\left(S_{1,t+1} + D_{1,t+1}\right)\right] = E_t^2\left[M_{2,t+1}\left(S_{1,t+1} + D_{1,t+1}\right)\right], \tag{A9}$$

$$\frac{1}{1-\Lambda_{1,t}}\, E_t^1\left[M_{1,t+1}\left(S_{2,t+1} + D_{2,t+1}\right)\right] = \frac{1}{1-\Lambda_{2,t}}\, E_t^2\left[M_{2,t+1}\left(S_{2,t+1} + D_{2,t+1}\right)\right], \tag{A10}$$

and the market-clearing conditions:[27]

$$\sum_{k=1}^{2} \theta_{k,0,t} = 0, \quad \text{and} \qquad \sum_{k=1}^{2} \theta_{k,n,t} = 1, \quad n \in \{1,2\}. \tag{A11}$$

[27] By Walras' law, clearing in the asset markets guarantees market clearing for the consumption good.

B Numerical Algorithm

We use the time-shift proposed by Dumas and Lyasoff (2012) to obtain a recursive system of equations characterizing equilibrium. That is, at date t, the "shifted" system of equations consists of the date-t kernel conditions (A8) and (A9), the date-t market-clearing conditions (A11), and the date-$t+1$ budget equations (A7):

$$C_{k,t+1,j} + \theta_{k,0,t+1,j} S_{0,t+1,j} + \sum_{n=1}^{2} \left(\theta_{k,1,t+1,j} - \theta_{k,1,t}\right) S_{1,t+1,j} \leq \theta_{k,0,t} + \sum_{n=1}^{2} \theta_{k,n,t}\, D_{n,t+1,j}, \forall k, j,$$

where the J future states (nodes) are denoted by $j = 1, \ldots, J$.[28] In total, we have a system of $2 \times J + 2 \times 3$ equations with $2 \times J + 2 \times 3$ unknowns: next period's consumption, $C_{k,t+1,j}$, for both investors and J states, and both investors' holdings in the three assets, $\theta_{k,n,t}$.

The system of equations is solved recursively, starting from $T - 1$. At each date t, we solve the equation system over the *grid of the state variables*. Next, when solving the system for date $t - 1$, we interpolate (over the grid) the optimal date-t portfolio positions, $\theta_{k,n,t}$ and corresponding security prices, $S_{n,t}$, using the terminal conditions $\theta_{k,n,T} = 0$ and $S_{n,T} = 0$, $\forall n, k$. After solving the shifted system for all dates $t \in \{0, ..., T-1\}$, one has solved all equations from the global system—except the date-0 budget equations, which have not been used because of the time shift. Thus, one only needs to solve the time-0 budget equations based on interpolating functions for the date-0 prices, $S_{n,0}$, and holdings, $\theta_{k,n,0}$. The endowed holdings $\theta_{k,n,-1}$ are exogenous to the system and reflect the incoming (endowed) wealth of the investors.

C Confidence Heterogeneity for Asset's Dividend Volatility

In the case of uncertainty regarding the less-familiar asset's cash-flow *volatility*, we assume that less-confident investors start at date $t = 1$ with a conjugate prior of an Inverse-Gamma distribution $\sigma_2 \sim \mathcal{IG}\left(\frac{\kappa_{1,t}}{2}, \frac{\kappa_{2,t}}{2}\right)$. This prior, combined with the dividend dynamics in (1), implies a time-t posterior density function $p(\sigma_2 | \Delta d_{2,1}, \ldots, \Delta d_{2,t}) = \mathcal{IG}\left(\frac{\kappa_{1,t+1}}{2}, \frac{\kappa_{2,t+1}}{2}\right)$, with the dynamics of $\kappa_{1,t+1}$ and $\kappa_{2,t+1}$ given by

$$\kappa_{1,t+1} = \kappa_{1,t} + 1, \;\; \kappa_{2,t+1} = \kappa_{2,t} + \frac{\left(\Delta d_{2,t+1} - \mu_2\right)^2}{1 + A_{k,t}}, \text{ and } A_{k,t+1} = \frac{1}{1/A_{k,t} + 1}.$$

[28] We approximate the joint dynamics of the dividends in (1) using a tree that has six nodes with growth realizations $\{(u_1, u_2), (u_1, m_2), (u_1, d_2), (d_1, u_2), (d_1, m_2), (d_1, d_2)\}$, where $u_n \equiv \mu_n + \sigma_n$, $m_n \equiv \mu_n$, and $d_n \equiv \mu_n - \sigma_n$ are chosen to match the expected dividend-growth rate and volatility of asset n. Under the less-confident investors' probability measure, the probabilities are set to match the less-confident investors' perceived dividend-growth rate, $\hat{\mu}_{2,t}$ and perceived dividend volatility $\sqrt{1 + A_{2,t}}\, \sigma_2$.

In particular, given distribution $\mathcal{IG}\left(\frac{\kappa_{1,t}}{2}, \frac{\kappa_{2,t}}{2}\right)$, the expected future cash-flow variance is given by $\kappa_{2,t}/(\kappa_{1,t} - 2)$.[29]

For the illustration in Section 5, we rely on the following parameter values $\kappa_{1,1} = 5$, $\kappa_{2,1} = \sigma_2^2(\kappa_{1,1} - 2)$ (guaranteeing initial beliefs that are unbiased), and $A_{2,k} = 20$ (as in the case of learning about the expected dividend-growth rate). Finally, we truncate the (quarterly) perceived cash flow variance at 0.0294^2 and 0.017^2 (i.e., plus/minus 50% of the true cash-flow variance of 0.024^2). All other parameter values are as specified in Table 1.

D Robustness Results

While our main mechanism requires only a preference for the early resolution of uncertainty, the magnitude of heterogeneity in investors' portfolios and investment returns *within our framework* naturally depends on the parameter values for investors' preferences and beliefs as well as the dividend dynamics. We now briefly describe the (quantitative) impact of variations in these parameter values.

Investors' Preferences: A reduction in the relative risk aversion of both groups of investors, γ, weakens investors' desire to smooth consumption across states, and hence, less-confident investors' precautionary-savings and their intertemporal-hedging demand decline (in absolute terms). Consequently, heterogeneity in investors' portfolios decreases. This naturally also implies a decline in investment-return heterogeneity, though the effect is partially offset by a further divergence in the risky assets' return moments. In particular, because less-confident investors trade more aggressively if they are less risk averse, excess volatility further increases, and accordingly, the risk premium of the less-familiar asset increases relatively more.

An increase in both investors' elasticity of intertemporal substitution (EIS), ψ, leaves less-confident investors' portfolio allocation largely unchanged. Indeed, the primary effect of a higher EIS is an increase in the return volatility and risk premium of the less-familiar asset. In particular, for EIS > 1, investors also substitute between the risk-free asset and the risky assets in response to changes in their beliefs regarding the less-familiar asset—instead of substituting only between the two risky assets. This further increases the less-familiar asset's return volatility and risk premium, but the impact on heterogeneity in portfolio returns is small. Overall, the effects of a change in EIS are rather small. These results are displayed in Figure D1.

[29] In the numerical solution approach, we then set the probabilities of the dividend tree to match this perceived dividend volatility.

Investors' Beliefs: Not surprisingly, an increase in less-confident investors' initial confidence, namely in the precision of their initial beliefs, $A_{2,1}$, reduces their demand for precautionary savings and intertemporal hedging (because their beliefs fluctuate less over time). As a result, heterogeneity in investors' portfolio holdings and investment returns declines. Reductions in the less-familiar asset's return volatility and risk premium (resulting from smaller fluctuations in the price-dividend ratio and investors' SDFs) further limit heterogeneity in investment returns.

Variations in the level of less-confident investors' initial beliefs, $\hat{\mu}_{2,k}$, also have intuitive effects. For example, if less-confident investors are initially optimistic regarding the less-familiar asset's dividend-growth rate, they allocate more capital to the asset, and hence, portfolio and investment-return heterogeneity declines. Opposite effects arise if less-confident investors are initially pessimistic about the less-familiar asset. These results are displayed in Figure D2.

Dividend Dynamics: Intuitively, the smaller is the less-familiar asset (captured by its initial dividend share $\delta_{2,1}$), the less relevant (for aggregate consumption) are fluctuations in its perceived dividend-growth rate. Consequently, both the precautionary-savings demand and the intertemporal-hedging demand decline, as do heterogeneity in investors' portfolios and investment returns. The decrease in the heterogeneity in investors' portfolios is partially offset by a relatively larger increase in the less-familiar asset's risk premium (because changes in less-confident investors' demand for this asset have to be absorbed by a smaller supply). As expected, an increase in the size of the less-familiar asset creates effects of opposite sign. These results are displayed in Figure D3.

Finally, a positive correlation between the dividend processes of the two risky assets has negligible effects, with heterogeneity in investors' asset demands and portfolio returns declining slightly because of the lower diversification benefits that the second (less-familiar) asset provides.

Figure D1: Robustness—Variations in Preferences

The figure illustrates how portfolio heterogeneity and investment returns vary with the preference parameters. Panels A and B plot the average proportion of wealth invested in the risk-free bond and Panels C and D plot the average proportion of wealth invested in the less-familiar asset—by less- and more-confident investors. Panel E depicts the *difference* in the average portfolio Sharpe ratio between less- and more-confident investors (Less-conf. − More-conf.). Portfolio returns are annualized and computed under the objective beliefs. Averages are calculated across 100, 000 simulation paths. All results are based on the parameter values described in Table 1, with the difference relative to our baseline case highlighted in the legend.

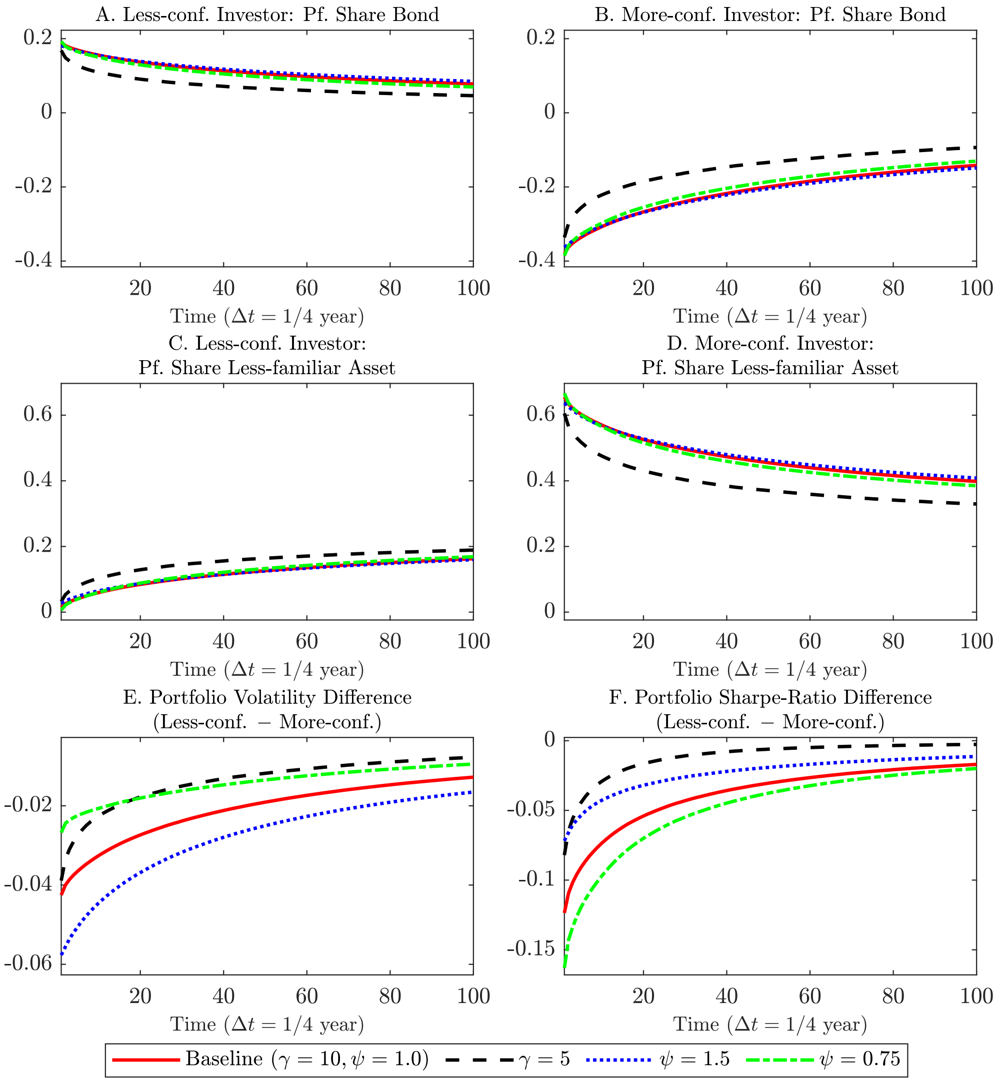

Figure D2: Robustness—Variations in Beliefs

The figure illustrates how portfolio heterogeneity and investment returns vary with the parameters governing beliefs. Panels A and B plot the average proportion of wealth invested in the risk-free bond and Panels C and D plot the average proportion of wealth invested in the less-familiar asset—by less- and more-confident investors. Panel E depicts the *difference* in the average portfolio Sharpe ratios of less- and more-confident investors (Less-conf. − More-conf.). Portfolio returns are annualized and computed under the objective beliefs. Averages are calculated across 100,000 simulation paths. All results are based on the parameter values described in Table 1, with the difference relative to our baseline case highlighted in the legend.

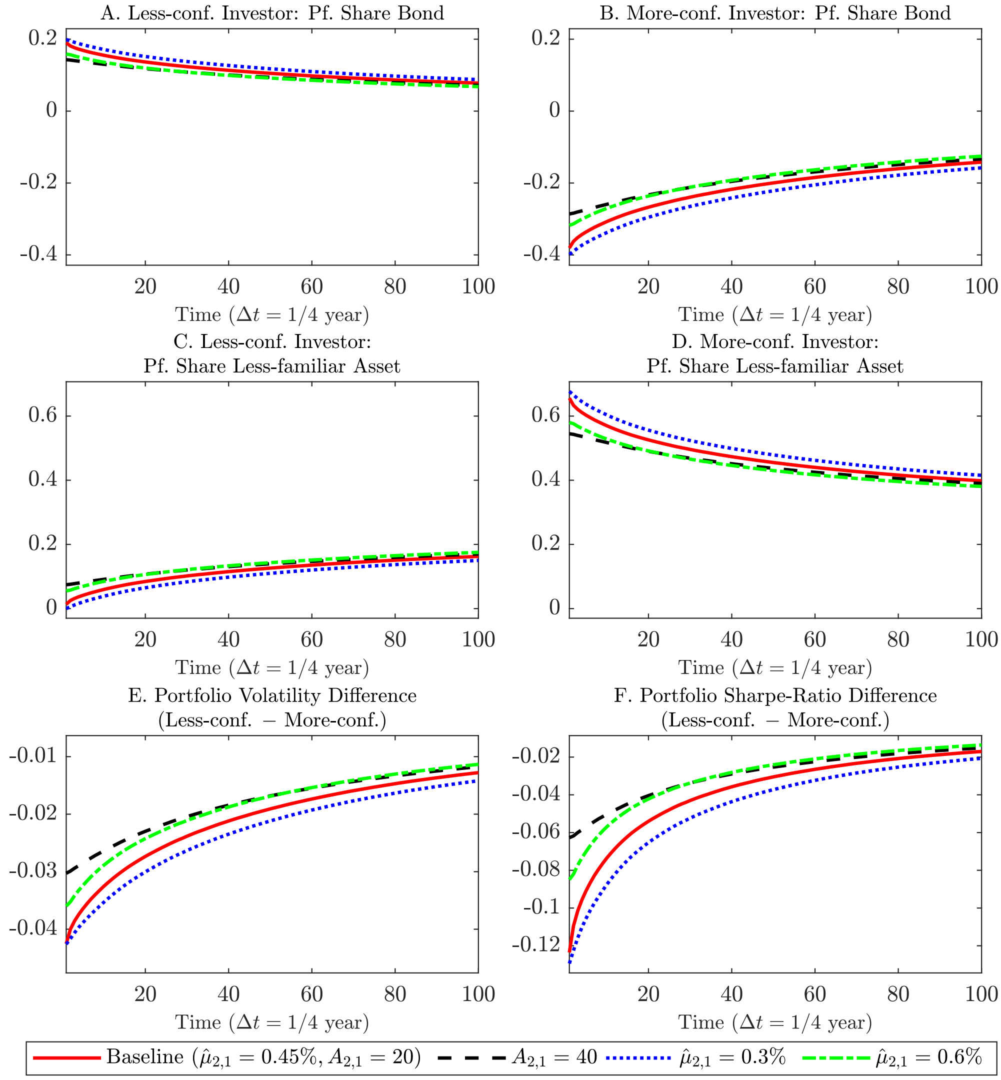

Figure D3: Robustness—Variations in Dividend Dynamics

The figure illustrates how portfolio heterogeneity and investment returns vary with the specifications of the dividend processes. Panels A and B plot the average proportion of wealth invested in the risk-free bond and Panels C and D plot the average proportion of wealth invested in the less-familiar asset—by less- and more-confident investors. Panel E depicts the *difference* in the average portfolio Sharpe ratios of the less- and more-confident investors (Less-conf. − More-conf.). Portfolio returns are annualized and computed under the objective beliefs. Averages are calculated across 100,000 simulation paths. All results are based on the parameter values described in Table 1, with the difference relative to our baseline case highlighted in the legend.

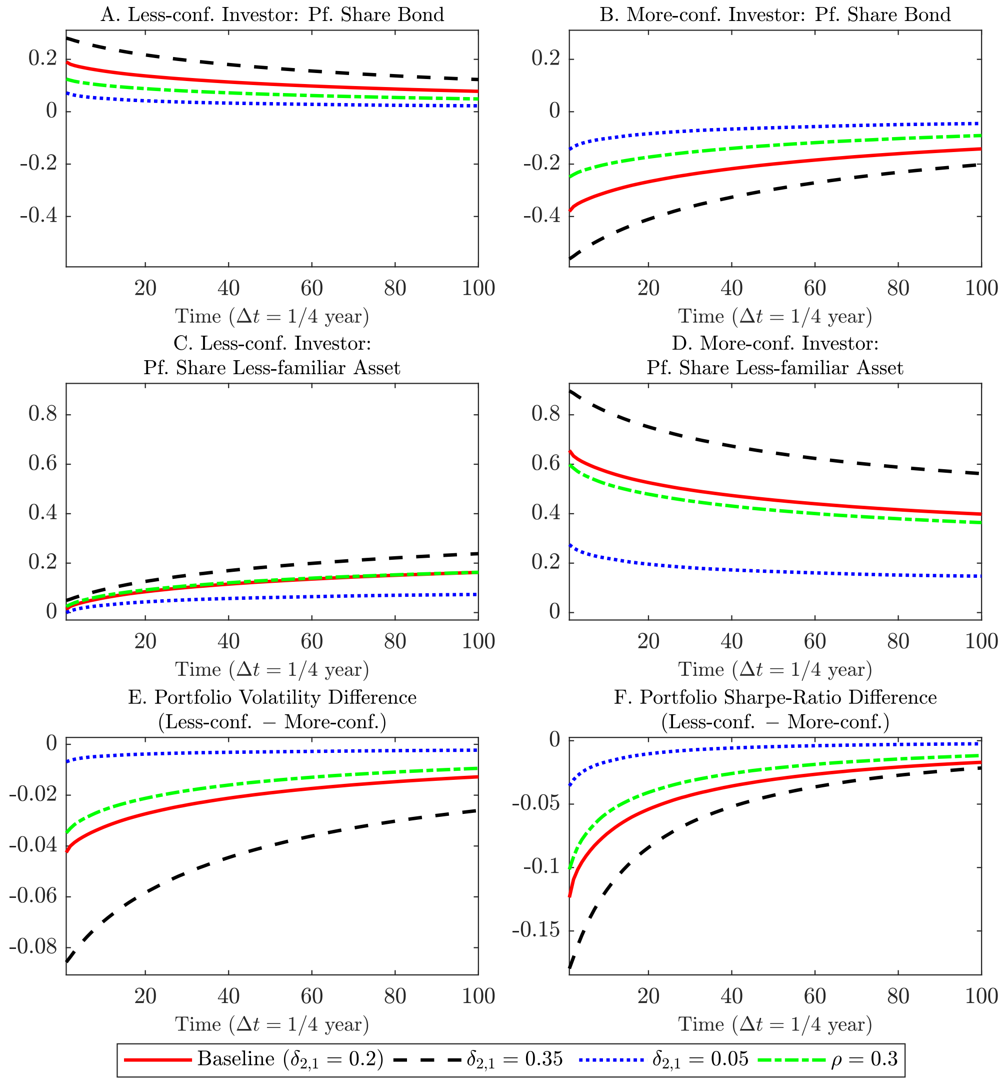

References

Ai, Hengjie, Ravi Bansal, Hongye Guo, and Amir Yaron, 2019, Identifying preference for early resolution from asset market data, Working paper, University of Minnesota.

Baker, Steven D., Burton Hollifield, and Emilio Osambela, 2016, Disagreement, speculation, and aggregate investment, *Journal of Financial Economics* 119, 210–225.

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

Borovička, Jaroslav, 2020, Survival and long-run dynamics with heterogeneous beliefs under recursive preferences, *Journal of Political Economy* 128, 206–251.

Brown, Keith C., Lorenzo Garlappi, and Cristian Tiu, 2010, Asset allocation and portfolio performance: Evidence from university endowment funds, *Journal of Financial Markets* 13, 268–294.

Brunnermeier, Markus, Emmanuel Farhi, Ralph S. J. Koijen, Arvind Krishnamurthy, Sydney C. Ludvigson, Hanno Lustig, Stefan Nagel, and Monika Piazzesi, 2021, Review article: Perspectives on the future of asset pricing, *Review of Financial Studies* 34, 2126–2160.

Campbell, John Y., and Luis M. Viceira, 2002, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors* (Oxford University Press, New York).

Chabakauri, Georgy, 2013, Dynamic equilibrium with two stocks, heterogeneous investors, and portfolio constraints, *Review of Financial Studies* 26, 3104–3141.

Chabakauri, Georgy, 2015, Asset pricing with heterogeneous preferences, beliefs, and portfolio constraints, *Journal of Monetary Economics* 75, 21–34.

Cochrane, John H., Francis Longstaff, and Pedro Santa-Clara, 2008, Two trees, *Review of Financial Studies* 21, 347–385.

Collin-Dufresne, Pierre, Michael Johannes, and Lars Lochstoer, 2016a, Asset pricing when 'this time is different', *Review of Financial Studies* 30, 505–535.

Collin-Dufresne, Pierre, Michael Johannes, and Lars Lochstoer, 2016b, Parameter learning in general equilibrium: The asset pricing implications, *American Economic Review* 106, 664–98.

Dumas, Bernard, Alexander Kurshev, and Raman Uppal, 2009, Equilibrium portfolio strategies in the presence of sentiment risk and excess volatility, *Journal of Finance* 64, 579–629.

Dumas, Bernard, and Andrew Lyasoff, 2012, Incomplete-market equilibria solved recursively on an event tree, *Journal of Finance* 67, 1897–1941.

Ehling, Paul, Alessandro Graniero, and Christian Heyerdahl-Larsen, 2018, Asset prices and portfolio choice with learning from experience, *Review of Economic Studies* 85, 1752–1780.

Epstein, Larry G., and Stanley Zin, 1989, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* 57, 937–969.

Gabaix, Xavier, and Ralph S. J. Koijen, 2021, In search of the origins of financial fluctuations: The inelastic markets hypothesis, Technical report, National Bureau of Economic Research.

Gallmeyer, Michael, and Burton Hollifield, 2008, An examination of heterogeneous beliefs with a short-sale constraint, *Review of Finance* 12, 323–364.

Goetzmann, William N., and Sharon Oster, 2013, Competition among university endowments, in *How the Financial Crisis and Great Recession Affected Higher Education*, 99–126 (University of Chicago Press, Chicago, IL).

Greenwood, Robin, and Stefan Nagel, 2009, Inexperienced investors and bubbles, *Journal of Financial Economics* 93, 239–258.

Johannes, Michael, Lars A. Lochstoer, and Yiqun Mou, 2016, Learning about consumption dynamics, *Journal of Finance* 71, 551–600.

Kimball, Miles S., 1990, Precautionary saving in the small and in the large, *Econometrica* 58, 53–73.

Koijen, Ralph S. J., and Motohiro Yogo, 2019, A demand system approach to asset pricing, *Journal of Political Economy* 127, 1475–1515.

Lerner, Josh, Antoinette Schoar, and Jialan Wang, 2008, Secrets of the academy: The drivers of university endowment success, *Journal of Economic Perspectives* 22, 207–222.

Lucas, Robert E., Jr., 1978, Asset prices in an exchange economy, *Econometrica* 46, 1429–1445.

Merton, Robert C., 1971, Optimum consumption and portfolio rules in a continuous-time model, *Journal of Economic Theory* 3, 373–413.

Morris, Stephen, 1995, The common prior assumption in economic theory, *Economics and Philosophy* 11, 227–253.

Panageas, Stavros, 2005, The neoclassical theory of investment in speculative markets, Working paper, University of Pennsylvania.

Panageas, Stavros, 2019, The implications of heterogeneity and inequality for asset pricing, *Foundations and Trends in Finance* (forthcoming).

Petajisto, Antti, 2009, Why do demand curves for stocks slope down?, *Journal of Financial and Quantitative Analysis* 44, 1013–1044.

Prieto, Rodolfo, 2013, Dynamic equilibrium with heterogeneous agents and risk constraints, Working paper, Boston University.

Scheinkman, Jose A., and Wei Xiong, 2003, Overconfidence and speculative bubbles, *Journal of Political Economy* 111, 1183–1219.

Weil, Philippe, 1990, Nonexpected utility in macroeconomics, *Quarterly Journal of Economics* 105, 29–42.

Xiong, Wei, and Hongjun Yan, 2010, Heterogeneous expectations and bond markets, *Review of Financial Studies* 23, 1433–1466.