DISCUSSION PAPER SERIES

DP15116
(v. 2)

Investor Sophistication and Portfolio Dynamics

Adrian Buss, Raman Uppal and Grigory Vilkov

FINANCIAL ECONOMICS

ISSN 0265-8003

Investor Sophistication and Portfolio Dynamics

Adrian Buss, Raman Uppal and Grigory Vilkov

Discussion Paper DP15116
First Published 31 July 2020
This Revision 04 February 2021

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

Investor Sophistication and Portfolio Dynamics

Abstract

We develop a simple dynamic general-equilibrium framework that can jointly rationalize many empirically observed features of household portfolios, investment returns, and wealth dynamics. The model differs from traditional models only along a single, natural dimension: households differ in their confidence about the return processes for risky assets. Less-confident households (but with unbiased beliefs) overinvest in safe assets, hold underdiversified portfolios concentrated in familiar assets, are trend chasers, and earn lower absolute and risk-adjusted investment returns. More confident households hold riskier positions and exhibit superior market-timing abilities. Despite Bayesian learning, this investment behavior persists for long periods, thereby exacerbating wealth inequality.

JEL Classification: D53, G11, G51, G53

Keywords: household finance, portfolio dynamics, Wealth Inequality, Belief formation, investors' expectations, trend chasing, market timing

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

Acknowledgements
We are grateful for helpful comments from Laurent Bach, Laurent Calvet, Bruce Carlin, Joao Cocco, Bernard Dumas, Lorenzo Garlappi, Francisco Gomes, Luigi Guiso, Michael Haliassos, Bernard Herskovic, Christian Heyerdahl-Larsen, Tullio Jappelli, Samuli Knüpfer, Alex Michaelides, Stavros Panageas, Anna Pavlova, Kim Peijnenburg, Joel Peress, Valery Polkovnichenko, Tarun Ramadorai, Nikolai Roussanov, Antoinette Schoar, Paolo Sodini, Dimitri Vayanos, 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, CEPR Household Finance Seminar Series, China International Conference in Finance, European Finance Association Meetings, 2nd World Symposium on Investment Research, 3rd European Workshop on Household Finance, Financial Intermediation Research Society (FIRS) Conference, Seventh Biennial McGill Global Asset Management Conference, and UBC Summer Finance Conference.

Investor Confidence and Portfolio Dynamics*

Adrian Buss Raman Uppal Grigory Vilkov

February 4, 2021

Abstract

We develop a simple dynamic general-equilibrium framework that can jointly rationalize many empirically observed features of household portfolios, investment returns, and wealth dynamics. The model differs from traditional models only along a single, natural dimension: households differ in their confidence about the return processes for risky assets. Less-confident households (but with unbiased beliefs) overinvest in safe assets, hold underdiversified portfolios concentrated in familiar assets, are trend chasers, and earn lower absolute and risk-adjusted investment returns. More confident households hold riskier positions and exhibit superior market-timing abilities. Despite Bayesian learning, this investment behavior persists for long periods, thereby exacerbating wealth inequality. (*JEL* D53, G11, G51, G53)

Keywords: household finance, portfolio dynamics, wealth inequality, belief formation, investors' expectations, trend chasing, market timing.

*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`. We are grateful for helpful comments from Tony Berrada (discussant), Harjoat Bhamra, Oliver Boguth (discussant), Laurent Bach, Laurent Calvet, Bruce Carlin, Georgy Chabakauri, João Cocco, Victor DeMiguel, Bernard Dumas, Lorenzo Garlappi, Francisco Gomes, Luigi Guiso, Michael Haliassos, Rawley Heimer, Bernard Herskovic (discussant), Christian Heyerdahl-Larsen (discussant), Tullio Jappelli, Samuli Knüpfer, Alex Michaelides, Stavros Panageas, Héctor Calvo Pardo, Anna Pavlova, Kim Peijnenburg, Joel Peress, Valery Polkovnichenko (discussant), Tarun Ramadorai, Thierry Roncalli, Lukas Roth, Nikolai Roussanov, Antoinette Schoar, Paolo Sodini, Dimitri Vayanos, Gyuri Venter (discussant), Pietro Veronesi, Lu Zhang (discussant), 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, CEPR Household Finance Seminar Series, China International Conference in Finance, European Finance Association Meetings, 2nd World Symposium on Investment Research, 3rd European Workshop on Household Finance, Financial Intermediation Research Society (FIRS) Conference, Seventh Biennial McGill Global Asset Management Conference, and UBC Summer Finance Conference.

1 Introduction and Motivation

Recent empirical evidence has documented striking patterns in households' investment choices. One group of households tilts its investments toward safe and familiar assets, trend chases, and earns lower investment returns, while another group of households holds riskier positions and exhibits superior market-timing abilities, consequently earning higher investment returns (Bianchi, 2018, Fagereng, Guiso, Malacrino, and Pistaferri, 2020). These differences in investment behavior persist for remarkably long periods and are a crucial determinant of the dynamics of wealth inequality (Piketty, 2014, Bach, Calvet, and Sodini, 2020, Fagereng et al., 2020). In addition, they have substantial welfare implications for individuals and also for society (Campbell, 2016, Bhamra and Uppal, 2019).

Our main contribution is to develop a novel but simple theoretical framework that can jointly rationalize *both* the static and dynamic properties of households' investment decisions and the resultant wealth dynamics observed empirically. Our model is driven by a single salient feature of financial markets—*differences in investor confidence* about the expected returns of risky assets.[1] Surprisingly, this *single* natural deviation from standard representative-agent models can explain the empirically observed *persistent* heterogeneity in household investments in safe and risky assets, in their portfolio returns, and their wealth dynamics. Just like in the data, in our model households who are less-confident (but whose beliefs on average are unbiased) overinvest in safe assets, hold underdiversified portfolios concentrated in familiar assets, are trend-chasers, and earn lower absolute and risk-adjusted investment returns. On the other hand, more-confident households hold riskier portfolios, time the market (i.e., take on more risk when expected returns are higher), and earn consistently higher investment returns. Notably, these patterns are a consequence of households' optimal choices (rather than behavioral biases or investment mistakes) that persist for long periods, thereby exacerbating wealth inequality.

Specifically, we develop a dynamic general-equilibrium model with two key features. First, in addition to a risk-free asset, there are *two* risky assets, namely a "traditional" risky asset and

[1] Extensive empirical evidence about heterogeneity in the financial capability of investors is provided by Hastings, Madrian, and Skimmyhorn (2013), Lusardi and Mitchell (2014), and Campbell (2016). Differences in confidence could arise from a variety of sources, such as: differences in literacy levels (Bianchi, 2018), sophistication (Fagereng et al., 2020, Calvet, Campbell, and Sodini, 2009b), skill (Bach et al., 2020), experience (Calvet et al., 2009b), financial knowledge (Lusardi, Michaud, and Mitchell, 2017), trust in capital markets (Guiso, Sapienza, and Zingales, 2008), and gender (Bucher-Koenen, Lusardi, Alessie, and Van Rooij, 2017).

a "less-familiar" risky asset, each modeled as a claim to a Lucas (1978) tree, with constant mean and volatility of dividend growth. Second, the model has *two* classes of households, both of which have a preference for early resolution of uncertainty, but differ in confidence, with some households having lower (initial) confidence about the less-familiar asset's dividend-growth rate. Intuitively, one can think of the traditional risky asset as (a portfolio of) assets that all households understand well (e.g., because their fundamentals are rather stable). In contrast, the less-familiar asset can be interpreted as (a portfolio of) assets whose fundamentals are more difficult to assess; for example, stocks from new sectors such as technology and biotechnology, international stocks, or a less-traditional asset class, such as real estate or commodities. Formally, 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 less-familiar asset's return distribution. Otherwise, the model is standard and kept as simple as possible to highlight the principal economic forces.

Many empirically observed patterns in households' (static and dynamic) investment choices arise naturally in the model. First, there is substantial heterogeneity in households' wealth allocation between safe and risky assets and also in the composition of their risky-asset holdings. In particular, less-confident households allocate more wealth to the risk-free asset, and their portfolio of risky assets is concentrated in familiar assets. In contrast, more-confident households overweight the less-familiar asset. These asset demands are a consequence only of differences in confidence about the less-familiar asset's return, which—coupled with Bayesian updating—gives rise to precautionary-savings (Kimball, 1990) and intertemporal-hedging demands (Merton, 1971, Campbell and Viceira, 2002). 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 households 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).

Second, there is strong *persistence* in the heterogeneity in asset demands, despite Bayesian learning by households. 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.

Third, our model generates predictions regarding *changes* in asset demands *conditional* on realized and expected returns. Less-confident households 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 households take on riskier positions when expected market returns are higher; that is, they have superior market-timing abilities.

Fourth, less-confident households earn lower investment returns than more-confident ones—in absolute terms but 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 fact that more-confident households take on riskier positions exactly when risk premia are higher, pushing up their investment returns.

Fifth, less-confident households' financial wealth increases more slowly than that of more-confident households, and hence, wealth inequality worsens over time—a natural result of the heterogeneity in investment returns. It is important to highlight that this worsening of wealth inequality is a consequence of a new channel that has not been studied before—limited confidence about the less-familiar asset's return.

Our model also has implications 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 traditional asset 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

[2]With a preference for early resolution of uncertainty, less-confident households 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).

is required to induce the more-confident households to hold a larger share of this asset—to compensate for the low demand from less-confident households.

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 households 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., with households 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 heterogeneous beliefs, which generally do not match the empirical observations regarding the dynamics of portfolio holdings and investment returns.

The model also allows us to obtain novel policy-relevant insights on the relation between investor confidence and wealth inequality. In particular, it highlights that to address wealth inequality it is important to improve investors' confidence, for example, through financial education, financial innovation, or financial technology.

Our model is motivated by the broad features of households' portfolio holdings and investment returns documented in the extensive empirical literature on household finance. In particular, the literature has documented substantial and persistent heterogeneity in asset demands and investment returns, which contributes significantly to wealth inequality; see, for instance, Campbell (2006), Calvet, Campbell, and Sodini (2007), Goetzmann and Kumar (2008), Calvet, Campbell, and Sodini (2009a), Calvet et al. (2009b), Campbell, Ramadorai, and Ranish (2015), Clark, Lusardi, and Mitchell (2017), Von Gaudecker (2015), Bianchi (2018), Bach et al. (2020), Fagereng et al. (2020).[3]

[3]While the focus of our paper is on households, our findings are consistent also with the empirical evidence for the behavior of institutional investors. For example, there is ample empirical evidence regarding heterogeneity in asset demands by institutions (Koijen and Yogo, 2019) and for persistent heterogeneity in the holdings and portfolio returns of university endowment funds (Lerner, Schoar, and Wang, 2008, Brown, Garlappi, and Tiu, 2010, Goetzmann and Oster, 2013). Moreover, less-confident institutional investors also exhibit "trend-chasing behavior" (Greenwood and Nagel, 2009, Goetzmann and Oster, 2013).

Our work is closest to the theoretical literature on households' investment choices. Campbell, Cocco, Gomes, Maenhout, and Viceira (2001), Gomes and Michaelides (2003, 2005), Cocco, Gomes, and Maenhout (2005), and Calvet, Celerier, Sodini, and Vallee (2020) study, in partial-equilibrium, how life-cycle considerations influence the asset-demand decisions of households. Lusardi et al. (2017) study a partial-equilibrium model in which individuals endogenously choose their investment in financial knowledge and show that differences in financial literacy are a key determinant of wealth inequality. Gomes and Michaelides (2008) develop a general-equilibrium model with incomplete risk sharing that can simultaneously explain a high equity premium and a low stock-market demand (limited participation).[4] In contrast to these models in which asset demands depend on information frictions or participation costs, we identify a new channel—confidence heterogeneity—to explain the empirically observed heterogeneity in the dynamics of asset demands, investment returns, and wealth inequality.

Our work is also related to the literature studying models with investors who have heterogeneous beliefs. Collin-Dufresne, Johannes, and Lochstoer (2016a), Ehling, Graniero, and Heyerdahl-Larsen (2018), and Malmendier, Pouzo, and Vanasco (2020) consider belief heterogeneity resulting from investors' "experience" and focus on the link between investor demographics (generations) and asset returns. For instance, both Collin-Dufresne et al. (2016a) and Ehling et al. (2018) study a single-asset equilibrium model with a "generational learning bias," in which young generations, when born, are endowed with more disperse beliefs than the dying old generation. Consistent with our paper, they report a positive link between past returns and future return forecasts. Nagel and Xu (2019), on the other hand, study asset prices in an economy where a representative agent learns with "fading memory." The key difference with these papers is that in our framework, learning is Bayesian, and our focus is on the dynamics of asset demands for multiple risky assets and the resulting *heterogeneity* in the dynamics of *portfolio returns and wealth* across investors, rather than on the asset-pricing implications of biased learning.[5]

[4] Comprehensive surveys of household finance are presented by Guiso, Haliassos, and Jappelli (2002), Haliassos (2003), Campbell (2006), Guiso and Sodini (2013), and Gomes, Haliassos, and Ramadorai (2020).

[5] 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.

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).[6] In Kacperczyk, Nosal, and Stevens (2019), wealth inequality arises for a distinctly different reason than in our model: improvement in aggregate information technology, which disproportionately benefits wealthy investors.

The rest of the paper is organized as follows. Section 2 introduces our economic framework. We show how heterogeneity in households' confidence affects portfolio dynamics in Section 3, the dynamics of asset returns in Section 4, and households' portfolio returns, wealth inequality, and welfare in Section 5. Section 6 contrasts our predictions with those from alternative model specifications. Section 7 concludes. Technical details are relegated to the appendices.

2 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.

2.1 The Model

We study a general-equilibrium model set in discrete time, with time interval Δt and a finite horizon T.[7]

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 $\Delta d_{n,t+1} \equiv \ln\left[D_{n,t+1}/D_{n,t}\right]$ is described by an IID-Normal model with expected dividend-growth rate μ_n and dividend-growth

[6] Panageas (2019) provides an excellent review of the literature on belief and preference heterogeneity.

[7] 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).

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.[8] We interpret the first risky asset as any portfolio of "familiar" (traditional) assets that are well understood by all households. The second risky asset represents the less-familiar asset (portfolio).

Households: The economy is populated by two groups of households, indexed by $k \in \{1, 2\}$. Households 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 household 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 households differ only in their level of confidence about the cash-flow dynamics of the less-familiar asset.[9] Intuitively, such heterogeneity could arise from differences in households' literacy levels (Bianchi, 2018), skill (Bach et al., 2020), experience (Calvet et al., 2009b), financial knowledge (Lusardi et al., 2017), trust in capital markets (Guiso et al., 2008), and gender (Bucher-Koenen et al., 2017).[10]

We model households' beliefs using parameter uncertainty combined with Bayesian updating (based on cash-flow news), which yields rich dynamics in households' subjective beliefs regarding the distribution of future cash flows of the less-familiar asset.[11] In particular, household k starts

[8] 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.

[9] 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 6.2.

[10] One might think that less-confident households could rely on financial advisers. However, Lusardi and Mitchell (2014) find that fewer than one-third of the respondents in the U.S. National Financial Capability Study consulted financial advisers. Intuitively, when households lack confidence, it is then also difficult for them to evaluate the financial adviser: Quis custodiet ipsos custodes?

[11] 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, Basak (2005) and Dumas, Kurshev, and Uppal (2009), and the papers cited therein.

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

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

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

Consequently, even though the dividend dynamics of the less-familiar asset are driven by an IID model with constant parameters, from the households' 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 households' initial precision, $A_{k,1}$, leads to (long-term) differences in households' confidence (and beliefs).

For expositional ease, we assume that there is no uncertainty about the parameters of the traditional asset's dividend process, and households 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 households' beliefs in Equations (3) and (4) that will be useful for understanding the portfolio dynamics. First, the perceived dividend-growth rate, $\hat{\mu}_{k,t}$, is a martingale, so revisions in subjective beliefs constitute permanent shocks to households' 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

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

The figure illustrates the interplay between households' 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.

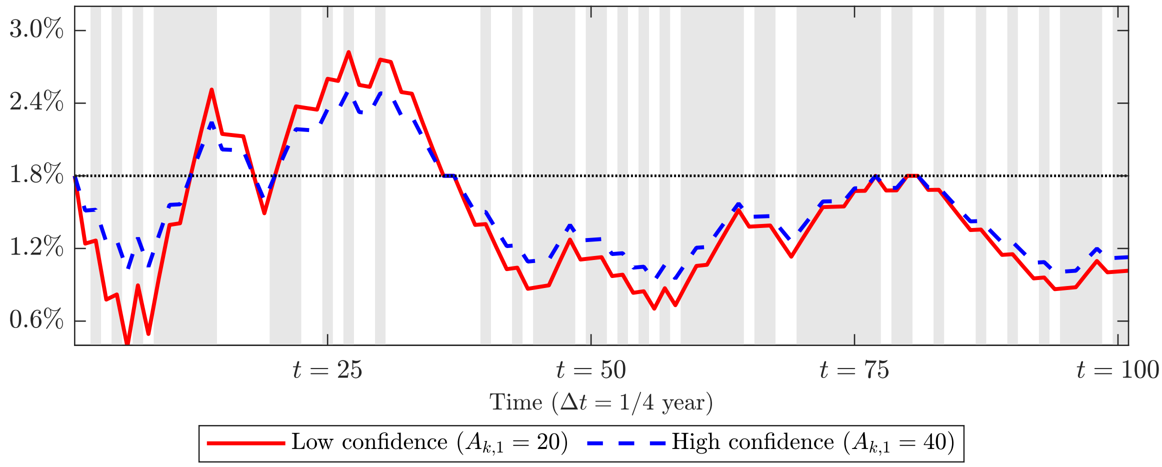

$\frac{A_{k,t-1}}{1+A_{k,t-1}} = \frac{A_{k,1}}{1+t\,A_{k,1}}$. The change in this expression is negative

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

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 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, households revise their beliefs about μ_2 upward and vice versa for lower-than-expected growth. As a result, households 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 households 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, households' estimates of the expected dividend-growth rate become more precise, and, thus, revisions in their beliefs decrease in size.

2.2 Households' Optimization Problem and Equilibrium

The objective of households 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

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

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 portfolio positions in 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: $\theta_{k,2,t} \geq 0, k \in \{1, 2\}$.[12]

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

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

where $M_{k,t+1}$ is household 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 $\xi_{k,t} \equiv E_t^k\left[\exp\big((1-\gamma)\, v_{k,t+1}\big)\right]^{(\gamma - 1/\psi)/(1-\gamma)}$.

[12] We constrain short sales because, in practice, it is not easy to short a risky asset, especially one with which households are not very familiar. Qualitatively, the constraint does not affect any of our results; quantitatively, the results are slightly stronger if we allow for short sales (see Section 6.3). For the traditional asset, the constraint is never binding (because there is no disagreement) and, hence, is omitted.

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, through its impact on households' 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 household 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 households, $\omega_{2,t} \in (0,1)$; 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 households, $\hat{\mu}_{k,t}$, $k \in \{1,2\}$, with its dynamics specified in (3); and the (deterministic) posterior variances of the beliefs of the two groups of households, $A_{k,t}\,\sigma_2^2$, $k \in \{1,2\}$, with the dynamics specified in (4).

2.3 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 households prefer early resolution of uncertainty, that is, $\psi > 1/\gamma$.[13] In Section 6 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

[13]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 households	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 households' prior distribution	0.45%
$A_{2,1}$	Initial precision of less-confident households' prior distribution	20
$\underline{\mu}_2, \bar{\mu}_2$	Truncation boundaries for beliefs of less-confident households	$[-0.55\%, 1.45\%]$

(less-familiar) asset's dividends, $\delta_{2,1}$, is set to a conservative value of 0.20. We specify identical processes for the two assets' dividends, calibrated to match jointly the historical mean and volatility of aggregate U.S. consumption growth (1.80% and 3.70% p.a., respectively).[14] Because most risky assets are 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 households knows the dividend-growth rate of the less-familiar asset (i.e., has infinite precision).[15,16] 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

[14]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.

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

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

as "more-confident households" and the second group ($k = 2$) as "less-confident households." 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).[17] Finally, we assume that less-confident households are initially endowed with 2/3 of the total wealth.[18]

Identifying the equilibrium is a non-trivial task. In particular, we extend the numerical solution approach proposed by Dumas and Lyasoff (2012) along several dimensions; for instance, we incorporate parameter uncertainty with Bayesian learning, multiple risky assets, and Epstein-Zin-Weil preferences.[19] 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 households have full confidence, which serves as a benchmark.

3 Portfolio Dynamics

In this section, we explain how heterogeneity in investor confidence gives rise to distinct portfolio patterns as observed in the data, including persistent portfolio heterogeneity and trend-chasing behavior.

[17]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.

[18]In Section 6.3, we demonstrate how our findings vary quantitatively with the truncation bounds and the initial wealth distribution.

[19]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.

3.1 Portfolio Dynamics

Figure 2 plots the proportion of wealth each household 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 households' allocations to the risk-free asset differ substantially. Less-confident households 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 households have, because of market-clearing, a short position in the safe asset.

Second, also the risky-asset components of households' portfolios differ markedly. Specifically, less-confident households persistently underweight the less-familiar asset (Panel C) and overweight the traditional asset (Panel E)—compared to more-confident households and also to the full-confidence market weights. As a result, less-confident households' portfolios are highly concentrated in the traditional risky asset, implying a loss in diversification. Market clearing implies that more-confident households are overinvested in the less-familiar asset (Panel D)—relative to less-confident households and relative also to the setting where all households are fully confident. The more-confident households' short position in the bond (Panel B), however, allows them to maintain their investment in the traditional asset 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.

The key economic force explaining this heterogeneity in households' risky-asset portfolios 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 households 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 households' portfolio holdings in the less-familiar asset into two components: (1) a myopic component and (2) an intertemporal-hedging component. Observe that the less-confident households' myopic demand for the less-familiar asset is large and positive (Panel A) and very similar to the myopic demand of the more-confident households

Figure 2: Portfolio Dynamics

The figure illustrates the portfolio shares of less-confident households (left column) and more-confident households (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, and Panels E and F in the traditional asset. 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.

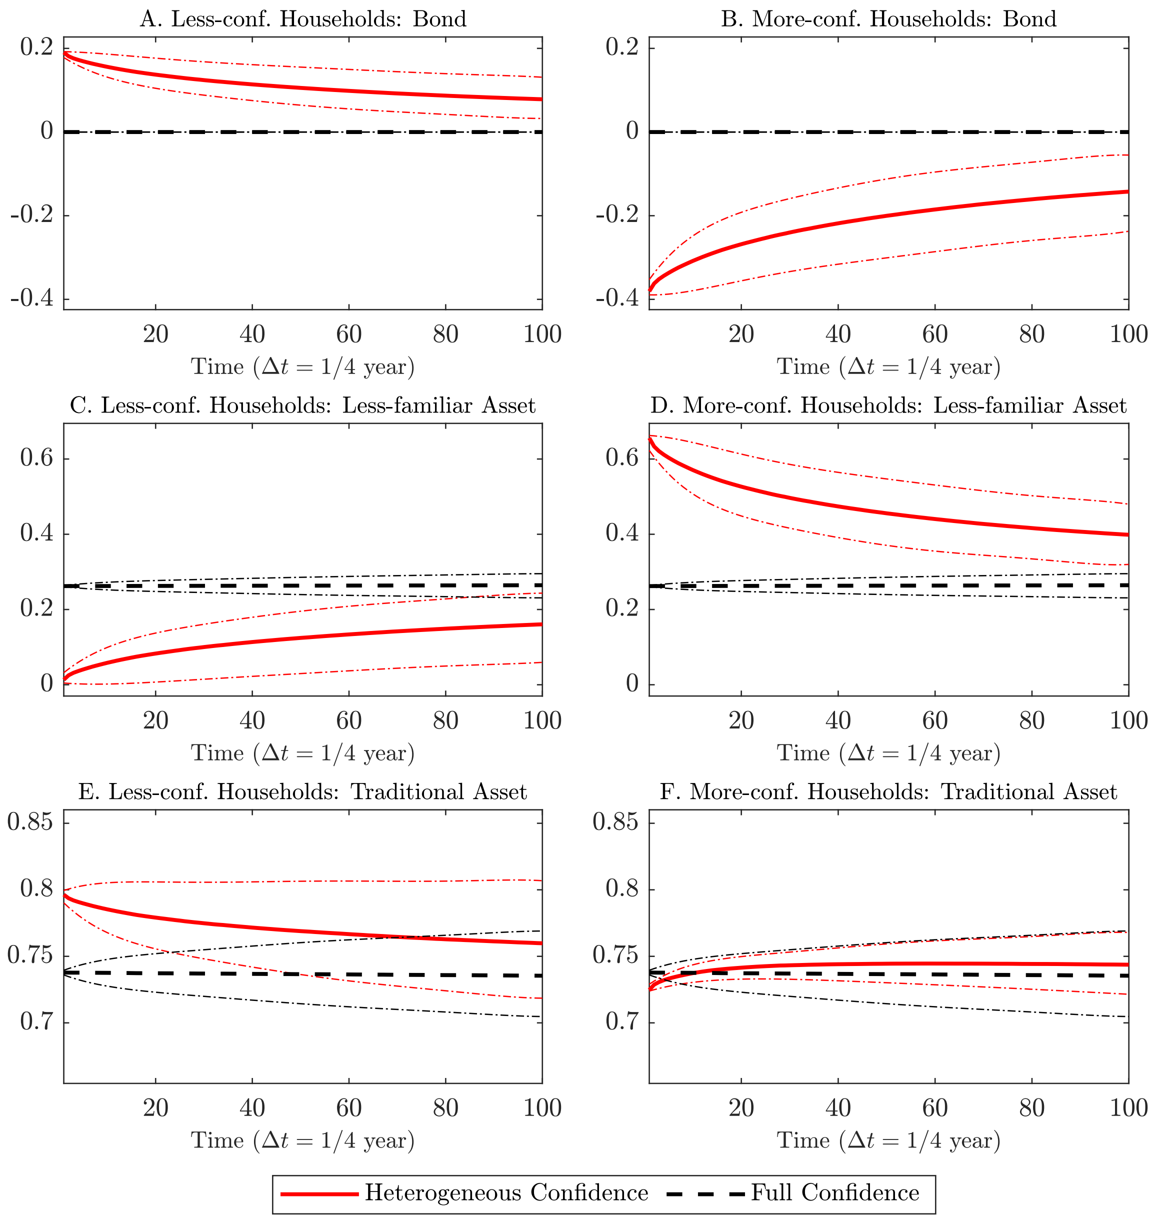

Figure 3: Less-familiar Asset: Portfolio-Share Decomposition

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 households (Panel A) and more-confident households (Panel B). Averages are calculated across 100,000 simulation paths. All graphs are based on the parameter values described in Table 1.

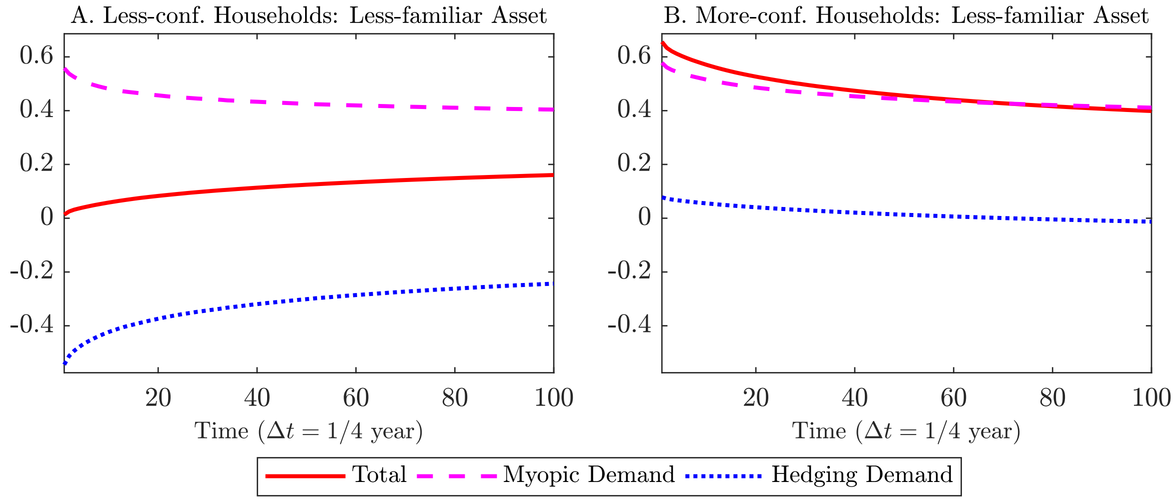

(Panel B).[20] However, there is a striking difference in the intertemporal-hedging demands of the two groups of households. While for less-confident households the intertemporal-hedging demand is strongly negative, more-confident households' intertemporal-hedging demand is quite small, arising mostly from (indirect) changes in their expected utility induced by changes in the beliefs of *less*-confident households.[21]

To understand the intuition for the less-confident households' 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 households' (long-term) consumption dynamics. Formally, one can decompose shocks to each household's subjective log-continuation utility (normalized by total output), $\hat{v}_{k,t+1}$, into the sensitivity of their utility with respect to the stochastic state variables

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

[21] More-confident households 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 households 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.

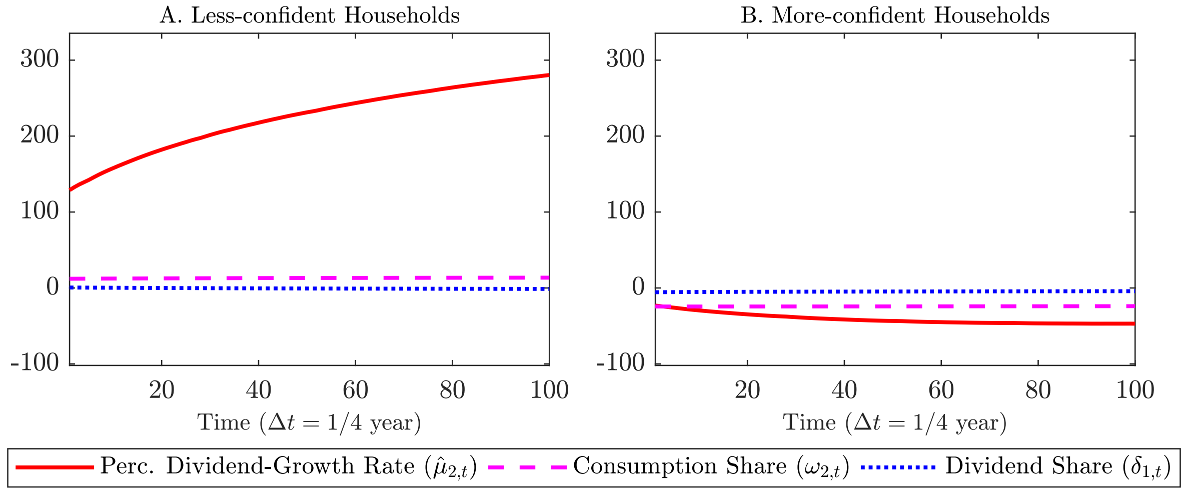

and the shocks to these state variables:

$$\hat{v}_{k,t+1} - E_t^k[\hat{v}_{k,t+1}] = \tag{10}$$
$$\left[\frac{\partial \hat{v}_{k,t}}{\partial \hat{\mu}_{2,t}}\left(\hat{\mu}_{2,t+1} - E_t^k[\hat{\mu}_{2,t+1}]\right) + \frac{\partial \hat{v}_{k,t}}{\partial \omega_{2,t}}\left(\omega_{2,t+1} - E_t^k[\omega_{2,t+1}]\right) + \frac{\partial \hat{v}_{k,t}}{\partial \delta_{1,t}}\left(\delta_{1,t+1} - E_t^k[\delta_{1,t+1}]\right)\right],$$

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 households, less-confident households' share of aggregate consumption, and the traditional asset's share of total dividends. Panel A of Figure 4 shows that the sensitivity of the less-confident households' 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 households' log-continuation utility to changes in beliefs of the less-confident households is much smaller and of opposite sign (Panel B of Figure 4).

In particular, because the sensitivity of less-confident households' 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 households' utility is low.

As is apparent from Figure 2, portfolio heterogeneity persists for long periods and weakens only gradually over time. These gradual changes are driven by less-confident households' learning. Specifically, as their confidence increases, the perceived dividend-growth rate settles down. Thus, their precautionary-savings motive weakens, and hence, the share of wealth allocated to the bond declines. Moreover, the revisions in less-confident households' 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 households' beliefs becomes smaller, as explained in Equations (5) and (6).

In summary, our model can simultaneously explain many of the patterns in households' allocations to safe and risky assets documented empirically in the large literature on household finance (Campbell, 2006, 2016, Calvet et al., 2007, Goetzmann and Kumar, 2008, Campbell et al., 2015, Clark et al., 2017, Von Gaudecker, 2015, Bianchi, 2018, Bach et al., 2020, Fagereng et al., 2020). Moreover, the portfolio 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.

Note also that the same economic mechanism described above would arise in a model with a single risky asset, say the "market portfolio," and heterogeneity in confidence about its cash-flow dynamics. In this case, less-confident households would have a negative intertemporal-hedging demand for the market portfolio. Thus, it could also help explain the low, and only very gradually increasing, stock-market participation of households observed empirically (Campbell, 2006), complementing explanations based on information frictions and participation costs (see, e.g., Gomes and Michaelides 2008).

Figure 5: Trend Chasing in the Demand for the Less-familiar Asset

Panel A plots the average change in the holdings of less-confident households (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 households 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.

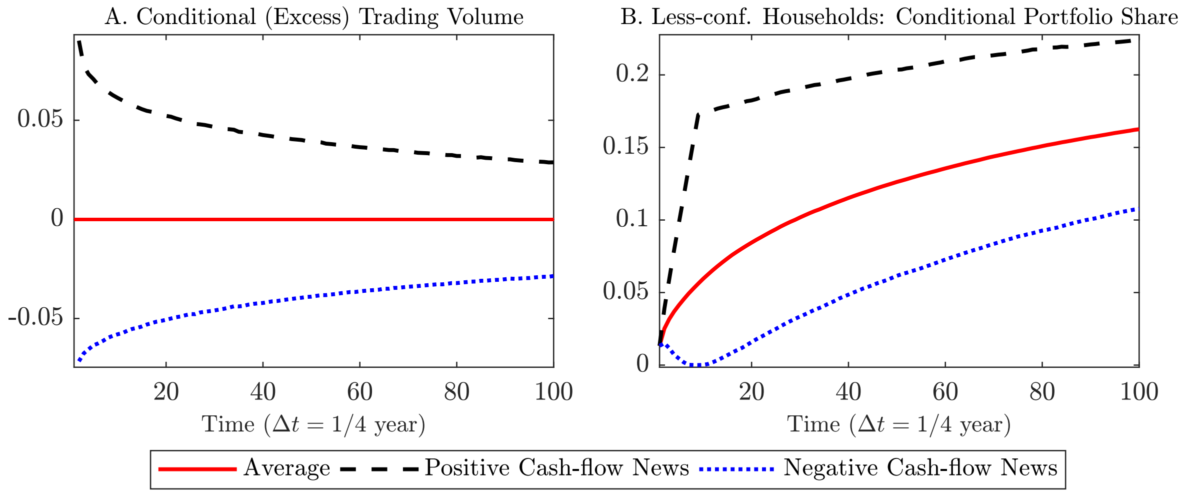

3.2 Trend-Chasing

The revisions in less-confident households' 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 households' current demand, as illustrated in Panel A of Figure 5. This trend chasing is consistent with the empirical data for households; see, among others, Benartzi (2001), Goetzmann and Kumar (2008), and Bianchi (2018).

Revisions in households' beliefs can also explain why less-confident households occasionally allocate substantial fractions of their wealth to assets that they are, in fact, not very confident about, as was the case, for example, for dot-com stocks. Intuitively, because less-confident households' beliefs are not very precise in the early periods, shocks have a substantial impact on the perceived dividend-growth rate. Accordingly, less-confident households 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 5, which plots the less-familiar asset's share in less-confident households' 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).

4 Dynamics of Asset Returns

Households' portfolio returns are driven by both their portfolio weights and the asset returns. We discussed the choice of portfolio weights in the preceding section. In this section, we briefly discuss the equilibrium implications of investor confidence for the short- and long-run dynamics of asset returns, which will help us to understand investment returns and wealth-inequality dynamics in the next section.[22]

4.1 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 households' stochastic discount factors (SDFs). Figure 6 illustrates the dynamics of the volatility of households' SDFs. Notably, with heterogeneity in confidence, the volatilities of the SDFs of *both* groups of households 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 households 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

[22]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 6: Stochastic Discount Factors

Panels A and B show the average conditional volatility of the stochastic discount factor for the less- and more-confident households, 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.

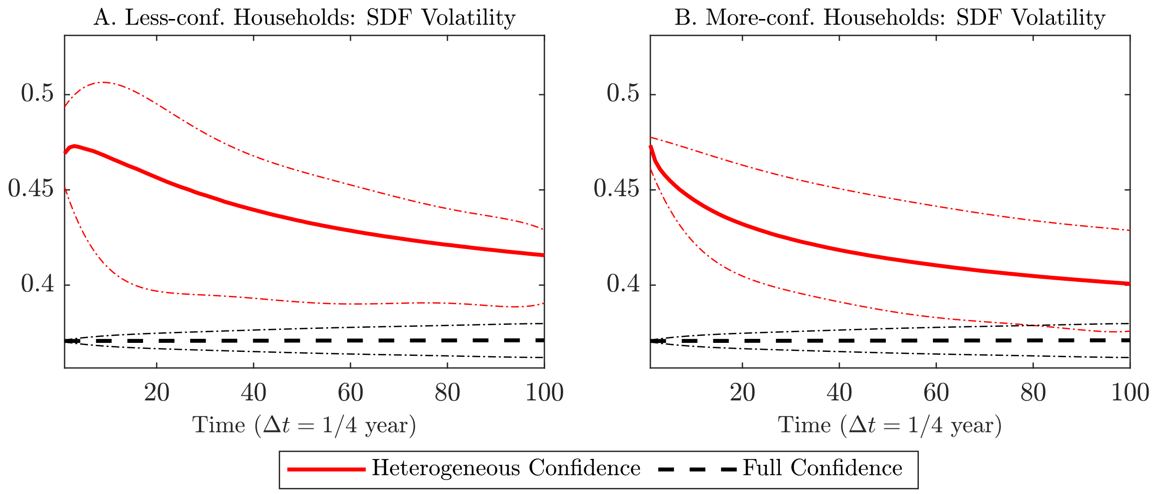

models) but also by shocks to the (forward-looking) subjective log-continuation utility $v_{k,t+1}$. As shown in Equation (10) above, the second term on the right-hand side of (11) can be further decomposed into (i) the sensitivity of households' subjective log-continuation utility with respect to the stochastic state variables and (ii) shocks to these state variables.

For less-confident households, the increase in their SDF's volatility in the first few periods (Panel A of Figure 6) comes from the continuation-utility component of the SDF. In particular, 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 households' SDF.

In contrast, the increase in the volatility of more-confident households' SDF (Panel B of Figure 6) results from higher consumption-growth volatility, which in turn results from the higher return volatility of their portfolios. In particular, because more-confident households, 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 households' 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 households' holding of the less-familiar asset and their short position in the bond, which explains the decline in the SDF volatility.

4.2 Asset Returns

Figure 7 illustrates the dynamics of the return moments of the two risky assets. In the presence of heterogeneity in confidence, the less-familiar asset's return volatility 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."[23,24] Over time, as the precision of less-confident households' 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 households' SDFs, also explains the increase in the less-familiar asset's risk premium, 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 households to hold a larger proportion of their wealth in the less-familiar asset, to compensate

[23]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).

[24]Note that while confidence heterogeneity increases the response of the less-familiar asset's price-dividend ratio to news (positive and negative), it has a negligible impact on its return autocorrelation (compared to the full-confidence case). Hence, there is neither momentum nor reversal in returns.

Figure 7: Dynamics of Asset Returns

The figure illustrates the dynamics of the return moments of the less-familiar and traditional risky asset. 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.

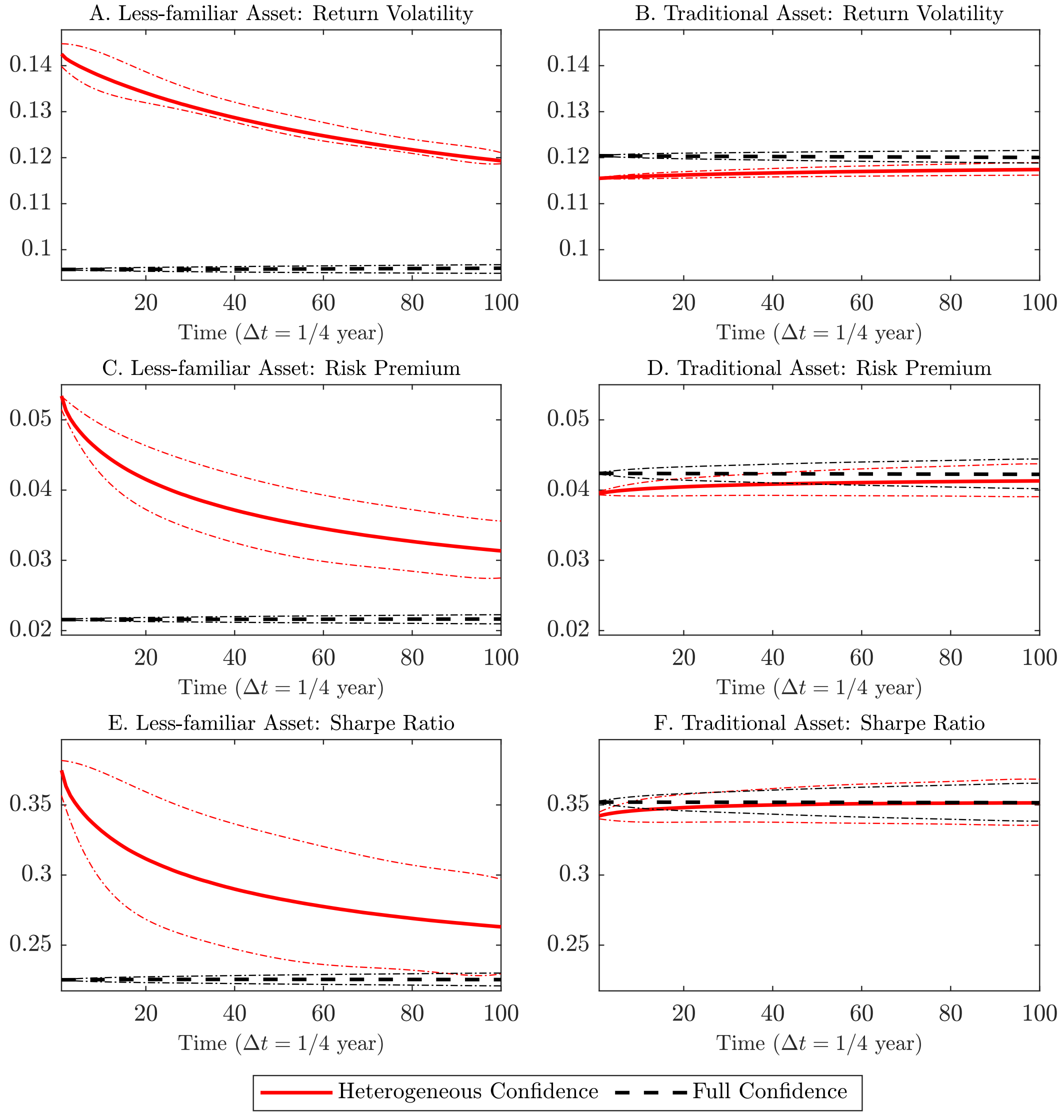

for the lower demand from less-confident households. The Sharpe ratio also declines only slowly over time.

Thus, the less-familiar asset's return volatility, risk premium, and Sharpe ratio are all endogenously higher than those of the traditional asset in the first few years (Panels B, D,

Figure 8: 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.

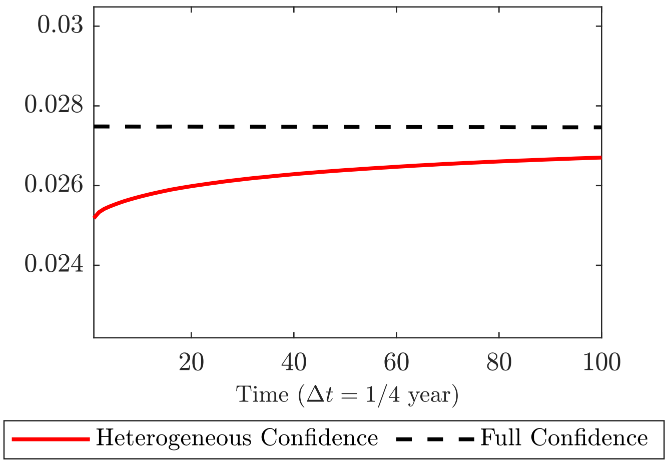

and F). These results are in sharp contrast to those from models with full confidence, in which the return volatility, risk premium, and Sharpe ratio of the less-familiar asset would be *lower* because of this asset's smaller share of aggregate dividends, and hence, smaller covariance with aggregate 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 households' 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 initially, as illustrated in Figure 8. Over time, as less-confident households gain confidence, their precautionary-savings demand declines, and, hence, the risk-free rate slowly moves toward the level in the full-confidence case.

5 Investment Returns, Wealth Inequality, and Welfare

We now describe the dynamics of households' portfolio returns and how they affect the evolution of wealth inequality and welfare.

5.1 Heterogeneity in Investment Returns

Figure 9 illustrates how the expected investment returns and volatilities differ across the two groups of households. Specifically, Panel A shows that less-confident households earn—on average—substantially lower absolute portfolio returns than more-confident households (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 households' 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 households, they invest only a small fraction of their risky-asset-portfolio in the less-familiar asset, thus benefiting less from its (initially) higher risk premium.

Because of the less-confident households' larger allocation to the safe asset, their investment returns are, on average, also considerably less volatile than those of more-confident households, 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 households' lower demand for the less-familiar asset, which initially is much more volatile than the traditional asset.

Notably, however, even adjusted for portfolio risk, less-confident households' investment returns are considerably lower than those of more-confident households, 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 households' 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 households' expected portfolio excess return declines disproportionately more than portfolio volatility.

The heterogeneity in investment returns between less- and more-confident households shows substantial persistence; that is, the difference in households' (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 households slowly become more confident, thus reducing their allocation to the safe asset and improving

Figure 9: Dynamics of Heterogeneity in Portfolio Returns

The figure illustrates the dynamics of the differences in households' 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 households (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 households 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.

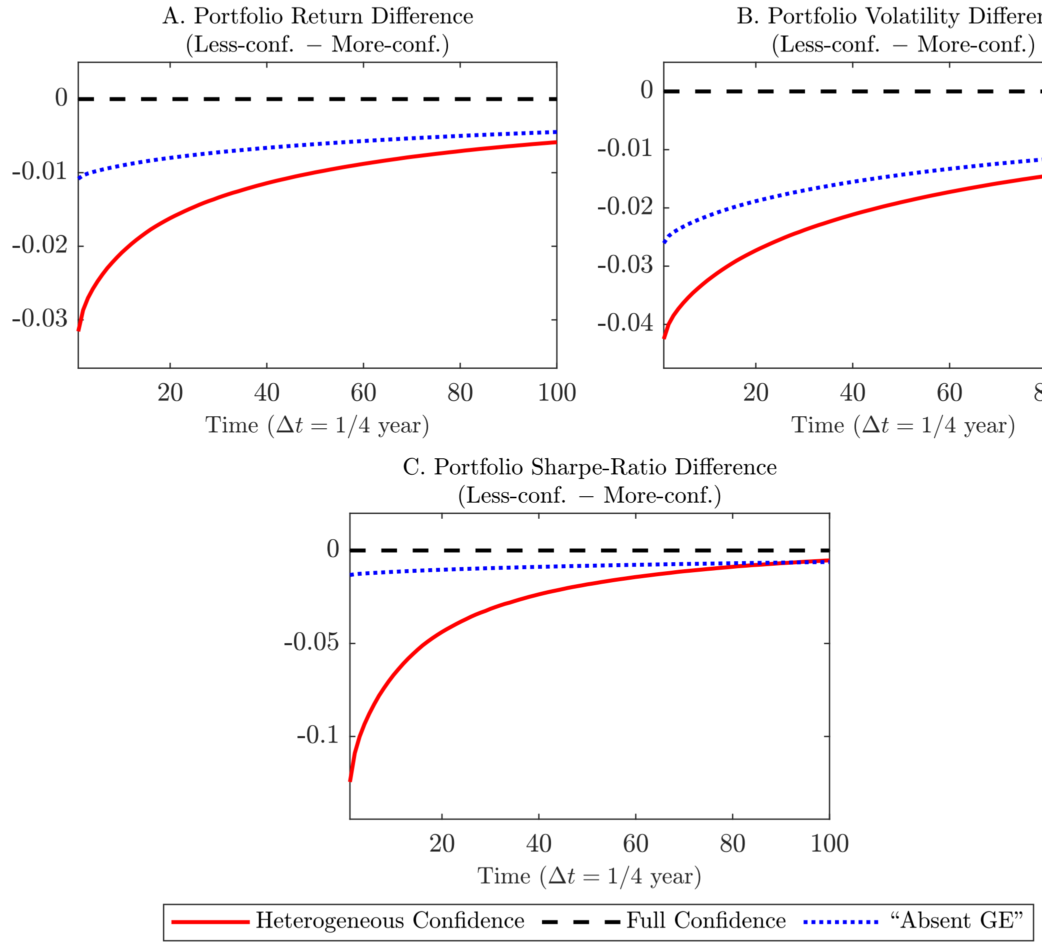

their portfolio's diversification, do their raw and risk-adjusted investment returns improve (in absolute terms but, more importantly, also relative to those of more-confident households).

It is important to highlight that the heterogeneity in households' 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 households to be less than fully confident about the returns of the less-familiar asset (see "Absent GE" in Figure 9). 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 the empirical literature (Campbell, 2006, Grinblatt, Keloharju, and Linnainmaa, 2011, Clark et al., 2017, Von Gaudecker, 2015, Bianchi, 2018, Bach et al., 2020, Fagereng et al., 2020). In particular, there is substantial heterogeneity in households' portfolio returns, with less-confident households earning, on average, substantially lower (risk-adjusted) returns compared to more-confident households. Like in the data, this is a consequence of less-confident households' larger allocation to the safe asset and their smaller allocation to riskier assets that have higher expected returns and their under-diversification. Moreover, these differences in investment returns persist for long periods.

5.2 Asset Demands Conditional on Asset Returns

The dynamic nature of our model also allows us to study households' risk-taking behavior *conditional* on asset returns. Bianchi (2018) documents empirically that more-confident households have riskier positions exactly when expected returns are higher. Figure 10 shows that the same pattern arises in our framework; that is, during periods with a high aggregate risk premium, more-confident households choose a considerably more volatile portfolio (Panel B), compared to periods with a low aggregate risk premium. In contrast, the volatility of less-confident households' 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 households' 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 households 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 households' portfolio volatility mostly unchanged. In response to the asset demands of less-confident households, market clearing requires more-

Figure 10: Portfolio Volatility and Asset Returns

The figure illustrates the relation between households' risk taking and asset returns. In particular, Panels A and B, respectively, plot the less- and more-confident households' 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.

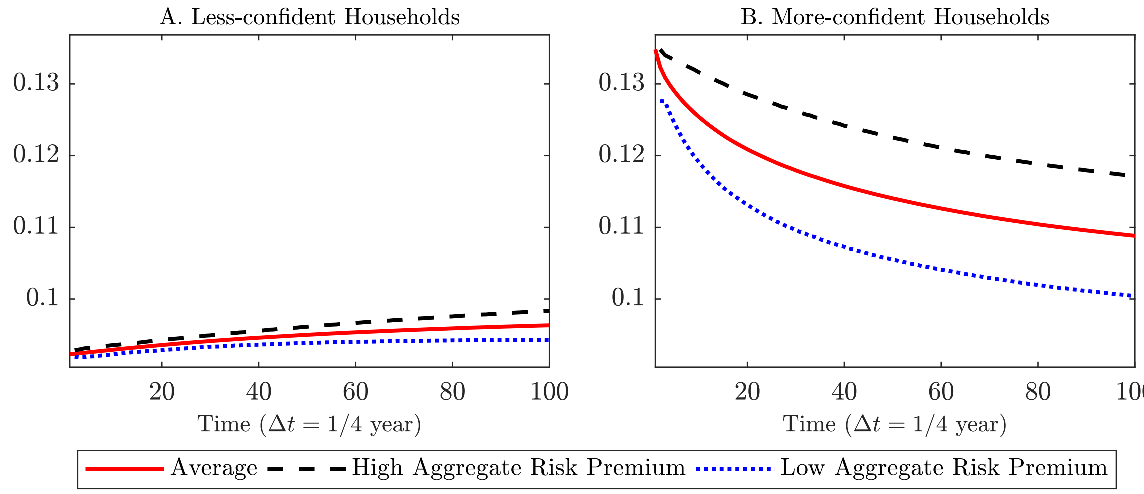

confident households to reduce their allocation to the safe asset and increase their holding of the less-familiar asset, *unambiguously* driving up their portfolio volatility. At the same time, the risk premium on the aggregate market goes up. That is, the lower demand of the less-confident households 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 traditional asset (whose expected return declines only marginally because of a slightly stronger demand).

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

5.3 Wealth Inequality and Welfare

Heterogeneity in households' portfolio returns plays a critical role in determining the dynamics of wealth inequality in the model and also empirically. Specifically, in the presence of confidence heterogeneity, more-confident households accumulate financial wealth at a considerably faster rate—because of their higher investment returns. Consequently, as illustrated in Panels A and B of Figure 11, the average share of aggregate financial wealth held by more-confident households increases over time, whereas that of less-confident households declines.[25] In other words, there is an endogenous worsening of wealth inequality over time.

In particular, note that, recently, Benhabib, Bisin, and Zhu (2011), Benhabib, Bisin, and Luo (2019), and Gabaix, Lasry, Lions, and Moll (2016) have provided theoretical models that, based on the assumption that individuals are exogenously endowed with persistent, idiosyncratic returns to wealth, can generate a wealth distribution with a thick right tail—as in the data. In our model, persistent heterogeneity in investment returns arises endogenously—potentially delivering a micro-foundation for the assumption of persistently different portfolio returns in the aforementioned papers.

The divergence in wealth shares is most pronounced in the early years when households' investment returns differ the most (cf. Figure 9). It slows down only in later years when less-confident households start to reduce their allocation to the safe asset and invest more in the less-familiar risky asset.

As Panels C and D of Figure 11 show, households' welfare—measured as the certainty-equivalent of consumption (normalized by total output)—inherits the patterns exhibited by households' wealth shares. That is, while less-confident households' welfare declines over time, that of more-confident households increases, with the (year-on-year) changes in welfare being most pronounced in the early years.[26]

To illustrate the crucial impact that financial innovation and financial education can have on the dynamics of households' wealth shares and welfare, Figure 11 also shows the results for some additional cases. First, it depicts the case in which less-confident households have no access to

[25] In the case of full confidence, both groups of households 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 households, respectively.

[26] This pattern arises when computing less-confident households' certainty-equivalent consumption under objective beliefs (recall, they are, on average, unbiased) and also under subjective beliefs. For further discussion of welfare in multi-asset economies with heterogeneous beliefs, see Fedyk, Heyerdahl-Larsen, and Walden (2012).

Figure 11: Dynamics of Wealth Shares and Welfare

Panels A and B plot the average share of aggregate wealth held by the less- and more-confident households. Panels C and D plot the average certainty-equivalent (CEQ) of consumption for the less- and more-confident households. 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 households have no access to the less-familiar asset. "$A_{2,1} = 40$" refers to the setting in which less-confident households 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.

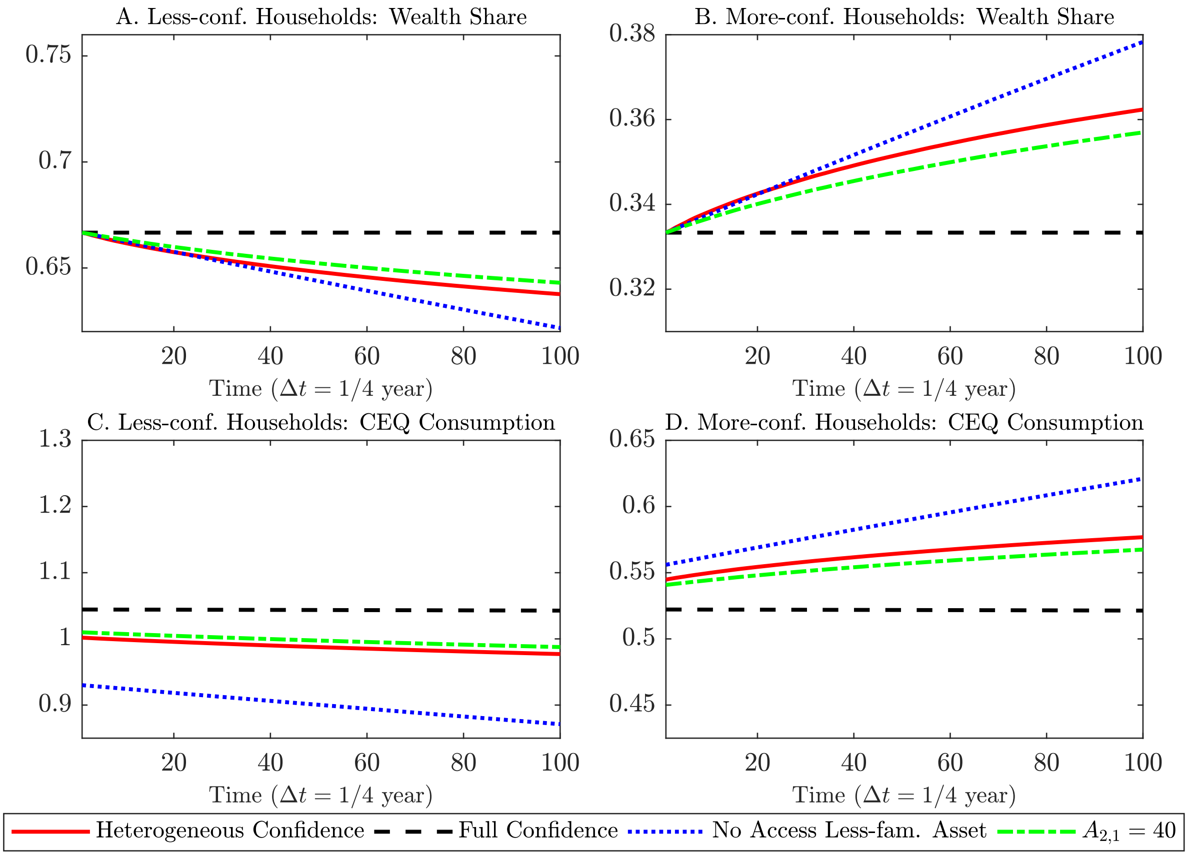

the less-familiar asset, whereas more-confident households can invest in this asset. Comparing this setting to that with confidence heterogeneity highlights that financial innovation, which gives less-confident households access to new assets, slows down the rise in wealth inequality and shifts up (down) the welfare of less-confident (more-confident) households—even though they are not fully confident about those assets.

Second, Figure 11 depicts the case in which less-confident households 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); for example, as a result of financial education or investment advice. Again, the consequence is a slower increase in wealth inequality and an improvement in less-confident households' welfare.

6 Comparison with Alternative Models

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

6.1 Alternative Belief Specifications

We start with an analysis of alternative belief specifications, the results of which are displayed in Figure 12. To highlight the differences to our baseline specification, we shut down confidence heterogeneity regarding the mean dividend-growth rate in each of these alternative specifications.

First, we consider the case in which, instead of lower confidence about the less-familiar asset's expected dividend-growth rate, less-confident households perceive the less-familiar asset to be riskier than it is. In particular, we assume that less-confident households 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, less-confident households allocate a slightly higher proportion of wealth to the risk-free asset. But, as Panel A of Figure 12 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 households. 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 households' myopic asset demand. These small differences in the holding of safe and risky assets are also the explanation for this model's failure to generate much heterogeneity in households' average investment returns and portfolio volatility.

Figure 12: Alternative Belief Specifications

The figure illustrates how portfolio 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 households. Panels E and F depict the *differences* in the average portfolio volatility and the portfolio Sharpe ratios of less- and more-confident households (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 households 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 households 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 households have time-invariant pessimistic beliefs (with $\hat{\mu}_{2,t} = 0.3\%, \forall t$), while the other class of households has unbiased beliefs. In all three cases, households 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.

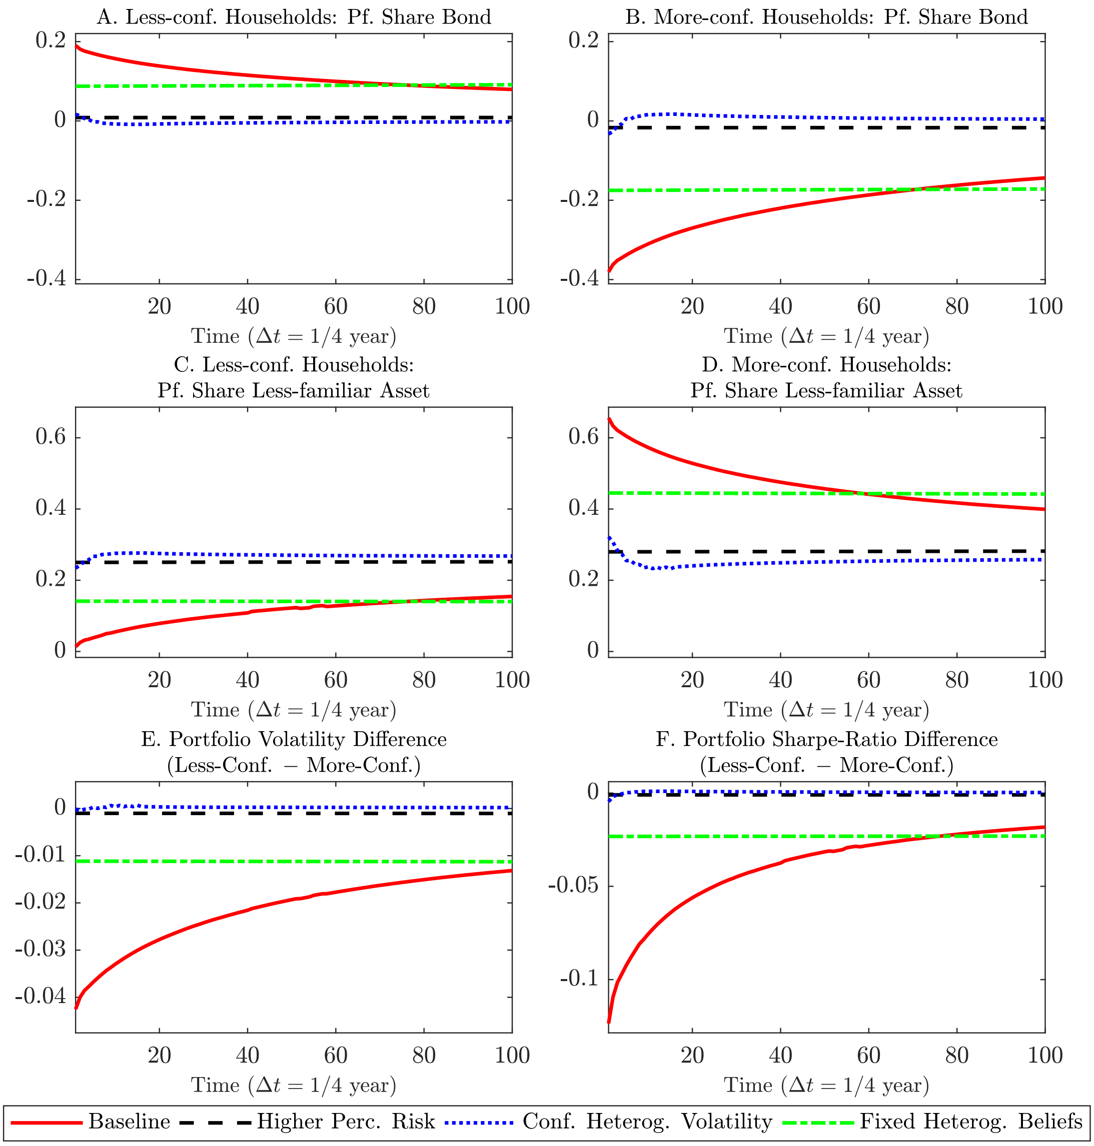

Second, we consider the case in which less-confident households are less confident about the *volatility* of the less-familiar asset (instead of the mean, as in our baseline model) but learn about it from realized dividend growth.[27] In this case, there is some heterogeneity in households' portfolios; in particular, less-confident households allocate more wealth to the safe asset and less wealth to the less-familiar asset initially. But, the magnitude of the portfolio 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, households' continuation utility is substantially less sensitive to these variations. Moreover, the little heterogeneity present in the early periods vanishes very quickly so that households are very soon holding the same portfolios; that is, there is no persistence in portfolio heterogeneity—in sharp contrast to the large persistence in our baseline case. The same is true for investment returns. The short-lived heterogeneity is a consequence of the fact that variance parameters are, in general, relatively easy to learn.

Third, we study the case in which "less-confident households" have time-invariant pessimistic beliefs, while the other group of households has unbiased beliefs. In this case, the pessimistic households 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). The risk-adjusted portfolio returns and volatility of less-confident households are also lower than those of more-confident households. While these patterns in average portfolio allocations and investment returns resemble most closely those in our baseline framework, time-invariant beliefs have, naturally, a difficult time rationalizing the *dynamics* of households' portfolio allocations and returns observed empirically.

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

6.2 Alternative Preference Specifications

We now turn to the case of alternative preference specifications. The results of these experiments are displayed in Figure 13.

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

Figure 13: Alternative Preference Specifications

The figure illustrates how portfolio 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 households. Panels E and F depict the *differences* in the average portfolio volatility and the portfolio Sharpe ratios of less- and more-confident households (Less-conf. − More-conf.). Portfolio returns are annualized and computed under the objective beliefs. "Log Utility" refers to a setting in which both households 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 households 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" households have a risk aversion of 7 whereas the other households have a risk-aversion of 13—in the absence of confidence heterogeneity. "Heterog. EIS" refers to a setting in which "less-confident" households have an EIS of 1.5 whereas the other households 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.

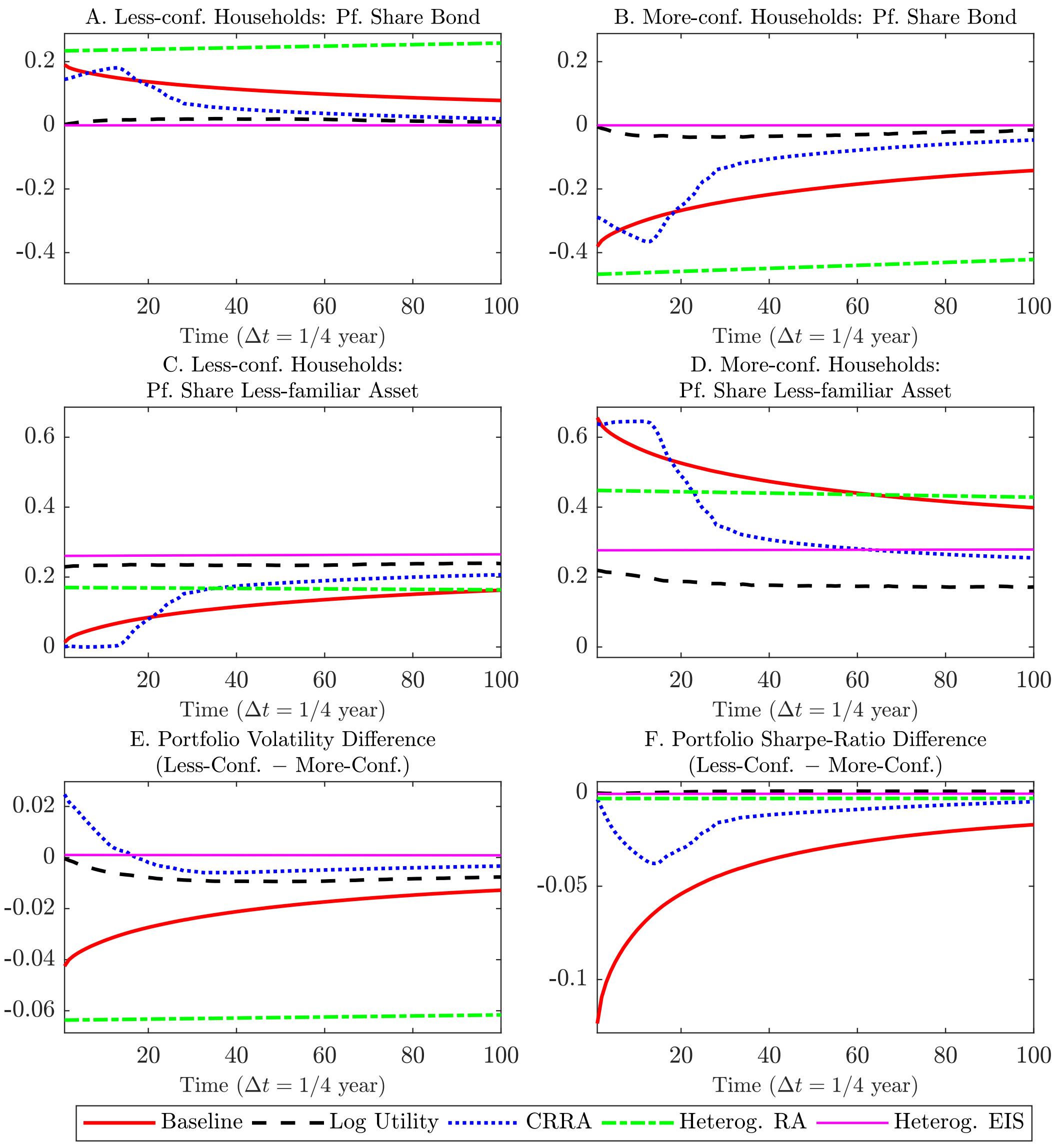

First, we consider the case in which households differ in their confidence regarding the mean dividend-growth rate of the less-familiar asset (exactly as in our baseline model) but assume that households have *log* utility instead of Epstein-Zin-Weil recursive utility. In this case, there is practically no heterogeneity in households' average portfolio holdings and investment returns despite heterogeneous confidence. That is, allocations of the less- and more-confident households to the safe asset are practically the same (Panels A and B). Moreover, as is well known, households 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, households' 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 portfolio dynamics observed in the data.

Next, we consider the case where households 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 portfolio holdings in the early periods are comparable to those in our baseline model, but there is much less persistence.[28] Moreover, the patterns in the risk and returns of households' portfolios are quite different from the baseline model and the data. In particular, less-confident households' portfolio returns are initially *more* volatile—because of the traditional asset's higher return volatility (compared to that of the less-familiar asset). Moreover, recall that, 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 households differ in their *preferences* but have the same level of (full) confidence. If the two groups of households differ only in their relative risk aversion, then the more risk-averse households (which, for simplicity, we dub "less-confident households" 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

[28]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 portfolio holdings. Specifically, because households' portfolios 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.

precautionary-savings demand. More risk-averse households 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 households 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 households are the same, because the decline in portfolio volatility offsets precisely the decline in the portfolio excess return.

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

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

Calvet, Campbell, Gomes, and Sodini (2019) document considerable heterogeneity in households' EIS and rate of time preference; however, our findings highlight that such heterogeneity is not sufficient to explain the empirical patterns in households' portfolio holdings and investment returns. In contrast, heterogeneity in risk aversion can explain some of the observed patterns, but Calvet et al. (2019) report only very only limited cross-sectional differences in risk aversion. Moreover, note that none of the models based on preference heterogeneity alone can generate the portfolio demands conditional on realized and expected returns that are observed empirically (trend-chasing and increased risk taking when the market risk premium is high).

6.3 Other Model Specifications

Figure 14 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 portfolio holdings and investment returns increases. Intuitively, when short-sales are permitted, less-confident households 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 households (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 households, $\hat{\mu}_{2,t}$, (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 households become more confident.

Third, reducing the initial wealth share of the less-confident households ($w_{2,1} = 1/2$ instead of $w_{2,1} = 2/3$) limits the heterogeneity in households' portfolio holdings and investment returns. Qualitatively, however, the results remain unchanged. Intuitively, the more pronounced impact of this change for the investment returns (Panels E and F) can be explained by the fact that asset prices become less sensitive to variations in the less-confident households' beliefs (as their share of wealth is smaller). This, in turn, limits the heterogeneity in households' portfolio returns.

Finally, as expected, increasing the "size" of the less-familiar asset, namely its initial dividend share ($\delta_{2,1} = 0.35$ instead of $\delta_{2,1} = 0.2$), strengthens the impact of confidence heterogeneity. That is, both the precautionary-savings demand and the intertemporal-hedging demand increase—widening the gap in households' portfolio holdings. As a result, the differences in households' investment returns are also more pronounced.

Figure 14: Other Model Specifications

The figure illustrates how portfolio 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 households. Panels E and F depict the *differences* in the average portfolio volatility and the portfolio Sharpe ratios of less- and more-confident households (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 households 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.

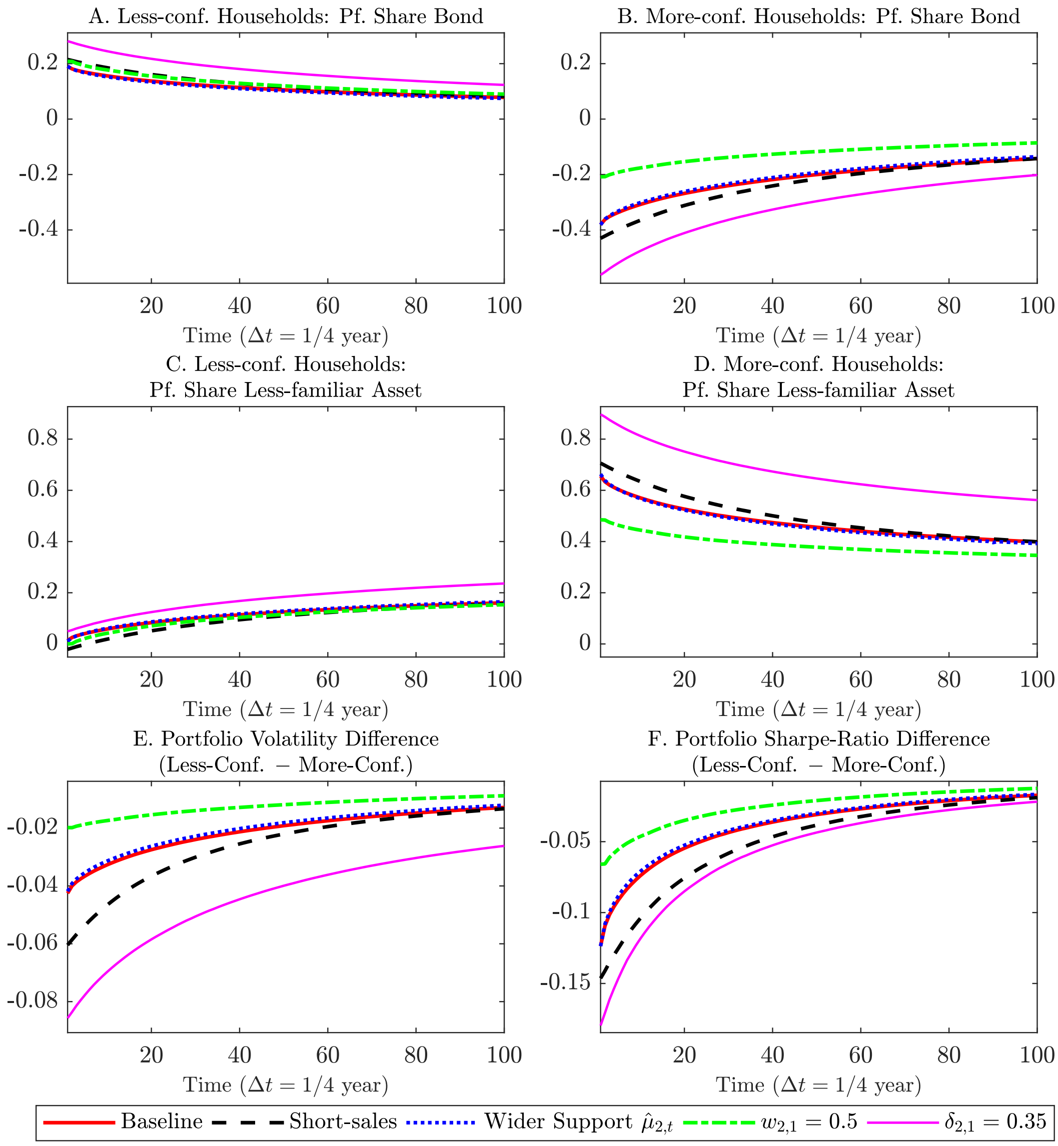

7 Conclusion

Recent empirical evidence has documented many intriguing patterns in households' asset demands and investment returns that persist for long periods (Bianchi, 2018, Fagereng et al., 2020). In this paper, we develop a dynamic general-equilibrium framework with multiple households and multiple risky assets that is driven by a single salient feature of financial markets—heterogeneity in households' confidence—that can explain these empirically observed features of the data. In particular, less-confident households tend to allocate more wealth to safe and familiar risky assets, engage in trend chasing, and earn lower (risk-adjusted) portfolio returns. The model also explains why the less-confident households' behavior, despite Bayesian learning, persists for long periods, thereby exacerbating wealth inequality. Importantly, we show that less-confident households' asset demands are a consequence of *optimal* choices that are driven by their lower confidence about asset returns, rather than behavioral biases or investment mistakes.

From a practical perspective, our work demonstrates that a critical step in reducing wealth inequality is to address the *heterogeneity in confidence* across households. For instance, new financial technologies, such as robo-advisors, which induce households to hold well-diversified portfolios (D'Acunto, Prabhala, and Rossi, 2019, Loos, Previtero, Scheurle, and Hackethal, 2020), could have a significant impact on leveling the playing field. Similarly, "nudges" in the form of default choices that offer well-diversified investment products in which less-confident households could invest with greater confidence could also help (Thaler and Sunstein, 2008). Solutions based on financial engineering—structured products that provide downside protection—are another way of inducing less-confident households to increase their investment in risky assets (Calvet et al., 2020). Of course, financial education should complement these measures (Hastings et al., 2013, Lusardi and Mitchell, 2014).

A Optimality Conditions and Equilibrium

The objective of each household 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 $\theta_{k,2,t} \geq 0$.[29]

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

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

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

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

$$\begin{aligned}\frac{\partial \mathcal{L}_{k,t}}{\partial \theta_{k,2,t}} &= \beta\, V_{k,t}^{\frac{1}{\psi}}\, E_t^k \left[V_{k,t+1}^{1-\gamma} \right]^{\frac{1-\frac{1}{\psi}}{1-\gamma}-1} E_t^k \left[V_{k,t+1}^{-\gamma} \frac{\partial V_{k,t+1}}{\partial \theta_{k,2,t}} \right] - \eta_{k,t}\, (1-\Lambda_{k,t}) S_{n,t} \equiv 0, \qquad \text{(A4)} \\ &\Lambda_{k,t}\theta_{k,2,t} \geq 0, \quad \theta_{k,2,t} \geq 0, \quad \Lambda_{k,t} \geq 0.\end{aligned}$$

[29] 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 household 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 12, 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 households:

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

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

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

and the market-clearing conditions:[30]

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

[30] 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 $j = 1, \ldots, J$.[31] 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 households and J states, and both households' 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 households.

C Confidence Heterogeneity Regarding Asset's Dividend Volatility

In the case of uncertainty regarding the less-familiar asset's cash-flow *volatility*, we assume that less-confident households 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

[31] 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 households' probability measure, the probabilities are set to match the less-confident households' perceived dividend-growth rate, $\hat{\mu}_{2,t}$ and perceived dividend volatility $\sqrt{1 + A_{2,t}}\, \sigma_2$.

dynamics of $\kappa_{1,t+1}$ and $\kappa_{2,t+1}$ given by

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

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 $\kappa_{2,t}/(\kappa_{1,t} - 2)$.[32]

For the illustration in Section 6, 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 households' portfolios and investment returns *within our framework* naturally depends on the parameter values for households' preferences and beliefs as well as the dividend dynamics. We now briefly describe the (quantitative) impact of variations in these parameter values.

Households' Preferences: A reduction in relative risk aversion of both groups of households, γ, weakens households' desire to smooth consumption across states and, hence, less-confident households' precautionary-savings and their intertemporal-hedging demand decline (in absolute terms). Consequently, heterogeneity in households' 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 households 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 households' EIS, ψ, leaves less-confident households' 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, households 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

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

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.

Households' Beliefs: Not surprisingly, an increase in less-confident households' 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 households' 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 households' SDFs) further limit the heterogeneity in investment returns.

Variations in the level of less-confident households' initial beliefs, $\hat{\mu}_{2,k}$, also have intuitive effects. For example, if less-confident households 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 households 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 households' portfolios and investment returns. The decrease in the heterogeneity in households' portfolios is partially offset by a relatively larger increase in the less-familiar asset's risk premium (because changes in less-confident households' 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 households' 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

Notes: 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.

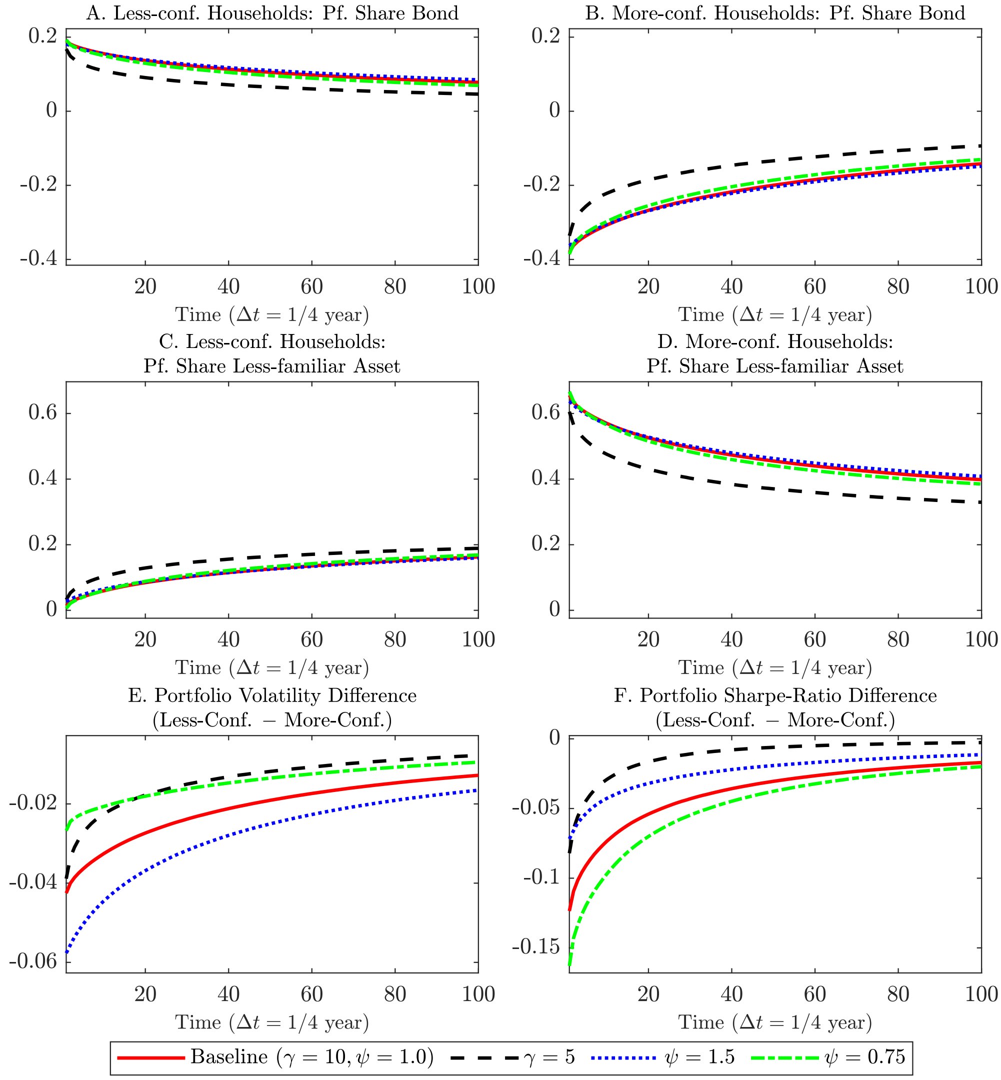

Figure D2: Robustness—Variations in Beliefs

Notes: 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.

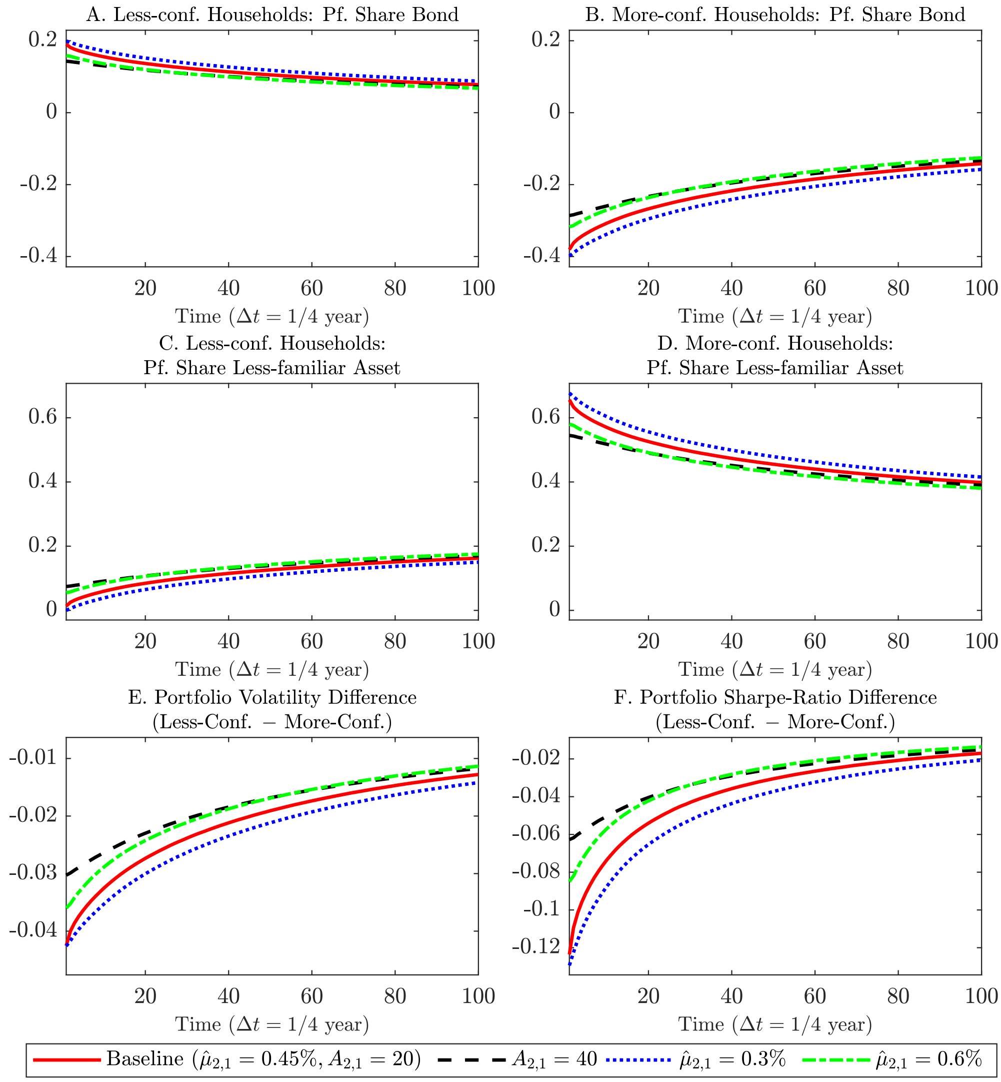

Figure D3: Robustness—Variations in Dividend Dynamics

Notes: 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.

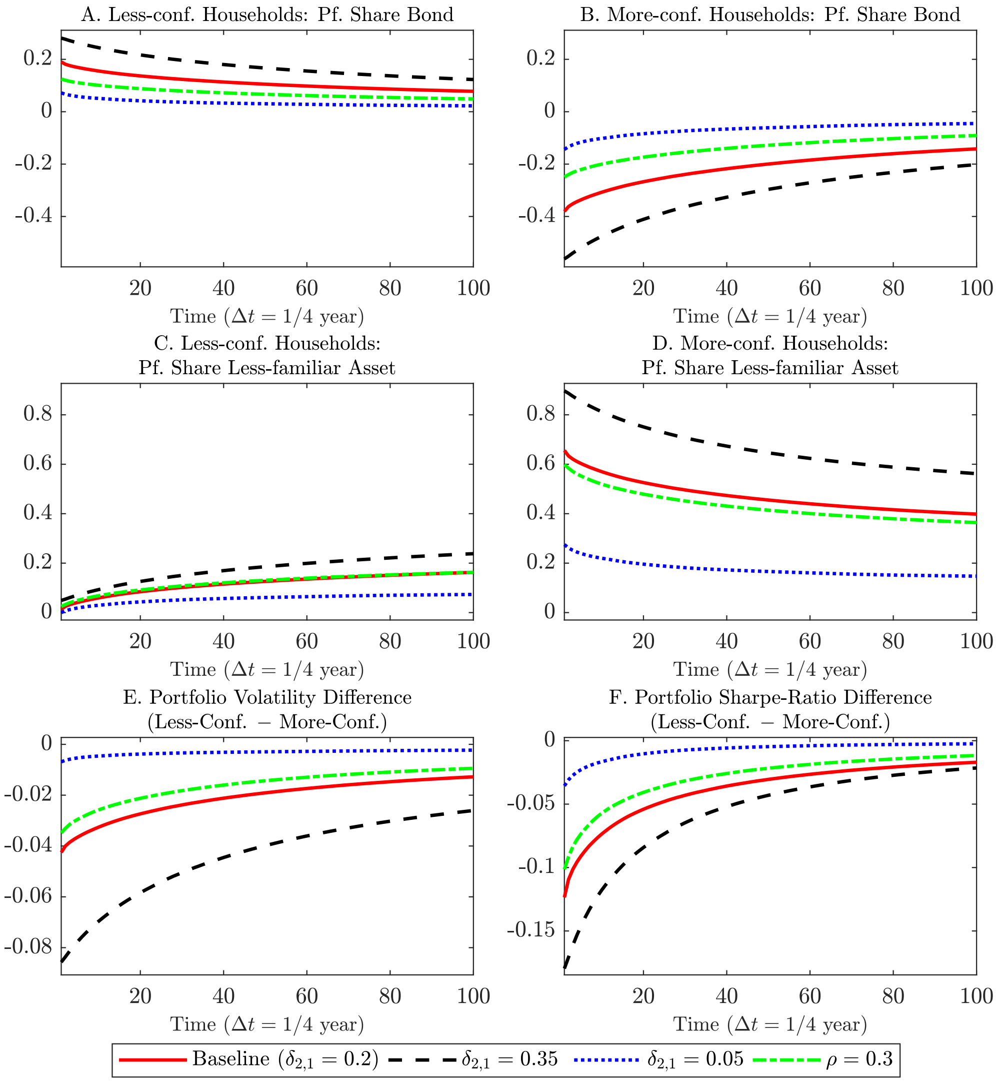

REFERENCES

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

Bach, Laurent, Laurent E. Calvet, and Paolo Sodini, 2020, Rich pickings? Risk, return, and skill in the portfolios of the wealthy, *American Economic Review* (forthcoming).

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.

Basak, Suleyman, 2005, Asset pricing with heterogeneous beliefs, *Journal of Banking and Finance* 29, 2849–2881.

Benartzi, Shlomo, 2001, Excessive extrapolation and the allocation of 401(k) accounts to company stock, *Journal of Finance* 56, 1747–1764.

Benhabib, Jess, Alberto Bisin, and Mi Luo, 2019, Wealth distribution and social mobility in the U.S.: A quantitative approach, *American Economic Review* 109, 1623–47.

Benhabib, Jess, Alberto Bisin, and Shenghao Zhu, 2011, The distribution of wealth and fiscal policy in economies with finitely lived agents, *Econometrica* 79, 123–157.

Bhamra, Harjoat S., and Raman Uppal, 2019, Does household finance matter? Small financial errors with large social costs, *American Economic Review* 109, 1116–54.

Bianchi, Milo, 2018, Financial literacy and portfolio dynamics, *Journal of Finance* 73, 831–859.

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.

Bucher-Koenen, Tabea, Annamaria Lusardi, Rob Alessie, and Maarten Van Rooij, 2017, How financially literate are women? an overview and new insights, *Journal of Consumer Affairs* 51, 255–283.

Calvet, Laurent, John Campbell, Francisco Gomes, and Paolo Sodini, 2019, The cross-section of household preferences, Working paper, Harvard University.

Calvet, Laurent E., John Y. Campbell, and Paolo Sodini, 2007, Down or out: Assessing the welfare costs of household investment mistakes, *Journal of Political Economy* 115, 707–747.

Calvet, Laurent E., John Y. Campbell, and Paolo Sodini, 2009a, Fight or flight? Portfolio rebalancing by individual investors, *Quarterly Journal of Economics* 124, 301–348.

Calvet, Laurent E., John Y. Campbell, and Paolo Sodini, 2009b, Measuring the financial sophistication of households, *American Economic Review: Papers and Proceedings* 99, 393–398.

Calvet, Laurent E., Claire Celerier, Paolo Sodini, and Boris Vallee, 2020, Can security design foster household risk-taking?, CEPR Working Paper No. 14955.

Campbell, John Y., 2006, Household finance, *Journal of Finance* 61, 1553–1604.

Campbell, John Y., 2016, Restoring rational choice: The challenge of consumer financial regulation, *American Economic Review: Papers & Proceedings* 106, 1–30.

Campbell, John Y., Joao Cocco, Francisco Gomes, Pascal Maenhout, and Luis M. Viceira, 2001, Stock market mean reversion and the optimal equity allocation of a long-lived investor, *European Finance Review* 5, 269–292.

Campbell, John Y., Tarun Ramadorai, and Benjamin Ranish, 2015, Getting better or feeling better? How equity investors respond to investment experiences, NBER Working Paper No. 20000.

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.

Clark, Robert, Annamaria Lusardi, and Olivia S. Mitchell, 2017, Financial knowledge and 401(k) investment performance: A case study, *Journal of Pension Economics and Finance* 16, 324–347.

Cocco, Joao, Francisco Gomes, and Pascal Maenhout, 2005, Consumption and portfolio choice over the life-cycle, *Review of Financial Studies* 18, 491–533.

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.

D'Acunto, Francesco, Nagpurnanand Prabhala, and Alberto Rossi, 2019, The promises and pitfalls of robo-advising, *Review of Financial Studies* 32, 1983–2020.

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.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri, 2020, Heterogeneity and persistence in returns to wealth, *Econometrica* 88, 115–170.

Fedyk, Yurii, Christian Heyerdahl-Larsen, and Johan Walden, 2012, Market selection and welfare in a multi-asset economy, *Review of Finance* 17, 1179–1237.

Gabaix, Xavier, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll, 2016, The dynamics of inequality, *Econometrica* 84, 2071–2111.

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 Alok Kumar, 2008, Equity portfolio diversification, *Review of Finance* 12, 433–463.

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).

Gomes, Francisco, Michael Haliassos, and Tarun Ramadorai, 2020, Household finance, forthcoming *Journal of Economic Literature*.

Gomes, Francisco, and Alexander Michaelides, 2003, Portfolio choice with internal habit formation: A life-cycle model with uninsurable labor income risk, *Review of Economic Dynamics* 6, 729–766.

Gomes, Francisco, and Alexander Michaelides, 2005, Optimal life-cycle asset allocation: Understanding the empirical evidence, *Journal of Finance* 60, 869–904.

Gomes, Francisco, and Alexander Michaelides, 2008, Asset pricing with limited risk sharing and heterogeneous agents, *Review of Financial Economics* 21, 415–448.

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

Grinblatt, Mark, Matti Keloharju, and Juhani Linnainmaa, 2011, IQ and stock market participation, *Journal of Finance* 66, 2121–2164.

Guiso, Luigi, Michael Haliassos, and Tullio Jappelli, 2002, *Household Portfolios* (MIT Press, Cambridge, MA).

Guiso, Luigi, Paola Sapienza, and Luigi Zingales, 2008, Trusting the stock market, *Journal of Finance* 63, 2557–2600.

Guiso, Luigi, and Paolo Sodini, 2013, Household finance: An emerging field, in George M. Constantinides, Milton Harris, and René M. Stulz, (eds.) *Handbook of the Economics of Finance* (Elsevier, Amsterdam).

Haliassos, Michael, 2003, Stockholding: Recent lessons from theory and computations, in *Stockholding in Europe*, 30–49 (Springer, New York).

Hastings, Justine S., Brigitte C. Madrian, and William L. Skimmyhorn, 2013, Financial literacy, financial education, and economic outcomes, *Annual Review of Economics* 5, 347–373.

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

Kacperczyk, Marcin, Jaromir Nosal, and Luminita Stevens, 2019, Investor sophistication and capital income inequality, *Journal of Monetary Economics* 107, 18–31.

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.

Loos, Benjamin, Alessandro Previtero, Sebastian Scheurle, and Andreas Hackethal, 2020, Robo-advisers and investor behavior, Working paper, University of Technology of Sidney.

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

Lusardi, Annamaria, Pierre-Carl Michaud, and Olivia S. Mitchell, 2017, Optimal financial knowledge and wealth inequality, *Journal of Political Economy* 125, 431–477.

Lusardi, Annamaria, and Olivia S. Mitchell, 2014, The economic importance of financial literacy: Theory and evidence, *Journal of Economic Literature* 52, 5–44.

Malmendier, Ulrike, Demian Pouzo, and Victoria Vanasco, 2020, Investor experiences and financial market dynamics, *Journal of Financial Economics* 136, 597–622.

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.

Nagel, Stefan, and Zhengyang Xu, 2019, Asset pricing with fading memory, NBER Working Paper No. 26255.

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).

Piketty, Thomas, 2014, *Capital in the Twenty-First Century* (Harvard University Press, Cambridge, MA).

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.

Thaler, Richard H., and Cass R. Sunstein, 2008, *Nudge: Improving Decisions about Health, Wealth, and Happiness* (Yale University Press, New Haven, CT).

Von Gaudecker, Hans-Martin, 2015, How does household portfolio diversification vary with financial literacy and financial advice?, *Journal of Finance* 70, 489–507.

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.