Asset Pricing with Heterogeneous Agents and Long-Run Risk*

Walter Pohl
University of Groningen

Karl Schmedders
University of Zurich

Ole Wilms
Tilburg University

May 12, 2020

Abstract

This paper shows that belief differences have strong effects on asset prices in consumption-based asset-pricing models with long-run risks. Belief heterogeneity leads to time-varying consumption and wealth shares of the agents. This time variation can resolve several asset-pricing puzzles, including the large countercyclical variation of expected risk premia, the volatility of the price–dividend ratio, the predictability of cash flows and returns, and the large predictability of returns in recessions. These findings show that belief differences, a widely observed attribute of investors, significantly improve the explanatory power of long-run risk asset-pricing models.

Keywords: asset pricing, belief differences, heterogeneous agents, long-run risk, recursive preferences.

JEL codes: G11, G12.

*We are grateful to an anonymous referee and the Editor, Ron Kaniel, for thoughtful reviews of earlier versions. Furthermore, we are indebted to Ravi Bansal, Nicole Branger, Michael Brennan, Ric Colacito, Lars Hansen, Ravi Jagannathan, Kenneth Judd, Joren Koëter, Lars Lochstoer, Michael Rockinger, Christian Schlag, and in particular our discussants Jarda Borovička, Bjarne Astrup Jensen, and Sergei Isaenko for helpful discussions on the subject. We thank Dave Brooks, Megan Payne, and Michelle Perrinjaquet for excellent editorial support on earlier versions. And we are grateful to audiences at the University of Zurich, Stanford University, the University of Münster, the University of St. Gallen, Tilburg University, NHH Bergen, the University of Groningen, Universität Hamburg, the European Central Bank, Aarhus University, the City University of Hong Kong, Cass Business School, the 2016 World Finance Conference (New York), the 2017 ASSA meetings, the SGF Conference 2017, the CEQURA Conference 2017, and the DGF Conference 2017 for comments.

1 Introduction

This paper shows that belief differences have large effects on asset prices in consumption-based asset-pricing models with recursive preferences. The simplifying assumption of a representative investor is standard in modern asset-pricing models with recursive preferences; see, for example, Bansal and Yaron (2004), Wachter (2013), Ju and Miao (2012), and many others. While this assumption greatly simplifies model solutions, it ignores an important degree of realism—namely, differences between investors. We demonstrate that for Epstein–Zin preferences, which are widely used in modern asset-pricing models, belief differences—a well documented feature on financial markets[1]—have long-lasting effects on prices. Even small differences in beliefs lead to time-varying consumption and wealth shares of the agents. This time variation can resolve a large number of asset-pricing puzzles, such as the excess volatility of the market portfolio (Shiller, 1981), the large countercyclical variation of expected risk premia (Martin, 2017), the predictability of cash-flows and returns (Beeler and Campbell, 2012), and the countercyclical variation of return predictability (Henkel, Martin, and Nardari, 2011; Dangl and Halling, 2012). Hence, investor heterogeneity can play a leading role in explaining financial market data in modern asset-pricing models.

The Bansal–Yaron long-run risk model (Bansal and Yaron, 2004) has emerged as one of the premier consumption-based asset-pricing models. It can generate many of the features of aggregate stock prices that have long been considered puzzles. The model generates a high equity premium by combining two mechanisms—investors with a taste for the early resolution of uncertainty, and very persistent shocks to the growth rate of consumption. For long-run risk to generate a high equity premium, the level of persistence must be very close to a unit root. The amount of persistence in the data is very difficult to measure, and arguments for a range of estimates have appeared in the literature; see Bansal, Kiku, and Yaron (2016), Schorfheide, Song, and Yaron (2018), and Grammig and Küchlin (2018). This literature suggests that there is considerable scope for disagreement over the true value.

The classical *market selection hypothesis* of Alchian (1950) and Friedman (1953) argues that this type of agent heterogeneity does not matter in the long run, because prices will ultimately be set by the agent with correct beliefs. This argument was confirmed for von Neumann-Morgenstern expected utility by the work of Sandroni (2000) and Blume and Easley (2006). This argument does not generalize to recursive utility, as shown by Borovička (2019) in the continuous-time setting with i.i.d. consumption growth and Dindo (2019) in a similar

[1] See, for example, Anderson, Ghysels, and Juergens (2005), Patton and Timmermann (2010), Buraschi, Trojani, and Vedolin (2014), and Carlin, Longstaff, and Matoba (2014).

setting in discrete time. So, given that agents can easily disagree over difficult-to-measure parameters and given that agents with wrong beliefs can persistently affect market prices, what are the consequences for asset pricing?

We show that a small amount of agent disagreement can significantly improve the realism of long-run risk as an explanation of asset-pricing puzzles. Starting with the Case I model of Bansal and Yaron (2004), which features long-run risk but does not introduce exogenous stochastic volatility to the model, we demonstrate that a small difference of opinion leads to endogenous fluctuations in the wealth distribution, which strongly affect prices. These price changes can explain the significant excess volatility of the market portfolio (Shiller, 1981), the large countercyclical variation of expected risk premia (Martin, 2017), the predictability of cash-flows and returns (Beeler and Campbell, 2012), and the countercyclical variation of return predictability (Henkel, Martin, and Nardari, 2011; Dangl and Halling, 2012).

We begin the paper by showing how to solve consumption-based asset-pricing models with heterogeneous agents and recursive preferences. Solving such models reveals a critical difference from the representative-agent model. Even for Markovian shocks, equilibrium allocations are no longer a function of the exogenous state alone. As a result, the standard solution methods from consumption-based asset pricing are not applicable. We employ a reformulation of the first-order conditions for the equilibrium that is recursive, through the device of introducing new endogenous state variables. These state variables have a clear interpretation in terms of time-varying weights in a social planner's problem. The weights capture the relative trend in an agent's consumption—an agent who has a declining share of consumption will have a declining weight.[2] To solve for the equilibrium, we propose a new numerical solution technique for heterogeneous-agent continuous-state asset-pricing models. The methodology is not limited to long-run risk models, but it can be applied to solve a broad class of models featuring heterogeneous agents, recursive utility, and continuous or discrete state processes.

Using the theoretical framework for the heterogeneous-agent asset-pricing model, we conduct an in-depth analysis of the effects of belief differences in long-run risk models. For this purpose, we first provide some helpful economic intuition on the different effects of belief heterogeneity in models with CRRA (constant relative risk aversion) preferences and in models with Epstein–Zin preferences. Most importantly, we show why the general result (Sandroni, 2000; Blume and Easley, 2006; Yan, 2008) for CRRA preferences—namely, that agents with

[2]Similar approaches have been used to solve models with multiple goods (Colacito and Croce, 2013), discrete-state models without growth and risk-sensitive preferences (Anderson, 2005), overlapping generation models with different preference parameters (Gârleanu and Panageas, 2015) and different beliefs (Collin-Dufresne, Johannes, and Lochstoer, 2016a), and models with i.i.d. consumption growth and belief differences (Borovička, 2019).

wrong beliefs do not survive in the long run—does not always hold for Epstein–Zin agents. For CRRA preferences, the changes in the consumption shares only depend on the subjective beliefs of the agent about the future state—we call this the speculation motive. Importantly, investors have no incentive to hedge long-term risks and in the long run only the investors with the correct beliefs survive. For Epstein–Zin utility with a preference for the early resolution of risks, agents do care about long-run consumption risks. In particular, agents believing in a higher persistence have an incentive to hedge low growth states. Investors believing in lower persistence are willing to provide the insurance for the bad states and collect a premium in the form of higher future wealth in return—we call this the risk-sharing motive. As long as risk premia in the economy are sufficiently large, this motive will, on average, transfer wealth to the agents who believe in a lower persistence irrespective of which investor holds the correct beliefs.

The interaction of the economic effects resulting from the two motives has strong economic implications. To analyze these implications we perform a comprehensive numerical analysis of the complete-markets heterogeneous-agent economy. As our baseline calibration of the model, we employ the calibration (without stochastic volatility) of Bansal and Yaron (2004, Case I) with the only exception being that we use two different persistence levels. We replace the original persistence value of the long-run risk process with two different values; compared to the original value, the first agent believes in a slightly larger and the second agent in a slightly smaller persistence level. We show that belief heterogeneity endogenously adds priced consumption risks to the model due to persistent changes in the wealth distribution. These risk premia are countercyclical and time varying and are consistent with the empirical findings recently reported in Martin (2017) and Martin and Wagner (2019). When there are negative shocks to the long-run risk component, the risk-sharing motive implies that the insurance provided by the investors with lower beliefs about the persistence pays off. Hence, when the economy enters a recession (multiple negative shocks to the long-run risk component), there is a wealth transfer to the investor who believes in a larger value for the persistence. As these are the investors that demand the larger risk premia, belief heterogeneity adds significant countercyclical variation in risk premia to the long-run risk model.

We report a standard deviation of expected risk premia of 5.73 percent in line with the values reported in Martin (2017). Standard representative-agent long-run risk models are not able to replicate this finding even when including exogenous stochastic volatility, a feature that is deliberately included to obtain time variation in risk premia (see Bansal and Yaron (2004)). Furthermore, the heterogeneous-belief model generates a large and significant equity premium and also addresses other empirical deficiencies of the representative-agent model, which have been emphasized by Beeler and Campbell (2012). Beeler and Campbell (2012) show that the long-run risk model cannot explain the large volatility of the price–dividend ratio observed in the data (a value of 0.45 compared to 0.18 in the model). In the heterogeneous-agent setup, the shifts in the wealth distribution increase the volatility of the price–dividend ratio to levels close to the data (0.38) as the impact of the different agents on asset prices varies over time. The variation in the wealth distribution also helps to address the predictability puzzle pointed out by Beeler and Campbell (2012). The endogenous variation in asset prices increases the predictability of returns while simultaneously decreasing the predictability of consumption and dividend growth. Furthermore, Henkel, Martin, and Nardari (2011) and Dangl and Halling (2012) provide evidence that return predictability is particularly high in recessions, while it is significantly lower in economic expansions. We show that our heterogeneous-agent model endogenously explains these predictability patterns. In bad economic times, the consumption shares of the agents believing in a higher persistence increase as the insurance against the low growth states pays off. This in turn increases risk prices and hence return predictability becomes larger in recessions. This endogenous mechanism to explain the countercyclical variation in return predictability is absent in the representative-agent model.

Our paper builds on the recent work by Colacito and Croce (2013) and Colacito, Croce, Liu, and Shaliastovich (2018b) who analyze the effects of risk sharing with recursive preferences in the context of two-country models. They suggest that the risk sharing can significantly affect marginal utilities and hence price dynamics. Due to exogenous shocks and a home bias, the agents in different countries have a strong motive for risk sharing, particularly since they have Epstein–Zin utility with a preference for the early resolution of risk. The reasons for risk sharing in these models are different than those in our economy. In their setting, agents in the two countries trade in response to idiosyncratic endowment shocks. In our setting, it is the subjective disagreement that motivates trade, but some of the intuition carries over, as explained in Section 3.2.

There are several other papers that analyze the influence of different kinds of heterogeneity in economies with Epstein–Zin preferences. For example, Gârleanu and Panageas (2015) analyze the influence of preference heterogeneity on asset prices. They show that combining preference heterogeneity with an overlapping-generations (OLG) setup can generate interesting asset-pricing dynamics. By assuming that one group of investors has a large degree of risk aversion together with a very low elasticity of intertemporal substitution (EIS) while the second group of investors has a significantly lower risk aversion and higher EIS, their OLG model can generate a high equity premium, volatile returns, and a low and smooth risk-free rate as well as long-run predictability of the price–dividend ratio.

Our results are complementary to the findings of Collin-Dufresne, Johannes, and Lochstoer (2016b) and Bidder and Dew-Becker (2016), which show that the asset-pricing implications of long-run risk can emerge endogenously from parameter uncertainty, even without long-run risk being present. Collin-Dufresne, Johannes, and Lochstoer (2016b) show that if investors learn the growth rate from the data, then innovations to expectations of growth rates are permanent. Agents then price in the risk from this permanent shock into their expected growth rates. Bidder and Dew-Becker (2016) show that ambiguity-averse investors will price in long-run risk if they cannot rule it out a priori. In our setup, neither investor suffers from model uncertainty, but despite this difference a clear picture of the effect of long-run risk emerges.

While in the present paper the agents agree to disagree about the long-run risks in the economy, Andrei, Carlin, and Hasler (2019a) provide an explanation of how this disagreement can arise from model uncertainty as market participants calibrate their models differently. They find that uncertainty about long-run risks can explain many stylized facts of stock return volatilities, such as large volatilities during recessions and booms and persistent volatility clustering. Andrei, Hasler, and Jeanneret (2019b) show how model uncertainty can lead to long-run-risk-like behavior in the presence of a noisy signal of the growth rate.

The remainder of this paper is organized as follows. In Section 2 we describe the general asset-pricing model with heterogeneous investors and recursive preferences. Section 3 explains the main economic mechanism of the two-agent economy in a stylized version of the model and Section 4 shows the effects of belief differences on asset prices. In Section 5 we examine the sensitivity of the model's predictions to alternative specifications of agents' beliefs and levels of risk aversion. Section 6 concludes. Online appendices containing a discussion of additional literature, the proofs of all theoretical results, a description of the numerical solution method, a report on numerical errors, and additional results complete the paper.

2 Theoretical Framework

We consider a standard infinite-horizon discrete-time endowment economy with a finite number of heterogeneous agents. Agents can differ with respect to both their utility functions and their subjective beliefs. We restrict our attention to the complete-markets setting, which allows us to reformulate the problem as a social planner's problem. Here we run into a critical difference from the representative-agent problem—even for a Markov economy, equilibrium allocations are no longer required to be functions of the exogenous state alone. This failure of recursiveness occurs for essentially economic reasons—even if aggregate consumption does not contain a trend, the individual consumption allocations can do so. This feature defeats most of the approaches to solving for equilibrium in an infinite-horizon asset-pricing model.

We present a reformulation of the first-order conditions for equilibrium that is recursive. This reformulation involves introducing new endogenous state variables. These state variables have a clear interpretation in terms of time-varying weights in the social planner's problem. The weights capture the relative trend in an agent's consumption—an agent who has a declining share of consumption will have a declining weight.

2.1 The Heterogeneous-Agents Economy

Time is discrete and indexed by $t = 0, 1, 2, \ldots$. Let y_t denote the exogenous state of the economy in period t. The state has continuous support and may be multidimensional. The economy is populated by a finite number of infinitely lived agents, $h \in \mathbb{H} = \{1 \ldots H\}$. Agents choose individual consumption at time t as a function of the entire history of the exogenous state, y^t, where $y^t = (y_0, \ldots, y_t)$. Let $C^h(y^t)$ be the individual consumption for agent h. Similarly, $C(y^t) \in \mathbb{R}_{++}$ denotes the aggregate consumption of all agents as a function of the history, y^t. The individual consumption levels satisfy the usual market-clearing condition,

$$\sum_{h=1}^{H} C^h(y^t) = C(y^t). \tag{1}$$

Agents have subjective beliefs about the stochastic process of the exogenous state. We denote the expectation operator for agent h at time t by E_t^h. Each agent has recursive utility. Let $\{C^h\}_t = \{C^h(y^t), C^h(y^{t+1}), \ldots\}$ denote the consumption stream of agent h from time t forward. The utility of agent h at time t, $U^h(\{C^h\}_t)$, is specified by an aggregator, $F^h(c, x)$, and a certainty equivalence, $G_h(x)$,

$$U^h(\{C^h\}_t) = F^h\left(C^h(y^t), R_t^h\left[U^h(\{C^h\}_{t+1})\right]\right), \tag{2}$$

with

$$R_t^h[x] = G_h^{-1}\left(E_t^h[G_h(x)]\right). \tag{3}$$

We assume that the functions F^h and G_h are both continuously differentiable. This preference framework includes both Epstein–Zin utility and discounted expected utility, for the appropriate choices of F^h and G_h. To simplify the analysis, we ensure that agents never choose zero consumption, in any state of the world, by imposing an Inada condition on the aggregator F^h; so, $F_1^h(c, x) \to \infty$ as $c \to 0$, where F_1^h denotes the derivative of F^h with respect to the first argument.

We also impose a condition on the agents' beliefs. Let $P_{t,t+1}^h$ be the subjective conditional distribution of y_{t+1} given y^t, and $P_{t,t+1}$ be the true conditional distribution. We assume that each agent's expectation can be written in terms of the true distribution as

$$E_t^h[x] = E_t\left[x \frac{\mathrm{dP}_{t,t+1}^h}{\mathrm{dP}_{t,t+1}}\right],$$

for some measurable function $\mathrm{dP}_{t,t+1}^h/\mathrm{dP}_{t,t+1}$. In mathematical terms, every agent's conditional distribution is absolutely continuous with respect to the true distribution. Then, by the Radon–Nikodym theorem (see Billingsley (1999, Chapter 32)) such a $\mathrm{dP}_{t,t+1}^h/\mathrm{dP}_{t,t+1}$ must exist. Accordingly, $\mathrm{dP}_{t,t+1}^h/\mathrm{dP}_{t,t+1}$ is known as the Radon–Nikodym derivative of $\mathrm{P}_{t,t+1}^h$ with respect to $\mathrm{P}_{t,t+1}$. We also assume that, vice versa, the true distribution is absolutely continuous with respect to every agent's subjective distribution.

To solve for equilibrium, we assume that markets are complete so that we can reformulate equilibrium as a social welfare problem (Mas-Colell and Zame, 1991). The social planner maximizes a weighted sum of the individual agents' utilities at $t = 0$. Let $\boldsymbol{\lambda} = \left(\bar{\lambda}^1, \ldots, \bar{\lambda}^H\right) \in \mathbb{R}_{++}^H$ be a vector of positive Negishi weights and let $\{\boldsymbol{C}\}_0 = \left(\{C^1\}_0, \ldots, \{C^H\}_0\right)$ be an H-vector of the agents' consumption processes. Then, the social planner maximizes

$$SP(\{\boldsymbol{C}\}_0; \boldsymbol{\lambda}) = \sum_{h=1}^{H} \bar{\lambda}^h U^h\left(\{C^h\}_0\right) \tag{4}$$

subject to the market-clearing equation (1). We denote an optimal solution to the social planner's problem for given Negishi weights $\boldsymbol{\lambda}$ by $\{\boldsymbol{C}\}_0^*$. For each agent $h \in \mathbb{H}$, let $U_t^h = U^h(\{C^h\}_t^*)$ be the utility in period t at the optimal solution. Also, for ease of notation, we suppress the state dependence of consumption and simply write C_t^h for $C^h(y^t)$.

Theorem 1. *The vector of consumption processes $\{\boldsymbol{C}\}_0^*$ solves the social planner's problem (4,1) for given Negishi weights $\boldsymbol{\lambda} = \left(\bar{\lambda}^1, \ldots, \bar{\lambda}^H\right)$ if and only if the consumption processes* *satisfy the following first-order conditions in each period $t \geq 0$:*

$$\lambda_t^h F_1^h(C_t^h, R_t^h\left[U_{t+1}^h\right]) = \lambda_t^1 F_1^1(C_t^1, R_t^1\left[U_{t+1}^1\right]), \tag{5}$$

where the weights λ_t^h satisfy

$$\lambda_0^h = \bar{\lambda}^h, \tag{6}$$

$$\frac{\lambda_{t+1}^h}{\lambda_{t+1}^1} = \frac{\Pi_{t+1}^h}{\Pi_{t+1}^1}\frac{\lambda_t^h}{\lambda_t^1}, \qquad t \geq 0, h \in \{2, \ldots H\}, \tag{7}$$

with Π_{t+1}^h given by

$$\Pi_{t+1}^h = F_2^h\left(C_t^h, R_t^h[U_{t+1}^h]\right) \cdot \frac{G_h'(U_{t+1}^h)}{G_h'(R_t^h[U_{t+1}^h])}\frac{\mathrm{dP}_{t,t+1}^h}{\mathrm{dP}_{t,t+1}}. \tag{8}$$

Appendix B contains the proof of this theorem as well as those of the theoretical results presented later in this section.

In each period t, the weights λ_t^h are only determined up to a scalar factor, so we are free to choose a normalization. For numerical purposes, the normalization requiring the weights λ_t^h to lie in the unit simplex in every period is convenient. From a conceptual point of view, an attractive choice is to let $\lambda_{t+1}^1 = \Pi_{t+1}^1 \lambda_t^1$, because then for all h, $\lambda_{t+1}^h = \Pi_{t+1}^h \lambda_t^h$.

If the aggregator F^h is additively separable, then the allocation of consumption in (5) depends only on the current value for the weights λ_t^h. Additive separability is the most common case in applications. Discounted expected utility is additively separable, while Epstein–Zin can be transformed to be so. In this particular case, the Negishi weights and individual agents' consumption allocations are closely linked. The following theorem provides an asymptotic result relating the limits of weights λ_t^h to the limits of consumption.

Theorem 2. *Suppose that F^h is additively separable for all $h \in \mathbb{H}$ and that the aggregate endowment is bounded, $C_t \in [\underline{C}, \overline{C}]$ for finite constants $\overline{C} \geq \underline{C} > 0$. If $\lambda_t^j/\lambda_t^i \to \infty$, then $C_t^i \to 0$. If $C_t^i \to 0$, then for at least one other agent j, $\limsup_t \lambda_t^j/\lambda_t^i = \infty$.*

Note that $\limsup_t \lambda_j/\lambda_i$ is a random variable—the limit can depend on the history. Theorem 2 generalizes a similar result by Blume and Easley (2006).

2.2 The Growth Economy with Epstein–Zin Preferences

We now consider the special case of our heterogeneous-agent economy in which aggregate consumption is expressed exogenously in terms of growth rates and agents have Epstein–Zin preferences (see Epstein and Zin (1989) and Weil (1989)). For this popular parametrization of asset-pricing models, we can sharpen the general results of Theorems 1 and 2. Here we state the equilibrium conditions for this model parametrization and refer any interested reader to Appendix B.2 for a proper derivation of those conditions.

If agent h has Epstein–Zin preferences, then

$$F^h(c,x) = \left[(1-\delta^h)c^{\rho^h} + \delta^h x^{\rho^h}\right]^{1/\rho^h} \tag{9}$$

$$G_h(x) = x^{\alpha^h} \tag{10}$$

with parameters $\rho^h \neq 0, \alpha^h < 1$. In this case, the equations are all homogeneous, so we can divide through by aggregate consumption and express the equilibrium allocations in terms of individual consumption shares, $s_t^h = C_t^h / C_t$. Market clearing (1) implies that

$$\sum_{h=1}^{H} s_t^h = 1. \tag{11}$$

Let V_t^h be agent h's value function. We also normalize this function by aggregate consumption, $v_t^h = V_t^h / C_t$. Let $c_t = \log C_t$ and $\Delta c_{t+1} = c_{t+1} - c_t$. The normalized value function of agent h satisfies the following fixed-point equation:

$$v_t^h = \left[(1-\delta^h)(s_t^h)^{\rho^h} + \delta^h R_t^h\left(v_{t+1}^h e^{\Delta c_{t+1}}\right)^{\rho^h}\right]^{\frac{1}{\rho^h}}, \qquad h \in \mathbb{H}, \tag{12}$$

with $R_t^h(x) = \left(E_t^h\left[x^{\alpha^h}\right]\right)^{\frac{1}{\alpha^h}}$. The parameter δ^h is the discount factor, $\rho^h = 1 - \frac{1}{\psi^h}$ determines the elasticity of intertemporal substitution (EIS), ψ^h, and $\alpha^h = 1 - \gamma^h$ determines the relative risk aversion, γ^h, of agent h.

To accompany the normalized value function we introduce a normalized Negishi weight, $\underline{\lambda}_t^h = \frac{\lambda_t^h}{(v_t^h)^{\rho^h - 1}}$. In Appendix B.2 we show that the consumption share s_t^h of agent h is given by

$$\underline{\lambda}_t^h (1-\delta^h)(s_t^h)^{\rho^h - 1} = \underline{\lambda}_t^1 (1-\delta^1)(s_t^1)^{\rho^1 - 1}. \tag{13}$$

Finally, the equations for $\underline{\lambda}_t^h$ simplify to

$$\begin{aligned} \frac{\underline{\lambda}_{t+1}^h}{\underline{\lambda}_{t+1}^1} &= \frac{\underline{\Pi}_{t+1}^h}{\underline{\Pi}_{t+1}^1} \frac{\underline{\lambda}_t^h}{\underline{\lambda}_t^1} \\ \underline{\Pi}_{t+1}^h &= \delta^h e^{\rho^h \Delta c_{t+1}} \frac{\mathrm{dP}_{t,t+1}^h}{\mathrm{dP}_{t,t+1}} \frac{\left(v_{t+1}^h e^{\Delta c_{t+1}}\right)^{\alpha^h - \rho^h}}{R_t^h\left(v_{t+1}^h e^{\Delta c_{t+1}}\right)^{\alpha^h - \rho^h}}, \qquad h \in \mathbb{H}^-. \end{aligned} \tag{14}$$

This simplification gives us $H-1$ nonlinear equations for the equilibrium. In our numerical calculation, we complete the system by requiring that $\sum \underline{\lambda}_t^h = 1$ when we solve for the weights, $\underline{\lambda}_t^h$, given by

$$
\begin{aligned}
\underline{\lambda}_{t+1}^h &= \frac{\underline{\lambda}_t^h \underline{\Pi}_{t+1}^h}{\sum_{h=1}^H \underline{\lambda}_t^h \underline{\Pi}_{t+1}^h} \\
\underline{\Pi}_{t+1}^h &= \underbrace{\delta^h e^{\rho^h \Delta c_{t+1}} \frac{\mathrm{dP}_{t,t+1}^h}{\mathrm{dP}_{t,t+1}}}_{\text{CRRA-Term}} \underbrace{\frac{\left(v_{t+1}^h e^{\Delta c_{t+1}}\right)^{\alpha^h-\rho^h}}{R_t^h \left(v_{t+1}^h e^{\Delta c_{t+1}}\right)^{\alpha^h-\rho^h}}}_{\text{Additional EZ-Term}}, \qquad h \in \mathbb{H}^-.
\end{aligned} \tag{15}
$$

Unlike in the discounted expected utility case, the dynamics of the weights $\underline{\lambda}_t^h$ depend on the value functions (12), which in turn depend on the consumption decisions (13). Hence, to compute the equilibrium we need to jointly solve equations (11)–(15). As there are—to the best of our knowledge—no closed-form solutions for the general model, we present in Appendix C.1 a numerical solution approach, which is based on projection methods as proposed in Pohl, Schmedders, and Wilms (2018) to approximate for the equilibrium functions. In Appendix C.2 we provide a detailed analysis of the accuracy of the solution approach.

In this setting, we can derive an improvement over Theorem 2—the limiting behavior for λ_t^h drives the limiting behavior for an agent's share of aggregate consumption. This result requires no assumptions on aggregate consumption, only that agents have utility in the Epstein–Zin family.

Theorem 3. *Suppose all agents in the economy have Epstein–Zin preferences. If* $\underline{\lambda}_t^j/\underline{\lambda}_t^i \to \infty$, *then* $s_t^i \to 0$. *If* $s_t^i \to 0$, *then for at least one agent,* j, $\limsup_t \underline{\lambda}_t^j/\underline{\lambda}_t^i = \infty$.

This completes our discussion of the theoretical framework for our analysis. Appendix B provides proofs for the three theorems in this section. Along the way, we derive a system of first-order conditions for Epstein–Zin preferences. This system constitutes the foundation for our numerical solution method (see Appendix C).

3 Heterogeneous Beliefs about Consumption Growth

In the following we use the theoretical framework from Section 2 to analyze the influence of belief differences in long-run-risk asset-pricing models. For this purpose, we use the standard model of Bansal and Yaron (2004), but without stochastic volatility so that risk premia would be constant in a model populated by a single representative agent. Log aggregate consumption growth, Δc_{t+1}, and log aggregate dividend growth, Δd_{t+1}, are given by

$$\begin{aligned} \Delta c_{t+1} &= \mu_c + x_t + \eta_{c,t+1} \\ x_{t+1} &= \rho_x x_t + \eta_{x,t+1} \\ \Delta d_{t+1} &= \mu_d + \Phi x_t + \eta_{d,t+1}. \end{aligned} \tag{16}$$

The key feature of the long-run risk model is that there are small but highly persistent shifts in the growth rate of consumption and dividends. The x_t process captures this long-run variation in the means of consumption and dividend growth. The shocks $\eta_{c,t+1}$, $\eta_{x,t+1}$, and $\eta_{d,t+1}$ are independent and normally distributed with mean 0 and standard deviations σ, $\phi_x\sigma$, and $\phi_d\sigma$, respectively. With a preference for early resolution of risks ($\gamma > \frac{1}{\psi}$), investors will dislike shocks in x_t and require a large premium for bearing the resulting risks. To generate a high enough equity premium, x_t must be very persistent and thus ρ_x must be close to one (0.979 in the original calibration of Bansal and Yaron (2004)). If investors must deduce the value of ρ_x from experience, it is reasonable to assume that investors disagree—at least slightly—about that value. In Section 3.1, we provide evidence that even for 500 years of data, point estimates of ρ_x show significant variation. Hence, as there are less than 100 years of data available, it is reasonable to assume that there are differences in the beliefs about ρ_x.

In our calibrations, we consider a model with $H = 2$ agents, denoted by A and B, in which agents disagree on ρ_x but agree on all other parameters. We denote by ρ_x^h the belief of agent h about ρ_x. As x_{t+1} conditional on time t information is normally distributed with mean $\rho_x^h x_t$ and variance $\sigma_x^2 = \phi_x^2\sigma^2$, agents' beliefs $\mathrm{dP}_{t,t+1}^h$ are given by

$$\mathrm{dP}_{t,t+1}^h = \frac{1}{\sqrt{2\pi}\sigma_x} \exp\left(-\frac{1}{2}\left(\frac{x_{t+1} - \rho_x^h x_t}{\sigma_x}\right)^2\right). \tag{17}$$

For $x_t = 0$ both agents have the same beliefs and the belief difference increases the further away x_t is from its long-run mean of 0. Hence, the state of the economy plays an important role for the beliefs and thus the equilibrium consumption shares. We can analyze the first-order effects of different beliefs about x_t on the consumption shares for $\psi^h = \psi$ for all $h \in \mathbb{H}$ directly by plugging equation (13) into equation (15). The equilibrium consumption shares are then given by[3]

$$\frac{s_{t+1}^A}{s_{t+1}^B} = \frac{s_t^A}{s_t^B} \underbrace{\left(\frac{\mathrm{dP}_{t,t+1}^A}{\mathrm{dP}_{t,t+1}^B}\right)^{\psi}}_{\text{CRRA-Term}} \underbrace{\left(\frac{\frac{\left(v_{t+1}^A\right)}{R_t^A\left(v_{t+1}^A\right)}}{\frac{\left(v_{t+1}^B\right)}{R_t^B\left(v_{t+1}^B\right)}}\right)^{1-\gamma\psi}}_{\text{Additional EZ-Term}}. \tag{18}$$

For $\gamma = \frac{1}{\psi}$ the second term vanishes and we obtain the standard case of CRRA preferences. In this case, the changes in the consumption shares only depend on the subjective beliefs of the agents about the future state. If agent A believes that a state tomorrow is more likely compared to the belief of agent B, her consumption share will increase if this state materializes. Hence, investors trade purely based on their subjective probabilities but do not care about the long-term effects of changes in x_t. Therefore, we call the motivation for trading based purely on the subjective beliefs the "speculation motive." As we show below, for CRRA preferences, both investors believe that their consumption share will increase in the future. Hence, the investor who holds the correct beliefs will dominate the economy in the long run; see, for example, Blume and Easley (2006) and Yan (2008).

This result does not hold true for Epstein–Zin preferences where we obtain an additional effect. As in Colacito and Croce (2013) and Borovička (2019), investors want to hedge states that they particularly dislike. Assume, for example, that a state materializes where the continuation utility of investor A, v_{t+1}^A, is low compared to her certainty equivalent $R_t^A\left(v_{t+1}^A\right)$. So, there is an unexpected bad shock that lowers agent A's utility. Furthermore, assume that investor B does not dislike this state as much, and so the deviation to his certainty equivalent is smaller. The equilibrium consumption shares in the model change such that investor A will obtain some compensation in the form of a higher consumption share from investor B to make up for the loss in her continuation utility. (The setup implies that the EZ-term inside the brackets is smaller than 1. If investors have a preference for the early resolution of risks ($\gamma > \frac{1}{\psi}$), the exponent is negative and thus the EZ-term will be larger than 1.) Put differently, the EZ-term alone would lead to an increase in the consumption share of investor A compared to the share of investor B. And so, investor A obtains some of the consumption share of investor B to compensate for her utility loss. In return, investor B will obtain compensation in states that investor A is less concerned about. We show later that in our economy with long-run risks, investors with a high belief about ρ_x particularly dislike states of low growth (small x_t). So investors who believe in a lower ρ_x are providing insurance for these low x_t

[3]We make use of the model property that investors disagree about the distribution of x_t but share the same beliefs about the distribution of shocks to Δc_{t+1}. As the agents have the same preference parameters, the term Δc_{t+1} cancels out in the equation.

states and earn an insurance premium in return. These motivations will transfer wealth on average to the investors who believe in a lower ρ_x. We call this reason for trading the "risk sharing motive."

We show that these two motivations for trading, the motive based on the subjective beliefs and the motive for risk sharing to insure against bad states, can be balanced such that both investors maintain significant consumption shares in long simulations. The changes in the wealth distribution over time then help explain several asset-pricing puzzles including the large countercyclical variation of expected risk premia, the volatility of the price–dividend ratio, and the predictability of cash flows and returns.

For the calibration of our model, we stick to the literature as close as possible in order to be able to isolate the influence of the heterogeneous-agent setup on asset prices. Hence, we use the standard calibration of Bansal and Yaron (2004, Case I)—the calibration without stochastic volatility—so $\psi^A = \psi^B = 1.5$, $\gamma^A = \gamma^B = 10$, $\delta^A = \delta^B = 0.998$, $\mu_c = \mu_d = 0.0015, \sigma = 0.0078, \Phi = 3, \phi_d = 4.5$, and $\phi_x = 0.044$. The only exception is the value for the long-run risk parameter, $\rho_x = 0.979$. We perturb this value so that agent A believes in a slightly larger value and agent B believes in a slightly smaller value. Specifically, we assume that investor A strongly believes in long-run risks ($\rho_x^A = 0.99$) while investor B is more skeptical about them ($\rho_x^B = 0.96$). Furthermore, we assume that investor A holds the correct beliefs, $\rho_x = \rho_x^A = 0.99$, so long-run risks indeed exist and are strong. We choose this particular calibration for several reasons. First, in the long run, both investors have a mean consumption share of about 50 percent, so both investors are important for asset prices. Second, in the finite samples we do not encounter issues with one type of investors being driven out of the market. Third, the equity premium in this economy is in line with the data—a necessary condition for a reasonable consumption-based asset-pricing model. And finally, the speculation and risk-sharing motives in the economy are strong and hence we observe interesting asset-pricing implications.

We emphasize that we do not attempt to find a perfect calibration but instead try to impose only minimal changes to the standard calibration to highlight the influence of the heterogeneous-agent setup. We stick to this calibration for the main parts of the paper and provide a detailed sensitivity analysis in Section 5, examining the influence of the value of the true persistence parameter and of the magnitude of the belief differences on the model predictions. To help motivate our results, we begin by motivating the plausibility of belief differences in the next section.

3.1 Plausibility of Belief Differences

Does it make sense for agents to disagree over the persistence of long-run risk? We provide two pieces of evidence that it does, one technical and one empirical. First we show that, if long-run risk is present, when agents try to estimate the persistence parameter from historical consumption data, it will take a long time for such an estimate to be reasonably accurate. So, agent disagreement would be consistent with the small samples of historical data that are available. Second, we use individual forecasts of real consumption growth to show that, in practice, agent forecasts do disagree significantly in terms of persistence.

Suppose the true persistence parameter of the long-run risk process is $\rho_x = 0.99$, which is the value we use for ρ_x^A in our main calibrations. Now suppose an investor does not know this parameter but estimates it from a finite sample. How long will it take before estimates of this parameter are guaranteed to be near the true value? To answer this question, we simulate 1,000 time series consisting of 500 years of monthly data and calculate estimates of the persistence after different time periods. Monthly U.S. consumption data has been reported since 1959, so there are currently about 60 years of data available. Annual data points start from 1930 onward. However, using annual instead of monthly data would significantly increase the standard errors in our estimation and the variation in the estimates of the persistence would be even larger. Nevertheless, we report estimates of the persistence after 60 years and 100 years as well as long-term estimates for 500 years.

As a first estimation approach, we assume that the investor directly observes x_t and simply estimates the AR(1) process

$$x_{t+1} = \mu_x + \rho_x x_t + \sigma_x \eta_{x,t+1}. \tag{19}$$

Note that in the long-run risk model $\mu_x = 0$. So as a second approach, we assume that the agent knows that $\mu_x = 0$. In both cases we estimate the model using least squares.

In the data, we do not directly observe x_t but only aggregate consumption growth Δc_{t+1}. Hence, we also consider the case when x_t is unobserved but must be inferred from the time series of consumption growth. To do this, we estimate the full state-space model (16) using the Kalman filter:

$$\begin{aligned} \Delta c_{t+1} &= \mu_c + x_t + \sigma \eta_{c,t+1} \\ x_{t+1} &= \rho_x x_t + \sigma_x \eta_{x,t+1}. \end{aligned} \tag{20}$$

Table 1 reports the mean point estimate of the three estimation approaches over the 1,000 simulated time series, as well as the 5% and 1% quantiles. We observe that for 60 years of

Table 1: Parameter Estimates from Simulated Data

		x_t observable		x_t unobservable
		with constant	w/o constant	all parameters
60 Years	$\hat{\rho}_x$	0.9836	0.9872	0.9613
	$\hat{\rho}_{x,0.05}$	0.9699	0.9750	0.9116
	$\hat{\rho}_{x,0.01}$	0.9583	0.9641	0.5120
100 Years	$\hat{\rho}_x$	0.9865	0.9886	0.9813
	$\hat{\rho}_{x,0.05}$	0.9764	0.9803	0.9566
	$\hat{\rho}_{x,0.01}$	0.9713	0.9756	0.9154
500 Years	$\hat{\rho}_x$	0.9893	0.9897	0.9887
	$\hat{\rho}_{x,0.05}$	0.9862	0.9865	0.9828
	$\hat{\rho}_{x,0.01}$	0.9844	0.9847	0.9783

The table shows the mean point estimates of ρ_x as well as the 5% and 1% quantiles after 60, 100, and 500 years obtained from simulating 1,000 monthly time series of data. In the first approach, Equation (19) is used for x_t, assuming the process is directly observable, and least squares is used to estimate the model parameters; we distinguish two cases of estimating the AR(1) model, the case with and the case without a constant. The second approach assumes that x_t is unobservable and the full state-space model (20) is estimated using the Kalman filter. For the data-generating process, we use $\mu_x = 0, \rho_x = 0.99, \sigma_x = 0.0003432, \mu_c = 0.0015$, and $\sigma = 0.0078$.

data there is the usual finite-sample downward bias in the mean of the point estimates $\hat{\rho}_x$ (see, for example, James and Smith (1998, Case I)). Kendall (1954, Case I) shows that the bias is approximately $-(1 + 3\rho_x)/T$. In our application, the resulting value is -0.0055 for the model with a constant and 60 years of data ($T = 720$), which is in accordance with the value we observe. (The investor can approximate the bias using the point estimate, $\hat{\rho}_x$, and the number of periods, T.) The table also reports the 5% and 1% quantiles of the point estimates from the 1,000 simulations. After 60 years, the range of estimates is still large with 5% quantiles of 0.9699 and 0.9750 for the case with and without the constant. Note that here we assumed that x_t is observable, which is not true in reality. If x_t is unobserved instead and needs to be inferred by the Kalman filter, the range of estimates increases dramatically with 5% and 1% quantiles of 0.9116 and 0.5120, respectively. For our benchmark calibration, we assume a value of $\rho_x^B = 0.96$ which is well within the 5% quantile. Furthermore, we observe that while for 100 years of data the quantile estimates become tighter, our calibrated value of 0.96 is still well within the 5% quantile and it takes a long time for the investor to obtain precise estimates for the highly persistent x_t process (see results after 500 years).

In light of the estimation results, we conclude that even if the investor might learn about the true data-generating process after 500 or more years, it is reasonable to assume that any nontrivial initial belief differences persist for at least 60 years, if not for much longer.

We can also find direct evidence of disagreement in persistence by considering the predictions for U.S. real consumption growth collected for the Survey of Professional Forecasters (SPF). This is a panel dataset that began collecting real-consumption-growth forecasts in 1981. Each survey, the forecaster makes a forecast for several quarters ahead, which allows us to measure the persistence in changes to consumption growth. We compare the forecast for 2 quarters ahead with the forecast for 1 quarter ahead. Note that both forecasts are made simultaneously in the same survey.

To test for the presence of heterogeneous beliefs, we estimate a simple AR(1) model for each forecaster. We restrict to forecasters who have provided predictions for at least 8 surveys. This restriction leaves a sample of 145 different forecasters. The forecasters can enter and leave the sample at different times, so to control for time variation we use a panel regression with quarterly fixed effects.[4] The overall model is

$$\Delta c'_{t,i} = A_i + \beta \Delta c_{t,i} + \beta_i \Delta c_{t,i} + K_t + e_{t,i}, \tag{21}$$

where $\Delta c_{t,i}$ is the log prediction made at time t of consumption growth by forecaster i one quarter ahead, and $\Delta c'_{t,i}$ is the log prediction two quarters ahead. The sum $\beta + \beta_i$ measures the total persistence for each forecaster i, so if the β_i are all zero, then the forecasters all share the same β. To test the joint hypothesis that the β_i are identically zero, we use an F-test to compare the model with and without individual β_i. The F-statistic is 14.65, so we can strongly reject the null even at the 0.1% level and hence we find significant evidence for differences in the persistence parameter.

Figure 1 shows the distribution of AR(1) coefficients $(\beta + \beta_i)$ for the group of forecasters. The spread of estimates is large, with a 5% quantile of -0.299 and a 95% quantile of 0.57 at the quarterly frequency. We, therefore, can conclude that in the forecasts by professionals there is significant disagreement about the persistence of real consumption growth.

In sum, forecasts from the U.S. SPF reveal large discrepancies in the persistence of real-consumption-growth forecasts among professional forecasters. Furthermore, under the assumption of long-run risk in consumption growth, statistical estimates of the persistence of such risk—relying on available data sets—likely contain nontrivial margins of error. In light of these observations, we believe that our assumption of heterogeneous beliefs about the per-

[4] We show the results of the individual regressions without fixed effects in Figure 17 in Section F of the Appendix. The results show a similar spread of regression coefficients as in the panel regressions. So the findings are robust with regard to the different specifications.

Figure 1: Individual Forecaster AR(1) Coefficients

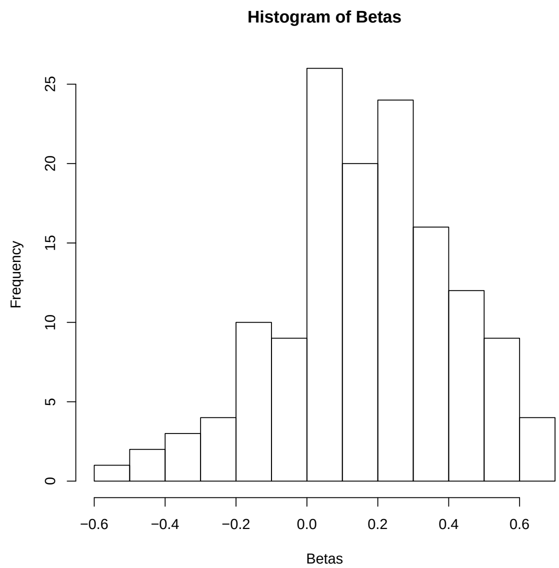

The figure shows a histogram of the AR(1) coefficients ($\beta + \beta_i$) for real-consumption-growth forecasts, see model (21), for a group of forecasters from the U.S. Survey of Professional Forecasters. Only forecasters with predictions in at least 8 surveys are included in the sample. Regressions have been performed with quarterly fixed effects. For the presentation of the histogram, the two highest and the two lowest coefficients have been removed.

sistence of long-run risk is well justified. In the following, we analyze the economic effects that disagreement on persistence generates in our two-agent economy.

3.2 Discrete-State Model

To illustrate the economic intuition of the two-agent economy, we first work in a simplified setting. Results for our full model appear in Section 4. Colacito and Croce (2013) provide a simple explanation of how Epstein–Zin investors behave in terms of the mean-variance trade-off of their continuation utilities. Let $\bar{v}_t^h = \frac{(v_t^h)^{\rho^h}}{\rho^h}$. The value functions of the investors (12) are then given by

$$\bar{v}_t^h = (1 - \delta^h) \frac{(s_t^h)^{\rho^h}}{\rho^h} + \delta^h E_t^h \left((\bar{v}_{t+1}^h)^{\frac{\alpha^h}{\rho^h}} e^{\alpha^h \Delta c_{t+1}} \right)^{\frac{\rho^h}{\alpha^h}}, \qquad h \in \mathbb{H}. \tag{22}$$

If $\bar{v}_{t+1}^h$ is lognormally distributed, then Colacito and Croce (2013) show that this can be approximated as

$$\bar{v}_t^h \approx (1-\delta^h)\frac{(s_t^h)^{\rho^h}}{\rho^h} + \delta^h k_t^h E_t^h(\bar{v}_{t+1}^h) - \delta^h k_t^h \frac{\rho^h - \alpha^h}{2\rho^h}\frac{Var_t^h(\bar{v}_{t+1}^h)}{E_t^h(\bar{v}_{t+1}^h)}, \qquad h \in \mathbb{H}, \tag{23}$$

where $k_t^h = E_t^h(e^{\alpha^h \Delta c_{t+1}})^{\frac{\rho^h}{\alpha^h}}$.

If an agent has CRRA preferences, then $\rho^h = \alpha^h$, and the third term in the equation vanishes. The agent is only interested in the expectation of future utility. In contrast, once an agent has a preference for the early resolution of risk, $\rho^h > \alpha^h$, the second term becomes negative, since $\rho^h E_t^h(\bar{v}_{t+1}^h)$ is positive. In that case, a subjectively higher variance of the continuation utility, $Var_t^h(\bar{v}_{t+1}^h)$, reduces welfare. Thus unlike the CRRA case, the agent is willing to trade off expected future utility, $E_t^h(\bar{v}_{t+1}^h)$, to reduce uncertainty about future utility, $Var_t^h(\bar{v}_{t+1}^h)$. As there is a monotonic relationship between continuation values and agent's wealth (Epstein and Zin, 1989), this means the agent is willing to give up future wealth to avoid future volatility. This trade-off generates an additional motivation to trade compared to the CRRA case as the investors have an incentive to share risks that affect continuation utility.

To illustrate how this trade-off affects equilibrium consumption shares in the long-run risk model, we consider a very simplified model where x_t follows a two-state Markov chain instead of a continuous AR(1) process. Long-run risk then simplifies into whether the economy is in a "good" state or a "bad" state, and the agents disagree on how persistent the current state is. Let z_t be the current state, $z_t \in \{1,2\}$. We assume that state 1 is the bad state and state 2 is the good state. The state x_t can take on one of two levels, x_1, x_2, where $x_1 < x_2$. We assume that $\rho_x^A = 0.99$ and $\rho_x^B = 0.96$ so that investor B believes in faster mean reversion.

We use the discretization method of Rouwenhorst (1995)[5] to discretize this highly persistent process. With this method, we obtain $x_{z_t} = \{-0.0024, 0.0024\}$ and the (subjective) state transition probabilities $P^h = p_{ij}^h$,

$$P^A = \begin{bmatrix} 0.995 & 0.005 \\ 0.005 & 0.995 \end{bmatrix}, \quad P^B = \begin{bmatrix} 0.98 & 0.02 \\ 0.02 & 0.98 \end{bmatrix}.$$

We compare the mean-variance trade-off in two different equilibria: the trade equilibrium, in

[5]The moment-matching technique of Rouwenhorst (1995) is better suited for highly persistent processes compared to standard techniques such as, for example, the method proposed by Tauchen and Hussey (1991); see Flodén (2008).

which the agents can trade freely in a complete market, and the no-trade autarky equilibrium, in which each agent receives a fixed share of the aggregate endowment and the agents are not allowed to trade future consumption claims. Figure 2 illustrates the trade-off for the bad state. (The corresponding figure for the good state is Figure 16 in Appendix F.)

Figure 2: Conditional Mean and Variance of Continuation Utilities

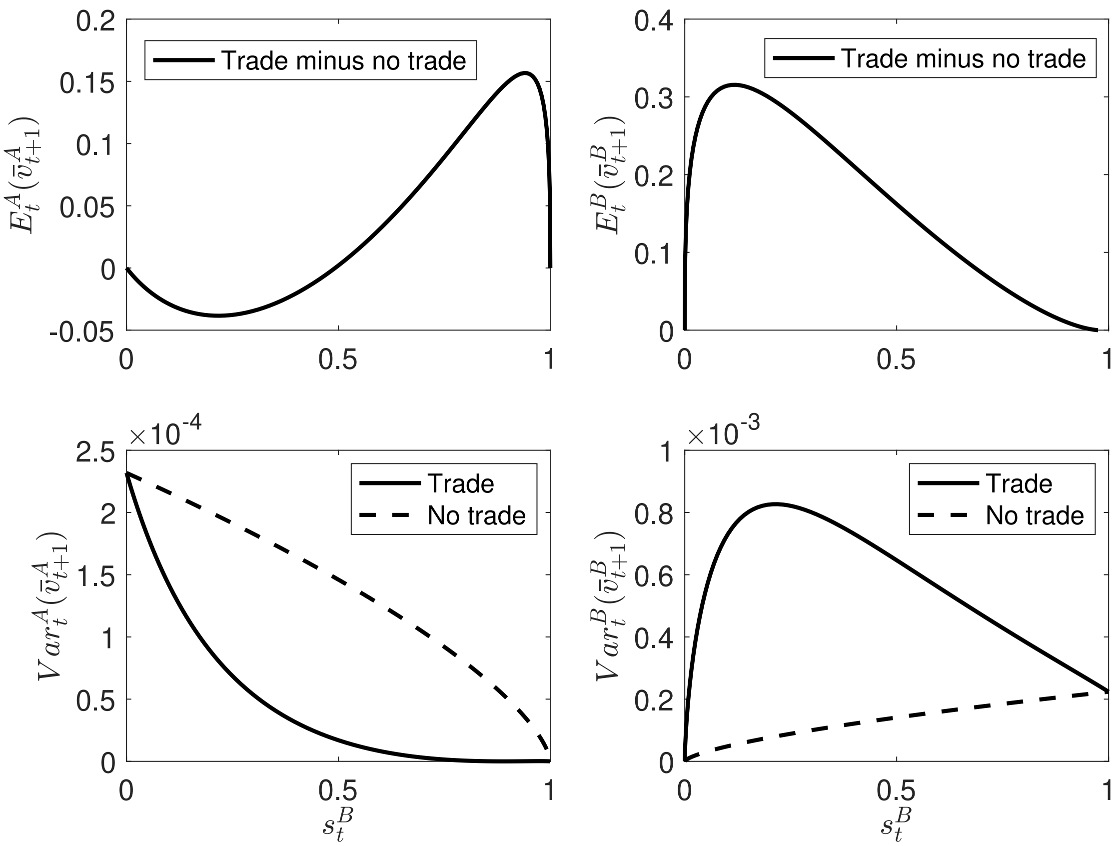

The figure plots the conditional mean and the variance of the continuation utilities as a function of the consumption share of investor B for the discrete-state economy. The top graphs show the difference in expected utility between an economy in which the agents are allowed to trade and an economy with no trade. The lower panel shows the conditional variance for both, the trade and no-trade case. The left graphs show the results for investor A, who believes in higher persistence and the right graphs show the results for investor B, who believes in lower persistence. The results are shown for economy being in the bad state 1.

The lower panel shows the conditional variances of the investors for the trade and no-trade setup. Investor B is willing to take on significantly more variance compared to the no-trade case. In return for taking on this additional variance, investor B subjectively receives additional future utility. So investor B is willing to take on additional risk in return for higher future wealth. We can think of this as if investor B is willing to provide insurance against the higher volatility in the bad state. On the flip side, investor A can significantly reduce the variance of her utility by trading with the other investor (left graphs). In return she is willing to give up part of her future utility if her consumption share is large (small s_t^B). For large s_t^B, there is so much supply of the insurance for the bad state that investor A can simultaneously reduce her variance and achieve larger expected continuation utility compared to the no-trade regime. The willingness to give up some of her utility to reduce future variance prevents the traditional market selection argument from always holding. Even if investor A has the correct beliefs, she is giving up future wealth to avoid volatility for large s_t^A. This prevents agent B from being driven out of the market. To demonstrate this mechanism, we explain how consumption shares change in terms of the two motives to trade—the speculation motive and the risk-sharing motive. This is crucial to understand the effects of trading on asset prices which we discuss in Section 4.

For this explanation, we first consider the case of CRRA preferences, for which the risk-sharing motive is absent and so we can analyze the speculation motive in isolation. This setup leads to the classical market selection result: in the long-run only the investors with correct beliefs survive. Recall that the changes in the consumption shares are given by

$$\frac{s_{t+1}^A}{s_{t+1}^B} = \left(\frac{\mathrm{dP}_{t,t+1}^A}{\mathrm{dP}_{t,t+1}^B}\right)^{\frac{1}{\gamma}} \frac{s_t^A}{s_t^B}, \tag{24}$$

see equation (18). As explained above, the changes in the consumption shares only depend on the subjective beliefs about the states in $t+1$ and the degree of risk aversion γ. Importantly, they do not depend on the value of the current state x_{z_t} itself, even though it affects the long-run growth rate of the economy. In other words, whether $x_{z_t} = \{-0.0024, 0.0024\}$, as we assume, or $x_{z_t} = \{-1000, 1000\}$ has no effect on the consumption sharing rule. The values of x_{z_t} will affect the value functions of the investors. But as the consumption shares do not depend on continuation utility for the CRRA case (see equation (18)), the values of x_{z_t} have no effect on the consumption shares. So the investors trade solely based on their subjective probabilities about the state in the next period. For $\gamma = 10$, the multiplicative effect on the consumption-share ratio, $\left(\frac{\mathrm{dP}_{t,t+1}^A}{\mathrm{dP}_{t,t+1}^B}\right)^{\frac{1}{\gamma}}$, is given by the corresponding value in the matrix

$$\begin{bmatrix} 1.0015 & 0.8706 \\ 0.8706 & 1.0015 \end{bmatrix}.$$

For example, when the economy is in the bad state 1 at time t and remains in that state at time $t+1$, then $\frac{s^A_{t+1}}{s^B_{t+1}} = 1.0015\frac{s^A_t}{s^B_t}$, and so the consumption share of agent A increases relative to that of agent B. When the economy is in state 1 at time t and transitions to the "good" state 2 at time $t+1$, then $\frac{s^A_{t+1}}{s^B_{t+1}} = 0.8706\frac{s^A_t}{s^B_t}$, and so the consumption share of agent A decreases relative to that of agent B. While we use specific numbers here to demonstrate the channel, it is straightforward to show that, as long as $\rho^A_x > \rho^B_x$, the multipliers are

$$\begin{bmatrix} >1 & <1 \\ <1 & >1 \end{bmatrix}.$$

As investor A believes in higher persistence and thus she always puts a larger weight on the economy staying in the same state relative to investor B, her consumption share always increases if the economy either remains in state 1 or in state 2, while the share of investor B increases if the economy transitions from one state to the other one.

Next we turn to the case of Epstein–Zin preferences. In this case, the investors have the additional incentive to share risks, which is reflected in the subjective mean-variance trade-off in consumption utility. For Epstein–Zin utility, the equilibrium consumption shares are given by

$$\frac{s^A_{t+1}}{s^B_{t+1}} = \frac{s^A_t}{s^B_t}\left(\frac{\mathrm{dP}^A_{t,t+1}}{\mathrm{dP}^B_{t,t+1}}\right)^{\psi}\left(\frac{\frac{\left(v^A_{t+1}\right)}{R^A_t\left(v^A_{t+1}\right)}}{\frac{\left(v^B_{t+1}\right)}{R^B_t\left(v^B_{t+1}\right)}}\right)^{1-\gamma\psi},$$

see again equation (18). Recall that for Epstein–Zin preferences the changes in the consumption shares do not only depend on the subjective probabilities, but also on the continuation utilities of the investors. This property introduces the new risk-sharing mechanism as investors would like to insure against states that they particularly dislike. Assume, for example, that $s^A_t = 0.5$. For the calibration of this example with EZ-utility the multiplicative effect on the consumption-share ratio,

$$\left(\frac{\mathrm{dP}^A_{t,t+1}}{\mathrm{dP}^B_{t,t+1}}\right)^{\psi}\left(\frac{\frac{\left(v^A_{t+1}\right)}{R^A_t\left(v^A_{t+1}\right)}}{\frac{\left(v^B_{t+1}\right)}{R^B_t\left(v^B_{t+1}\right)}}\right)^{1-\gamma\psi},$$

takes the corresponding value in the matrix

$$\begin{bmatrix} 1.0005 & 0.6699 \\ 1.0363 & 1.0010 \end{bmatrix}.$$

For example, when the economy is in the bad state 1 at time t and transitions to the good state 2 at time $t+1$, then $\frac{s^A_{t+1}}{s^B_{t+1}} = 0.6699\frac{s^A_t}{s^B_t}$, and so the consumption share of agent A decreases substantially relative to that of agent B. The two off-diagonal elements are quite different for EZ-utility than for CRRA utility. The investors have a preference for the early resolution of risk, so they dislike shocks that affect the long-run growth rate (see Bansal and Yaron (2004)). This preference is especially strong for investor A who believes in a higher persistence. Assume for example that the economy is in the good state. Investor A believes in higher persistence of the low growth state 1 compared to agent B, so she particularly dislikes if the economy enters this state (her continuation utility v^A_{t+1} is low relative to the certainty equivalent in the good state $R^A_t\left(v^A_{t+1}\right)$). When the economy enters the bad state, investor B is willing to forego some of his consumption share and compensates the utility loss of investor A (the multiplier is 1.0363 and hence there is a significant increase compared to the multiplier of 0.8706 for CRRA utility). On the other hand, investor A also believes in higher persistence of the good state 2. So she believes that periods of high growth are longer lasting. If the economy goes from the bad state 1 to the good state 2, her continuation utility will be higher compared to the utility of agent B and hence she is willing to forego some of her consumption in this case (the multiplier is 0.6699 compared to the multiplier of 0.8706 for CRRA utility).

There is also a striking difference in the *subjective* experience of the agents in the Epstein–Zin case compared to the CRRA case. For CRRA utility, both investors believe that their consumption share is going to increase on average (see Appendix D for an analytic derivation of this result). For the Epstein–Zin case, investor A is deliberately willing to give away part of her consumption share to obtain compensation whenever the economy enters a recession state. Figure 3 shows the expected changes in the consumption shares of the two investors. The left graph shows the change in the consumption share of investor A under her own subjective beliefs and the right graph shows the change in the consumption share of investor B under his beliefs.

For CRRA utility (dotted lines), both investors believe that their consumption share is going to increase on average. This property is the reason why under CRRA preferences the agents make "speculative bets." Which investor then becomes wealthier over time is purely determined by the true distribution of the underlying process. If investor A (B) holds the correct beliefs, she (he) will drive investor B (A) out of the market.

For the Epstein–Zin case (solid lines) and small s^B_t under both, the subjective beliefs of investor A and those of investor B, the consumption share of investor B is increasing on average. In other words, investor A is deliberately giving away part of her consumption share, on average. She is willing to accept this trade, as she obtains an "insurance" payment from

Figure 3: Expected Change in the Consumption Shares in the Discrete-State Model

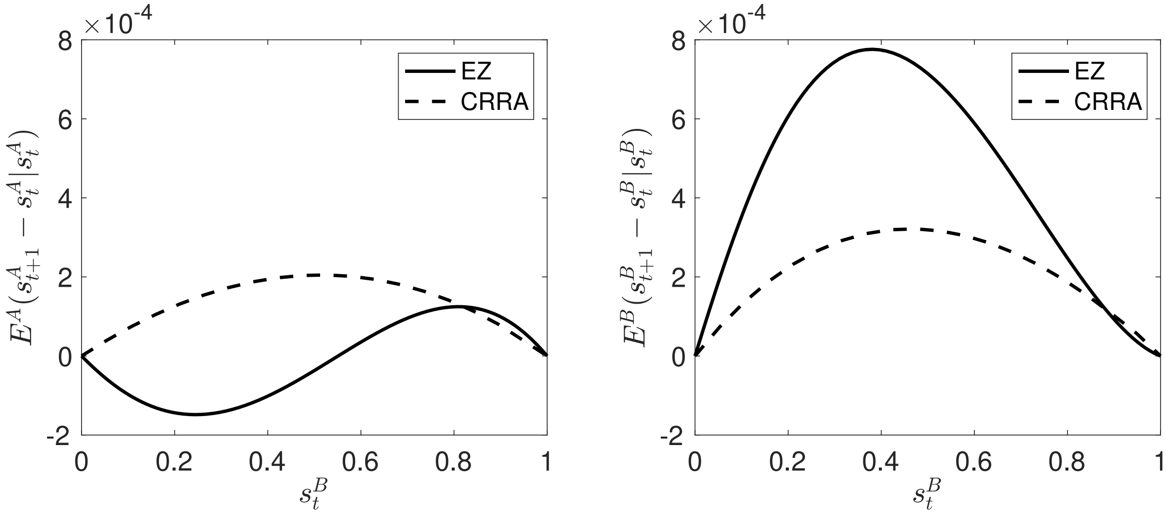

The figure shows the expected changes in the consumption shares $s^h_{t+1} - s^h_t$, $h \in \{A, B\}$, as a function of s^B_t for the stylized economy with a two-state Markov chain for x_t. The left graph shows the expected change in the consumption share of agent A under the subjective beliefs of agent A. The right graph shows the expected change in the consumption share of agent B under the subjective beliefs of agent B. The solid lines show the expected changes for Epstein–Zin preferences and the dashed lines show the corresponding results for CRRA preferences.

investor B, whenever the economy transitions from the good state 2 to the bad state 1. The larger s^B_t, the larger is the supply and the lower is the demand for the insurance. Hence, the insurance becomes cheaper and so the influence of the speculation motive increases. So for large s^B_t, investor A believes that her consumption share increases on average also for the Epstein–Zin case.

This concludes the analysis on why and how the investors trade in the two-investor two-state economy. In the following, we transfer the findings to the continuous-state models and analyze the effects of shifting consumption shares and asset trading on asset prices.

4 Belief Differences and Asset Prices

In the following we analyze the influence of belief differences for the fully stochastic infinite-horizon economy with the exogenous processes given in system (16). We first briefly explain how the speculation and risk-sharing motives affect the equilibrium outcomes and then analyze the asset-pricing implications of the belief differences.

4.1 Speculation and Risk-Sharing Motives

This section examines how the consumption shares change over time in the full model. We rely on the discussion of speculative and risk-sharing motives from Section 3.2. Figure 4 shows the change in the consumption share of investor B, $s_{t+1}^B - s_t^B$, as a function of s_t^B. The three graphs display results for $x_t = 0$ and for ± 1 unconditional standard deviation of x_t around its unconditional mean of $x_t = 0$. The blue lines depict the case of a negative shock in x_{t+1} ($x_{t+1} - \rho_x x_t = -0.0005$), the yellow lines that of a positive shock in x_{t+1} ($x_{t+1} - \rho_x x_t = 0.0005$), and the red lines show the average over all shocks.

Figure 4: Changes in the Consumption Shares for $\rho_x = \rho_x^1 = 0.99$ and $\rho_x^B = 0.96$

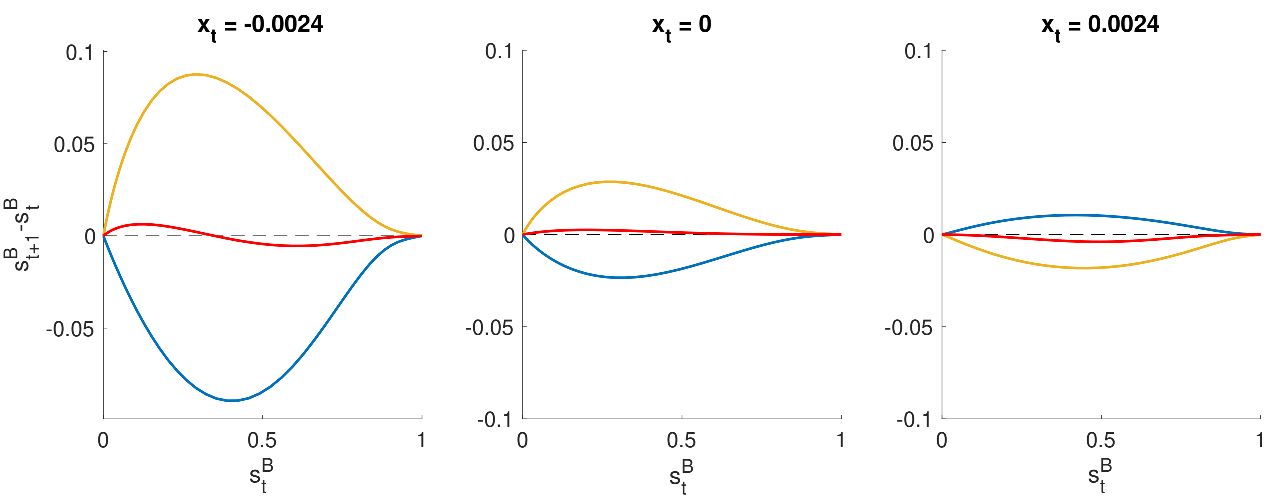

The figure shows the change in the consumption share $s_{t+1}^B - s_t^B$ as a function of s_t^B. Results are shown for $x_t = 0$ and for ± 1 unconditional standard deviations of x_t around its unconditional mean of $x_t = 0$. The blue lines depict the case of a negative shock in x_{t+1} ($x_{t+1} - \rho_x x_t = -0.0005$), the yellow lines that of a positive shock in x_{t+1} ($x_{t+1} - \rho_x x_t = 0.0005$), and the red lines show the average over all shocks. Agent A has the correct beliefs with $\rho_x^A = \rho_x = 0.99$ and agent B has the belief $\rho_x^B = 0.96$.

In the special case of $x_t = 0$ (center panel), both investors agree on the distribution of x_{t+1} and hence there is no speculation motive. In a model with identical CRRA preferences, investors would therefore not be willing to trade and there would be no changes in the consumption shares; see also the closed-form solutions derived in Appendix D. As shown in Section 3.2, for Epstein–Zin preferences, investor A ($\rho_x^A = 0.99$) has a stronger aversion than investor B to negative shocks to x_t. A low x_{t+1} implies a relatively lower utility of agent A compared to agent B. Investor B is willing to absorb some of this utility risk in return for higher future wealth; recall also Figure 2. Hence, the consumption share of investor B increases on average (red line). In case a bad shock materializes (low future x_t), investor A obtains compensation from investor B for her utility loss, hence the consumption share of type A investors increases (blue line). In return, investor B receives a larger consumption share in future good states (yellow line).

Next, consider the case of a positive x_t (right panel), when the economy is in a good state. In this case, the demand for risk sharing is much smaller than for $x_t \leq 0$; compare also Figure 2 to Figure 16 in the appendix. Now low values of x_{t+1} are significantly less likely due to the high persistence of the long-run-risk shocks. However, the speculation motive increases as investors now disagree on the state tomorrow. Investor B believes (incorrectly) in faster mean reversion—that is, he has stronger beliefs in smaller values for x_{t+1} than investor A does. Put differently, investor B bets on smaller values of x_{t+1}. Therefore, his consumption share increases for a negative shock to x_t. As investor A has the correct belief, her consumption share increases, on average. So in this case, the speculation motive dominates the risk-sharing motive.

Finally, we turn to the case of a negative x_t (left panel), when the economy is in a bad state. Now both the risk-sharing and the speculation motives are present and strong. Investors disagree on the future state of the economy (speculation motive) and, as the economy is in a bad state, investor A would like to hedge against a further decline into an even worse state (risk-sharing motive). The speculation motive implies that investor B bets money on more positive shocks to x_{t+1} as he believes in faster mean reversion; recall Section 3.2. So in this case, both the risk-sharing and the speculation motive imply that the consumption share of investor A increases for a negative shock to x_{t+1} while, on the contrary, the share of investor B increases for a positive shock to x_{t+1}. Hence, the magnitude of the changes in consumption shares is much larger compared to the other cases. Which effect dominates on average? For an explanation of the average effect, we need to take into account how the consumption share influences the risk-sharing motive. Recall from Figure 2 that a small amount of type B investors is sufficient to induce a lot of risk sharing, while the amount of risk sharing decreases for large s_t^B. Hence, for small s_t^B, type B investors are compensated by higher future wealth in return for their willingness to take on additional risk. Therefore, the consumption share of investor B increases on average. For large shares s_t^B, there is less demand from type-A investors for risk sharing and, therefore, investor B takes on less of the variation in future wealth and also obtains less compensation in future wealth; recall again Figure 2. The speculation motive transfers wealth on average to investor A who holds the correct beliefs, which dominates the effect of the risk-sharing motive for sufficiently large values of s_t^B. Hence, for large s_t^B, the consumption share of investor B decreases, on average, in the left panel.

The aggregate dynamics of these effects induce changes in the wealth distribution which in return affect asset prices. In the following we analyze this mechanism and show that it helps to resolve several financial-market patterns that have been puzzling from the viewpoint of theoretical models.

4.2 Portfolio Allocations and Asset Prices

In this section we analyze the portfolio allocations and asset-pricing moments in the heterogeneous-agent economy. Since we consider a discrete-time economy with continuous states and complete markets, the agents' portfolio choices involve positions in an infinite set of Arrow–Debreu securities. Clearly, it is impossible to describe continua of portfolio holdings. Instead, as proposed by Collin-Dufresne, Johannes, and Lochstoer (2016a), we examine the changes in investors' wealth shares to obtain an intuition about investors' security holdings. (Broadly speaking, we would expect that an investor who has a more optimistic outlook about the state of the economy is willing to invest in more risky assets and, therefore, her or his wealth share increases (decreases) in good (bad) future states of the economy state.)

To analyze the change in the wealth shares, we define individual wealth, W_t^h, as the claim to the future consumption stream of the respective agent. Total wealth in the economy, W_t, is given by the sum of the individual wealth levels. For Epstein–Zin utility, the individual wealth-consumption ratio is given by

$$\frac{W_t^h}{C_t^h} = \frac{1}{1-\delta^h}\left(\frac{V_t^h}{C_t^h}\right)^{\rho^h} \tag{25}$$

from which we can compute the wealth shares $\frac{W_t^h}{W_t}$ of the investors. Figure 5 shows the change in the wealth share $\frac{W_{t+1}^B}{W_{t+1}} - \frac{W_t^B}{W_t}$ as a function of the shock to x, $\eta_{x,t+1}$, for different wealth share levels $\frac{W_t^B}{W_t}$. Consider first the center panel which depicts the case of $\frac{W_t^B}{W_t} = 0.5$. For $x_t = -0.0024$ (blue line), investor B is more optimistic than investor A, since he believes in faster mean reversion. Hence, he invests in riskier assets which pay off in good future states; as a result, his wealth share increases for positive $\eta_{x,t+1}$. For the good state ($x_t = 0.0024$, yellow line), the opposite holds true. Investor B believes that the expansion is shorter lived compared to investor A's beliefs, and thus he invests to insure against a downturn. Then his wealth increases if a bad shock hits the economy ($\eta_{x,t+1} < 0$). In Section 3.2, see the explanations for Figure 2, we have seen that a small fraction of type-B investors is sufficient for the strong risk-sharing motive to become visible. Hence, for $\frac{W_t^B}{W_t} = 0.1$, the changes in the wealth shares are of similar magnitude compared to the $\frac{W_t^B}{W_t} = 0.5$ case. On the contrary, for

Figure 5: Changes in the Wealth Shares for $\rho_x = \rho_x^1 = 0.99$ and $\rho_x^B = 0.96$

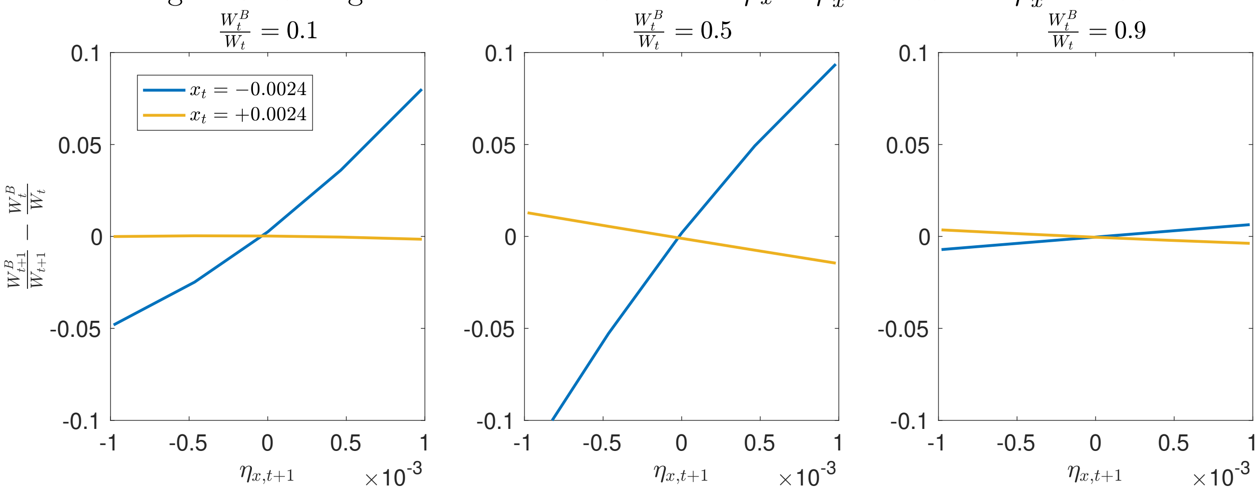

The figure shows the change in the wealth share of investor B, $\frac{W_{t+1}^B}{W_{t+1}} - \frac{W_t^B}{W_t}$, as a function of the shock to x_t, $\eta_{x,t+1}$. The three graphs show the results for $\frac{W_t^B}{W_t} \in \{0.1, 0.5, 0.9\}$, respectively. The blue lines depict the case of a bad x-state ($x_t = -0.0024$) and the yellow lines that of a good x-state ($x_t = 0.0024$). Agent A has the correct beliefs with $\rho_x^A = \rho_x = 0.99$ and agent B has the belief $\rho_x^B = 0.96$.

a large fraction of type-B investors, see the right panel, the possible risk sharing decreases significantly and thus the effect on the wealth shares becomes much smaller.

Next we analyze how the changes in the consumption shares induced by the risk-sharing and speculation motives affect asset prices. The stochastic discount factor for Epstein–Zin preferences is given by

$$M_{t+1}^h = \delta^h e^{(\rho^h - 1)\Delta c_{t+1}} \left(\frac{s_{t+1}^h}{s_t^h} \right)^{\rho^h - 1} \frac{\left(v_{t+1}^h e^{\Delta c_{t+1}} \right)^{\alpha^h - \rho^h}}{R_t^h \left(v_{t+1}^h e^{\Delta c_{t+1}} \right)^{\alpha^h - \rho^h}}, \tag{26}$$

which we can use to price the claim to aggregate dividends. Figure 6 plots the log price-dividend ratio $p_t - d_t$ as a function of x_t and s_t^B. Prices increase in x_t as the investors have a preference for the early resolution of risks; see Bansal and Yaron (2004). Furthermore, as $\rho^A > \rho^B$, investor A requires a larger return than investor B for holding the risky dividend stream and hence the price-dividend ratio increases with s_t^B. As investor B is willing to take on more risk in the economy, even if he has only a small consumption share, the increase in the price-dividend ratio is particularly strong for small s_t^B and negative x_t. Hence, introducing a small fraction of type-B investors—with their belief in lower persistence and faster mean reversion of long-run risk—to the model has large effects on asset prices. In particular, the heterogeneous-agent setup adds additional countercyclical movements and excess volatility to

Figure 6: Log Price-Dividend Ratio for $\rho_x = \rho_x^1 = 0.99$ and $\rho_x^B = 0.96$

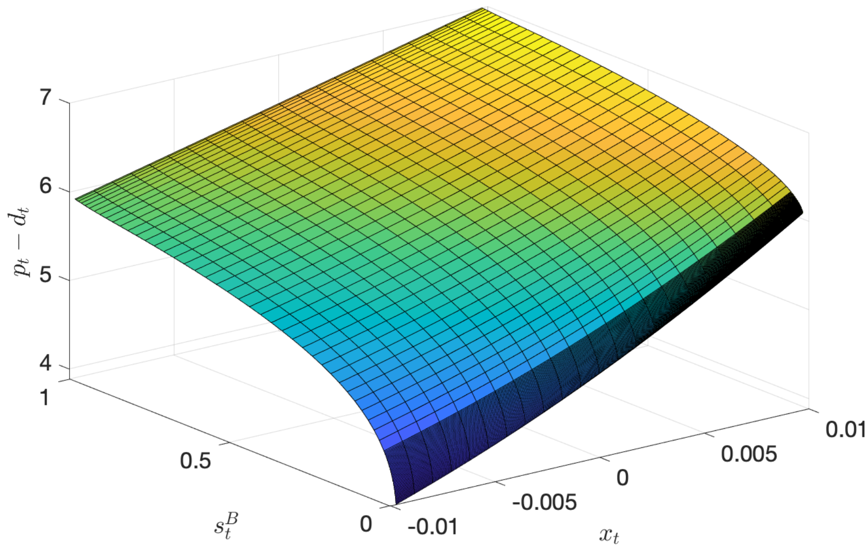

The figure plots the graph of the log price-dividend ratio $p_t - d_t$ as a function of x_t and s_t^B. Agent A has the correct beliefs with $\rho_x^A = \rho_x = 0.99$ and agent B has the belief $\rho_x^B = 0.96$.

the price-dividend ratio. In the representative-agent economy, a bad shock to x_t also leads to lower prices. In the two-agent economy, this effect is amplified as a bad shock implies that the insurance pays off and hence the consumption share of investor A increases, see also again Figure 4. This, in turn, lowers the demand of investor B for the risky dividend stream and hence prices decrease even further. Next we show that this mechanism is quantitatively strong and can lead to significant excess volatility.

We simulate 1,000 samples each containing 77 years of data initialized at the long-run mean of $s_0^B = 0.5$.[6] The number of years is chosen to make the results comparable to the empirical estimates that we take from Bansal, Kiku, and Yaron (2012). As described above, the economy is calibrated such that the long-run mean consumption share is $s_t^B = 0.5$; we observe an average standard deviation of the share s_t^B of about 0.18—that is, the consumption shares display quite significant variations in response to exogenous shocks due to the agents' risk-sharing and speculation motives.

Table 2 reports the annualized asset-pricing moments for the two-agent economy as well as for the representative-agent economies populated by either of the two investors.[7] All three first moments for the heterogeneous-agent economy lie between the boundary cases of the

[6] We do not encounter any survival issues in this setup. Even after 1,000 years of simulated data, the consumption shares along all paths remain significantly positive with a long-run mean share of about $s_t^B = 0.5$.

[7] We assume that in the representative-agent economies, the respective investor has the correct beliefs.

Table 2: Annualized Asset-Pricing Moments

	$E(p_t - d_t)$	$\sigma(p_t - d_t)$	$AC1(p_t - d_t)$	$E\left(R_t^m - R_t^f\right)$	$E\left(R_t^f\right)$	$\sigma(R_t^m)$	$\sigma\left(R_t^f\right)$
$s_t^B = 0$	2.36	0.28	0.79	12.52	1.72	18.66	1.65
Two-Ag.	3.52	0.38	0.80	5.42	2.29	21.79	1.92
$s_t^B = 1$	3.97	0.11	0.38	1.78	2.93	14.08	0.88
Data	3.36	0.45	0.87	7.09	0.57	20.28	2.86

The table shows selected moments from 1,000 samples—each containing 77 years of simulated data—starting with an initial share of $s_0^B = 0.5$. It reports the mean, the standard deviation, and the first-order autocorrelation of the annualized log price–dividend ratio as well as the mean and the standard deviation of both the annualized equity premium and the risk-free return. Agent A has the correct beliefs with $\rho_x^A = \rho_x = 0.99$; agent B has the belief $\rho_x^B = 0.96$. All returns are shown in percentages, so a value of 1.5 is a 1.5% annualized figure. The estimates from the data are taken from Bansal, Kiku, and Yaron (2012).

representative-agent economies. In line with the data, we observe a large and significant equity premium for the two-agent economy of 5.42 percent.[8] More interestingly, the volatilities of prices and returns are significantly larger for the two-agent economy than for both representative-agent economies. As argued above, this significant excess volatility is generated by the shifts in the wealth distribution in response to the exogenous shocks. As the wealth distribution shifts consumption weights back and forth between the two types of agents, the second moments increase in response to the time variation in the wealth shares. So instead of adding exogenous features to the model as for example habit preferences (Campbell and Cochrane, 1999), time varying disaster risks (Wachter, 2013) or stochastic volatility (Bansal and Yaron, 2004), the excess volatility is generated endogenously by the agents' belief differences which arise naturally given the properties of the unobserved x_t process (see Section 3.1).

But even with exogenous stochastic volatility, Beeler and Campbell (2012) argue that the long-run risk model is not able to generate as much volatility of the price–dividend ratio as observed in the data. They report a value of 0.18 for the calibration of Bansal, Kiku, and Yaron (2012), which deliberately increases the influence of stochastic volatility, compared to 0.45 observed in the financial market data. Our results show that differences in beliefs can resolve this puzzle, since they lead to a significant increase in excess volatility close to the

[8] The level of the risk-free rate is slightly too high in the two-agent economy compared to the reported data estimate. Recall that we have deliberately taken all parameters from Bansal and Yaron (2004) and only varied the belief in the persistence parameter. A lower risk-free rate could be achieved by slightly increasing the subjective discount factor δ.

value observed in the data.

4.3 Time Variation in Expected Risk Premia

In the next step of our analysis, we demonstrate that the model generates significant countercyclical time variation in expected risk premia. Figure 7 shows the annualized expected risk premium as a function of s_t^B and the expectations are calculated under the true beliefs, $\rho_x = 0.99$. We find that the expected risk premium changes significantly with the consumption share. More precisely, the expected risk premium is large when investor A holds a larger consumption share, while it is significantly smaller when agent B dominates. This effect is particularly strong for bad states—that is, for negative values of x_t. Why do risk-premia move

Figure 7: Time Variation in Expected Risk Premia

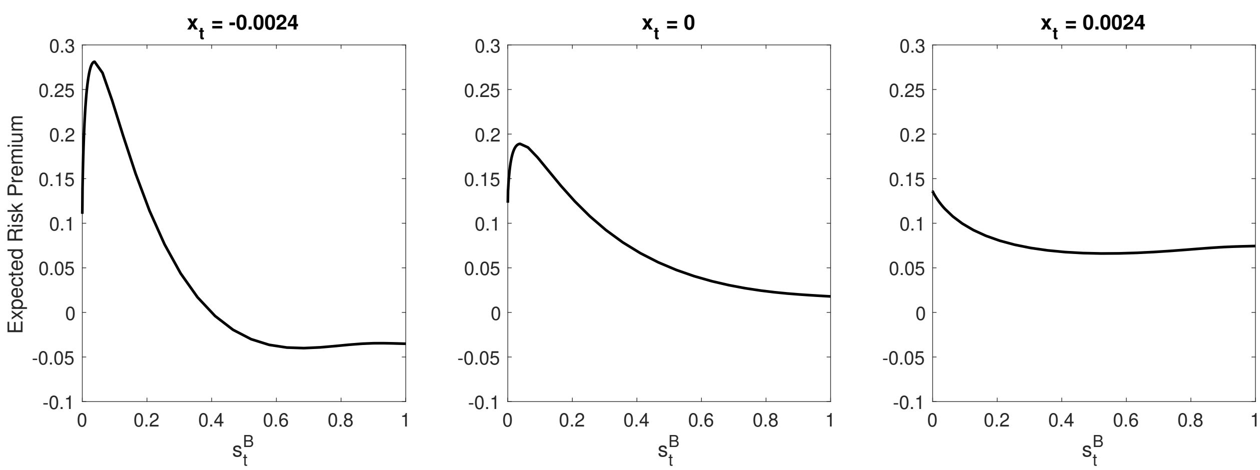

The figure shows the annualized expected risk premium as a function of s_t^B. Results are shown for $x_t = 0$ and for ± 1 unconditional standard deviations of x_t around its unconditional mean of $x_t = 0$. Agent A has the correct beliefs with $\rho_x^A = \rho_x = 0.99$ and agent B has the belief $\rho_x^B = 0.96$.

with the consumption shares of the investors? Consider first the case of $x_t = 0$ or negative x_t. In Section 4.1, we have shown that for small s_t^B the influence of the risk-sharing motive is large (see, for example, Figure 4). So if s_t^B is small, the consumption share of investor B is expected to increase. As an increase in s_t^B leads to larger prices (see Figure 6), the expected increase in s_t^B implies an increase in the expected risk premium. Consistent with the results from Section 4.1, this effect is even stronger for $x_t = -0.0024$, when risk prices are larger and, hence, the consumption share of investor B is expected to increase even faster. For $x_t = 0.0024$ the speculation motive is stronger, consumption shares move less (see Figure 4) and prices are less sensitive to changes in the consumption shares (see Figure 6). Hence, the effect vanishes and the expected risk premium is relatively flat as a function of s_t^B.

Furthermore belief heterogeneity can increase the model-implied expected risk premium compared to the representative-agent cases. The changes in the consumption shares add an additional source of risk for the investors which can potentially lead to a larger risk premium (see Gârleanu and Panageas (2015)). The changes in the wealth distribution are highly persistent due to the high persistence of x_t (we obtain a persistence of s_t^B of 0.9926 and a standard deviation of 0.18). In the case of CRRA preferences, the persistent changes have a negligible effect as investors only care about the change in the consumption share in the subsequent period (see equation (26)). For EZ preferences with a preference for the early resolution of risks, the effect of the changes in the consumption shares on continuation utility is large (see Figure 2) and hence investors care about the persistent changes in the wealth distribution. As the risky asset is exposed to these changes, investors require an additional premium for holding the stock. Figure 4 shows that for small s_t^B, the expected changes in the wealth distribution are particularly large. Hence, for small but positive values of s_t^B, the premium is larger than in the representative-agent case of $s_t^B = 0$. Put differently, belief heterogeneity endogenously generates priced consumption risk—that is, investors require a premium for the expected changes in the wealth distribution.

In the following we argue that the variation in risk-premia is countercyclical and economically significant. Recall from Section 4.1 that in normal and bad times negative shocks to x_t will increase the consumption share of investors of type A, who believe in a larger value for ρ_x, as these investors buy insurance against such bad shocks (see Figure 4). So when the economy enters a recession—that is, for a series of negative shocks to x_t—the consumption share of investor A increases. As a result, we observe a quite dramatic increase in the expected risk premium for negative values of x_t and small consumption shares of investor B; see Figure 7. Hence, the endogenous movements in the wealth distribution lead to countercyclical variations in the expected equity risk premium. This variation is economically large in our calibrated economy.

Table 3 reports the mean and the standard deviation of the annualized expected risk premia for the simulated data. Martin (2017) constructs a lower bound for the expected risk premium in terms of its risk-neutral variance and shows that there is significantly more time variation in the premium than previous studies have shown. We test the ability of the heterogeneous-beliefs model to explain this finding and compare it to the standard long-run risk model. Martin (2017) reports a mean expected risk premium of 5% per year with a standard deviation of 4.6%. The heterogeneous-agent setup implies similar numbers, with a mean of 5.42% and a standard deviation of 5.73% (see Table 3). In the standard long-run risk model without stochastic volatility, expected risk premia are constant as shown by Bansal and Yaron (2004).[9] They propose to add stochastic volatility to the model to generate countercyclical time variation in risk premia as observed in the data. Table 3 shows that while the model of Bansal and Yaron (2004) is able to generate a large mean premium, the standard deviation is considerably smaller (1.13%) than in the data. In the new calibration of Bansal, Kiku, and Yaron (2012), the influence of the stochastic volatility channel is increased in order to match second-order moments. This adjustment should also generate more volatility in expected risk premia. However, while the value of Bansal and Yaron (2004) is slightly improved, it is still considerably lower (1.61%) than in the data. In sum, we observe that

Table 3: Expected Risk Premia

	Data	Het. Agents	BY (2004)	BKY (2012)
Mean	5.00	5.42	5.59	6.38
Std. dev.	4.60	5.73	1.13	1.61

The table shows the annualized mean and volatility of the expected risk premium. The first column shows the empirical values reported in Table 1 of Martin (2017). The second column shows the results for the heterogeneous-agent setup with $\rho_x^A = \rho_x = 0.99$ and $\rho_x^B = 0.96$ for the correct beliefs. The data is obtained from 1,000 samples—each containing 77 years of simulated data—starting with an initial share of $s_0^B = 0.5$. Columns three and four show the results for the models of Bansal and Yaron (2004, Case II) and Bansal, Kiku, and Yaron (2012), respectively, using the same size for the simulated data set. All returns are shown in percentages, so a value of 1.5 is a 1.5% annualized figure.

belief heterogeneity in the long-run risk model provides a solution that accounts for the large variation in expected risk premia reported in the data.

4.4 Predictability of Returns, Consumption, and Dividends

Next we examine the implications of the two-investor economy for the predictability of returns and cash flows. Beeler and Campbell (2012) argue that in the long-run risk model, the price–dividend ratio has too much predictive power for consumption and dividend growth while the predictability of returns is too low compared to the values observed in the data. Bansal, Kiku,

[9] Risk premia in the standard long-run risk model without stochastic volatility are only constant when solving the model using a log-linearization. Solving the model with a global method shows that risk premia are time-varying and pro-cyclical. But the variation is rather small for standard calibrations; see Pohl, Schmedders, and Wilms (2016).

and Yaron (2012) propose a solution by increasing the importance of the stochastic volatility channel and decreasing the importance of the growth-rate channel (x_t). Beeler and Campbell (2012) acknowledge that the statistical rejections of the model are less extreme for this calibration, but it requires extremely persistent movements in the volatility process. We show that the heterogeneous-agent economy helps to resolve the predictability puzzle. Furthermore, the heterogeneous-agent model can reconcile the evidence on the countercyclical variation of return predictability (see Henkel, Martin, and Nardari (2011) and Dangl and Halling (2012))—that is, higher predictability in economically bad times and lower predictability in good times. The time-variation in the consumption shares endogenously generates the countercyclical pattern in the heterogeneous-agent model, while such an effect is absent in the representative-agent framework.

Table 4 reports R^2 statistics and regression coefficients from regressing cumulative log excess returns, consumption growth, and dividend growth on the lagged log price–dividend ratio. Statistics are shown for the annualized time series with one-, three-, and five-year horizons. We observe that the predictability of returns is low in the representative-agent economies. In the two-agent setup, prices become more volatile and revert to the mean of the stationary distribution. Therefore, there is more return predictability, which also increases with the horizon. For consumption and dividends, we obtain the opposite pattern. For the representative-agent economy, predictability increases significantly with the persistence level. The two-agent setup reduces the predictability compared to the high persistence representative-agent economy due to the endogenous variation in the consumption shares. Overall the predictability of cash flows is still too large compared to the data but the results show that investor heterogeneity can contribute to resolving the predictability deficiencies of the long-run risk model.

Furthermore, Henkel, Martin, and Nardari (2011) and Dangl and Halling (2012) provide empirical evidence that the predictability of returns by, for example, the price-dividend ratio is significantly larger in recessions than it is in expansions. The authors argue that these countercyclical variations in the risk premium can be explained by time-varying degrees of risk aversion as, for example, in the external habit model of Campbell and Cochrane (1999). Our results from Section 4.3 suggest that investor heterogeneity can add significant time variation in risk premia to the long-run risk model—if the consumption share of the investor with the larger beliefs about ρ_x increases, risk premia increase. To demonstrate this mechanism, we again run the predictability regressions for excess returns, but now including a slope dummy

Table 4: Predictability of Excess Returns, Consumption, and Dividends

	R^2			β		
	1Y	3Y	5Y	1Y	3Y	5Y
$\sum_h^H (r_{m,t+h} - r_{f,t+h}) = \alpha + \beta(p_t - d_t) + \epsilon_{t+H}$						
$s_t^B = 0$	0.0070	0.0190	0.0277	-0.0134	-0.0496	-0.1116
Two-Ag.	0.0129	0.0350	0.0540	-0.0563	-0.1676	-0.2639
$s_t^B = 1$	0.0067	0.0113	0.0150	-0.0226	-0.0824	-0.1215
Data	0.04	0.19	0.31	-0.09	-0.27	-0.43
$\sum_h^H (\Delta c_{t+h}) = \alpha + \beta(p_t - d_t) + \epsilon_{t+H}$						
$s_t^B = 0$	0.5121	0.5738	0.5075	0.0829	0.2187	0.3124
Two-Ag.	0.3043	0.3346	0.2825	0.0495	0.1283	0.1805
$s_t^B = 1$	0.1522	0.1438	0.1012	0.0930	0.1941	0.2165
Data	0.060	0.01	0.000	0.01	0.01	0.00
$\sum_h^H (\Delta d_{t+h}) = \alpha + \beta(p_t - d_t) + \epsilon_{t+H}$						
$s_t^B = 0$	0.4881	0.4903	0.4393	0.3044	0.7074	0.9705
Two-Ag.	0.3088	0.2868	0.2486	0.1882	0.4233	0.5910
$s_t^B = 1$	0.4283	0.2051	0.1305	0.6557	0.9187	0.9977
Data	0.09	0.06	0.04	0.07	0.11	0.09

The table reports R^2 statistics and regression coefficients from regressing cumulative log excess returns, consumption growth, and dividend growth on the lagged log price–dividend ratio from 1,000 samples—each containing 77 years of simulated data—starting with an initial share of $s_0^B = 0.5$. Statistics are shown for the annualized time series with one-, three-, and five-year horizons. The estimates from the data are taken from Bansal, Kiku, and Yaron (2012).

variable for the economy either being in a recession or being in an expansion:

$$\sum_{h}^{H}(r_{m,t+h} - r_{f,t+h}) = \alpha + \beta_1(p_t - d_t) + \beta_2 d_{\{exp,rec\}}(p_t - d_t) + \epsilon_{t+H}, \tag{27}$$

where H denotes the horizon for the cumulative excess returns and $d_{\{exp,rec\}}$ takes the value of 1 if the economy is in a recession (expansion) and zero otherwise. So we run the regressions twice, once only including the recession slope dummy and once only with the expansion slope dummy. We define a recession (expansion) as the long-run risk state, x_t, being below (above) one standard deviation around its unconditional mean.[10]

Table 5 shows the regression results. First consider the representative-agent cases, $s_t^B = 0$ and $s_t^B = 1$. In these cases the R^2 are almost the same for the regressions including the recession slope dummy and the regressions including the expansion slope dummy. Hence, the predictability of risk premia does not change with the business cycle. For the heterogeneous-agent economy, the R^2 values are significantly larger for the regressions including the recession slope dummy compared to the regressions with the expansion slope dummy. The coefficients β_2 for the recession (expansion) slope dummy are negative (positive), so the aggregate slope coefficient of the log price-dividend ratio, $\beta_1 + \beta_2$, is larger in absolute value in recessions compared to normal times and expansions. Put differently, risk premia are more sensitive to changes in the price-dividend ratio in recessions than they are in expansions. This countercyclical variation is generated by time-varying risk prices. A recession is characterized by a series of negative shocks to x_t, which will increase the consumption share of investor A who believes in stronger persistence (see Section 4.3). As investor A has a larger aversion against long-run risks compared to investor B, risk prices in the economy rise. So belief heterogeneity endogenously generates countercyclical variation in risk premia and hence can explain the larger degree of predictability in recessions compared to expansions found in the data. This mechanism is absent in the representative-agent model.

4.5 Endogenous vs. Exogenous Time Variation in the Persistence

A key implication of the heterogeneous-agent model is that—as the wealth distribution in the economy shifts—the agent with beliefs in higher persistence has sometimes a greater weight in the social welfare function, and sometimes a lesser weight. A natural comparison would be to

[10] To go from the simulated monthly frequency to the annual frequency, we define a recession (expansion) year as a year with at least six recession (expansion) months. The regression results are robust with regard to changes in these threshold values.

Table 5: Predictability in Recessions and Expansions

	R^2			β_1			β_2		
	1Y	3Y	5Y	1Y	3Y	5Y	1Y	3Y	5Y
					Recession Dummy				
$s_t^B = 0$	0.0180	0.0393	0.0538	-0.0085	-0.0377	-0.0876	-0.0035	-0.0122	-0.0058
Two-Ag.	0.0475	0.0880	0.1182	-0.1016	-0.2465	-0.3884	-0.0259	-0.0396	-0.0426
$s_t^B = 1$	0.0190	0.0287	0.0357	-0.0217	-0.0761	-0.1121	-0.0001	-0.0003	-0.0005
					Expansion Dummy				
$s_t^B = 0$	0.0188	0.0395	0.0538	-0.0075	-0.0224	-0.0666	-0.0017	-0.0066	-0.0104
Two-Ag.	0.0245	0.0597	0.0941	-0.0621	-0.1806	-0.2898	0.0084	0.0210	0.0336
$s_t^B = 1$	0.0199	0.0303	0.0378	-0.0190	-0.0773	-0.1233	-0.0003	-0.0004	-0.0021

The table reports R^2 statistics and regression coefficients from regressing cumulative log excess returns on the lagged log price–dividend ratio including a recession (expansion) slope dummy, $\sum_h^H (r_{m,t+h} - r_{f,t+h}) = \alpha + \beta_1(p_t - d_t) + \beta_2 d_{\{exp,rec\}}(p_t - d_t) + \epsilon_{t+H}$. Median values are shown for 1,000 samples—each containing 77 years of simulated data—starting with an initial share of $s_0^B = 0.5$. Statistics are shown for the annualized time series with one-, three-, and five-year horizons.

modify the representative-agent model to directly model time-varying shifts in persistence.[11] In the following section, we consider a representative agent who receives exogenous shocks to persistence and highlight the key differences to the heterogeneous-agent model. In Section 4.5.2, we offer some more advantages of the heterogeneous-agent model over a representative-agent model with additional exogenous processes.

4.5.1 A Representative-Agent Model with Exogenous Persistence Shocks

We consider a representative-agent long-run risk model in which $\rho_{x,t}$ is time varying,

$$\begin{aligned} \Delta c_{t+1} &= \mu_c + x_t + \sigma \eta_{c,t+1} \\ x_{t+1} &= \rho_{x,t} x_t + \sigma_x \eta_{x,t+1} \\ \rho_{x,t+1} &= \mu_\rho (1-\nu) + \nu \rho_{x,t} + \sigma_{\rho,x} \eta_{x,t+1} + \sigma_\rho \eta_{\rho,t+1} \\ \Delta d_{t+1} &= \mu_d + \Phi x_t + \phi_d \sigma \eta_{d,t+1}, \end{aligned} \tag{28}$$

where $\eta_{c,t+1}$, $\eta_{x,t+1}$, $\eta_{\rho,t+1}$, and $\eta_{d,t+1}$ are independent and normally distributed with mean 0 and standard deviation 1. The persistence $\rho_{x,t}$ follows an AR(1) process that exhibits the time variation observed in the heterogeneous-agent economy. For this purpose we define the

[11] We thank an anonymous referee for suggesting this version of the representative-agent model as an alternative.

aggregate persistence in the heterogeneous-agent economy as the weighted average of the persistence levels of the two agents, where the weights are given by the endogenous consumption shares,

$$\rho_{x,t}^{Agg} = s_t^B \rho_x^2 + (1 - s_t^B)\rho_x^1. \tag{29}$$

For the benchmark calibration in the current section, we observe the following moments for the aggregate persistence level in the heterogeneous-agent economy:

$$\begin{aligned} E(\rho_{x,t}^{Agg}) &= 0.9742, \\ \sigma(\rho_{x,t}^{Agg}) &= 0.0056, \\ AC1(\rho_{x,t}^{Agg}) &= 0.9926, \\ Corr(\rho_{x,t}^{Agg}, x_t) &= -0.2740. \end{aligned}$$

We calibrate the representative agent economy to match these numbers, and so we obtain the following parameters: $\mu_\rho = 0.9742, \nu = 0.9926$, $\sigma_{\rho,x} =$ -2.1760e-04 and $\sigma_\rho =$ 6.4425e-04. For the remaining parameters we use the standard calibration from the heterogeneous-agent model. We solve the model by ensuring that the process $\rho_{x,t}$ varies in the interval $[0.96, 0.99]$, so the state space is bounded by the two values of the heterogeneous-agent model. This ensures that the model is well behaved and $\rho_{x,t}$ is bounded well below 1. Figure 8 shows annualized expected risk premia in the model as a function of $\rho_{x,t}$.

Figure 8: Time Variation in Expected Risk Premia in the Model with Time-Varying Persistence and a Negative Correlation

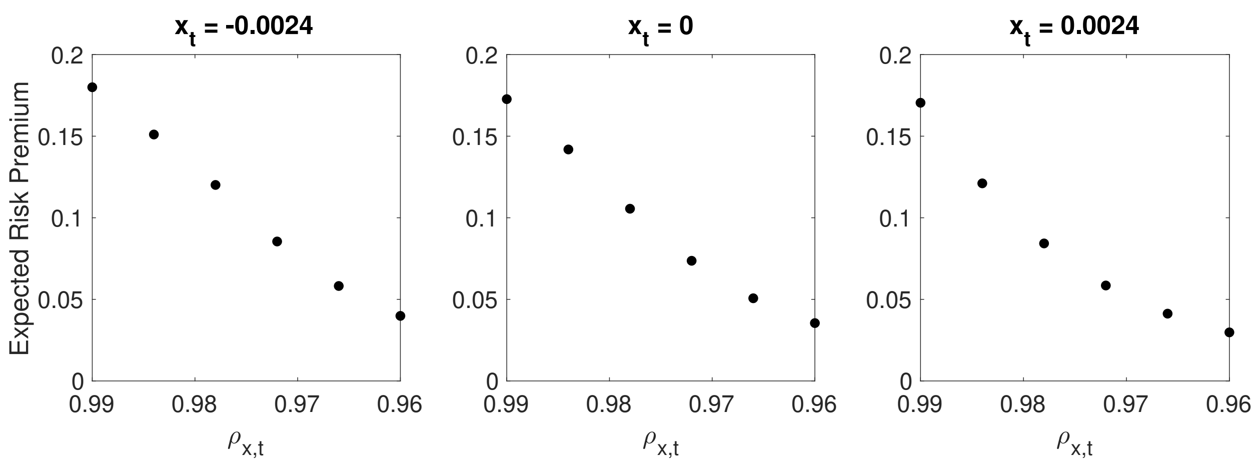

The figure shows the annualized expected risk premium as a function of $\rho_{x,t}$ for the representative-agent model with time-varying persistence and a negative correlation between x_t and $\rho_{x,t}$. Results are shown for $x_t = -0.0024$, $x_t = 0$ and $x_t = 0.0024$.

As expected, risk premia increase the higher $\rho_{x,t}$. However, in contrast to the heterogeneous-agent model, risk premia change only very little with the state x_t. What explains this key difference between the two models? In the heterogeneous-agent model, the state x_t has a large influence on future consumption shares and hence on the future aggregate persistence in the economy. Assume, for example, that the aggregate persistence in the economy is high (low s_t^B). If simultaneously x_t is low, the consumption share of investor B is expected to increase strongly due to the strong risk-sharing motive; see Figure 4. An increase in s_t^B implies higher future prices and so risk premia are expected to be high. Consider next the same scenario with a low s_t^B but now assume that x_t is large. In this case the consumption share is expected to vary only very little due to the increased influence of the speculation motive; see again Figure 4. Hence, risk premia are much lower due to the low influence of expected changes in s_t^B.

This influence of x_t on the distribution of the aggregate persistence is absent in the representative agent model with a time-varying persistence. Put differently, $E_t(\rho_{x,t+1}|x_t)$ is independent of x_t. The shock $\eta_{x,t+1}$ to x_{t+1} negatively affects next-period persistence, $\rho_{x,t+1}$. However, the level of the long-run risk state, x_t, affects neither same-period persistence, $\rho_{x,t}$, nor next-period persistence, $\rho_{x,t+1}$. Hence, Figure 8 shows only very little variation in risk premia with regard to x_t. This difference between the two models has a large influence on the second moments of asset prices. Table 6 reports the annualized moments of the expected risk premia. It shows that the annualized standard deviation in the representative-agent model with a time-varying persistence is significantly lower compared to the heterogeneous-agent economy. Hence, the representative-agent model cannot explain the large variation in expected risk premia. We find similar differences for the annualized volatility of the price-dividend ratio: the volatility of the heterogeneous-agent model is 0.38 while its only 0.20 in the representative-agent economy with the time-varying persistence.

The results show that the presented representative-agent models with time-varying persistence cannot match the empirical predictions generated by the trading in the heterogeneous-agent economy. Possibly, adding more and more exogenous features to the model could help to improve the model. While this might eventually lead to a better explanation of the financial-market data, we believe that a simple model with small belief differences—as proposed in this paper—is a more compelling argument to explain asset-price variation. In the following section, we provide economic arguments in favor of the heterogeneous-belief model.

Table 6: Expected Risk Premia in the Model with Time-Varying Persistence

	Data	Het. Agents	TV Pers.
Mean	5.00	5.42	8.80
Std. dev.	4.60	5.73	2.49

The table shows the annualized mean and volatility of the expected risk premium. The first column shows the empirical values reported in Table 1 of Martin (2017). The second column shows the results for the heterogeneous-agent setup with $\rho_x^A = \rho_x = 0.99$ and $\rho_x^B = 0.96$ for the correct beliefs. Column three shows the results for the representative-agent model with time-varying persistence. All returns are shown in percentages, so a value of 1.5 is a 1.5% annualized figure.

4.5.2 Advantages of the Heterogeneous-Belief Model

Detecting the persistent component x_t in the consumption data is already rather challenging since we cannot observe this variable directly. Therefore, additionally identifying a time-varying persistence (another unobservable) will be very difficult, particularly with the limited amount of data we have. Adding additional exogenous variation requires additional parameters that must be matched to fit the data. It also carries with it that all agents in the model agree on the values of these difficult-to-estimate parameters. In Section 3.1 we provided evidence on the plausibility that agents could naturally disagree over such a parameter. Any such disagreement automatically enriches the model dynamics endogenously.

In general, a sufficiently complex representative-agent model could capture the same effects as the heterogeneous-belief model. (Ultimately, we could use the social planner's objective as the utility of a representative agent.) Thus a logical criterion for comparing models is their parsimony. Our model is a very parsimonious extension of the simple Case I model of Bansal and Yaron (2004). We introduce exactly one extra parameter (the persistence parameter of the second agent), and that by itself is sufficient to generate several effects visible in the data, such as excess volatility, and increased return predictability. Belief heterogeneity gives rich asset-pricing dynamics endogenously. To get excess volatility in Case II of Bansal and Yaron (2004) requires adding an additional stochastic volatility shock, three additional parameters, and assuming that the effects of these shocks are extremely persistent. While the representative-agent framework has long been the workhorse of asset pricing, it leads to a paradigm where rich dynamics are built into the model through postulating additional exogenous shocks. This derives from the fact that we shut off the most natural channel in economics, which is trade among agents. In a setting where market heterogeneity is no longer relevant, trade again becomes an important channel, one that works endogenously.

Skepticism in the field about the persistence parameter in long-run risk is easy to come by—many finance researchers object to how close to a unit root the long-run risk process must be. With our new calibration we show that paradoxically this can be a strength of the model. With a moderate amount of disagreement, agents trading on their skepticism naturally create empirically plausible outcomes which are absent in the model with a representative agent.

5 Sensitivity to Agents' Beliefs about Persistence

In the following we conduct several robustness checks, which highlight the role of belief differences about the persistence of long-run risks on the risk sharing between the agents.

5.1 Long-Run Simulations

In the discussion of our main model in Section 4, we emphasized that none of the agents were driven out of the market in the finite sample simulations each containing 77 years of data. We now show that this property was a result of the carefully chosen difference between the agents' persistence parameters, $\rho_x^A = 0.99$ and $\rho_x^B = 0.96$. For this purpose, we now examine simulations of the two alternative model specifications for longer time periods.

Figure 9 displays the consumption share of agent B in a model with $\rho_x^A = 0.985$ and $\rho_x^B = 0.975$ over time for different initial shares $s_0^B = \{0.01, 0.05, 0.5\}$. We report the median, 5%, and 95% quantile paths using 1,000 samples each consisting of 500 years of simulated data. To minimize the influence of the initial value of x_t, we initialize each simulated path by running a "burn-in" period of 1,000 years before using the output. The left panel shows the results for $\rho_x = \rho_x^A = 0.985$ (agent A has the correct beliefs) and the right panel those for $\rho_x = \rho_x^B = 0.975$ (agent B has the correct beliefs).

The figure shows that in all cases the consumption share of agent B strongly increases over time. This can be rationalized by the findings from Section 3.2, which considers the same calibration for the discretized model, and shows that for Epstein–Zin preferences the investor with the high belief about ρ_x is deliberately giving away future wealth to insure against low growth states; see Figures 2 and 3. The increase in the share of agent B occurs faster if agent B has the correct beliefs (Figure 9, right panel), but the increase is almost as strong if agent A has the correct beliefs (left panel) due to the small influence of the speculation motive for small belief differences. Hence, given a small difference in the beliefs and independent of whether agent A or agent B has the correct beliefs, in the long run the agent who believes in a smaller ρ_x will dominate the economy. So the influence of the risk-sharing motive dominates the influence

Figure 9: Consumption Shares—Simulations

(a) $\rho_x = 0.985, s_0^B = 0.5$

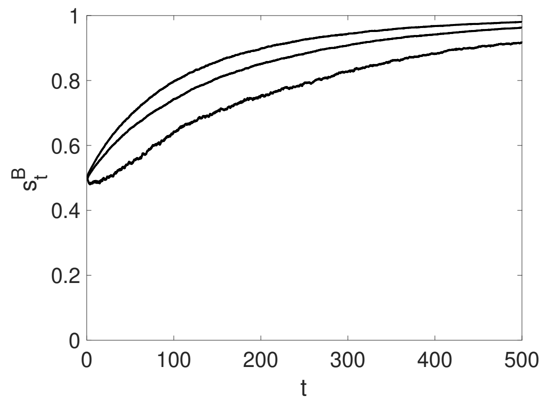

(b) $\rho_x = 0.975, s_0^B = 0.5$

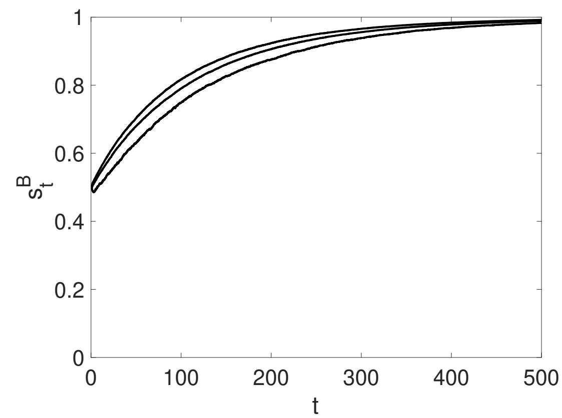

(c) $\rho_x = 0.985, s_0^B = 0.05$

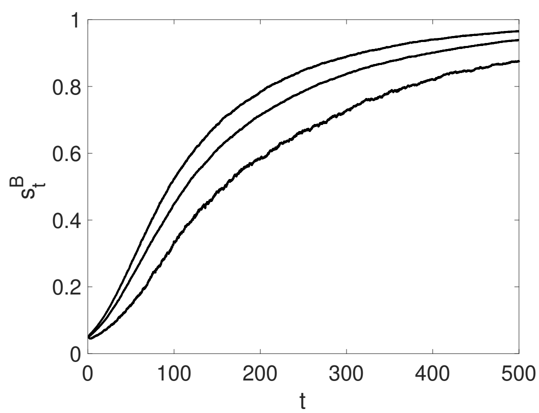

(d) $\rho_x = 0.975, s_0^B = 0.05$

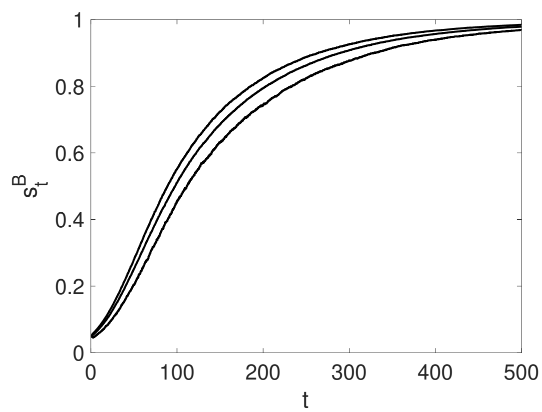

(e) $\rho_x = 0.985, s_0^B = 0.01$

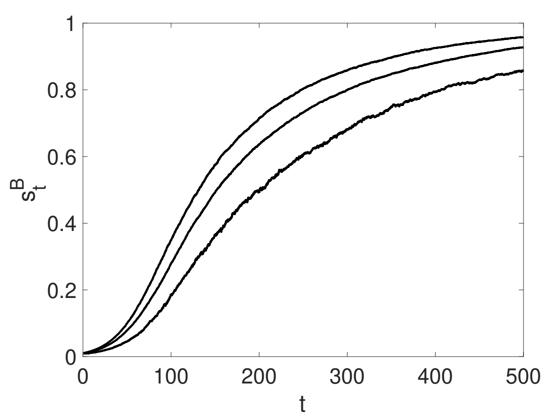

(f) $\rho_x = 0.975, s_0^B = 0.01$

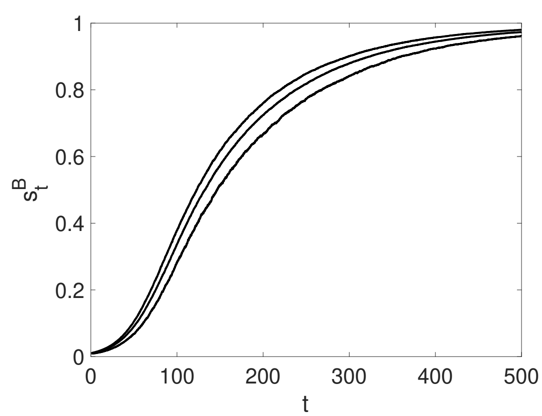

The figure shows the median, 5%, and 95% quantile paths of the consumption share of agent B for 1,000 samples each consisting of 500 years of simulated data. Agent B believes that $\rho_x^B = 0.975$ and agent A believes that $\rho_x^A = 0.985$. Results are shown for different initial consumption shares ($s_0^B = \{0.01, 0.05, 0.5\}$). The left panel depicts the case where agent B has the wrong beliefs about the long-run risk process ($\rho_x = 0.985 = \rho_x^A$) and in the right panel agent B has the right beliefs ($\rho_x = 0.95 = \rho_x^B$).

of the speculation motive. If the economy is initially almost entirely populated by agent A, $s_0^B = 0.01$, her consumption share decreases sharply, leading the agent to lose significant share in a short amount of time. Again, this is in line with Figure 2 which shows that risk sharing is especially strong for small s_t^B. So for the calibration with small belief differences about the persistence, the risk-sharing motive will quickly shift wealth to the investors who believe in a lower persistence, independent of the true data generating process.

Figure 10 shows the corresponding results for $\rho_x^B = 0.95$ and an initial allocation of $s_0^B = 0.01$. The right panel shows the results for $\rho_x = 0.95$ (agent B has the correct beliefs). In this case, both the speculation and risk-sharing motives will transfer wealth on average to agent B and hence we observe a strong increase in his consumption share. The left panel shows the

Figure 10: Consumption Shares for $\rho_x^B = 0.95$—Simulations

(a) $\rho_x = 0.985$ (b) $\rho_x = 0.95$

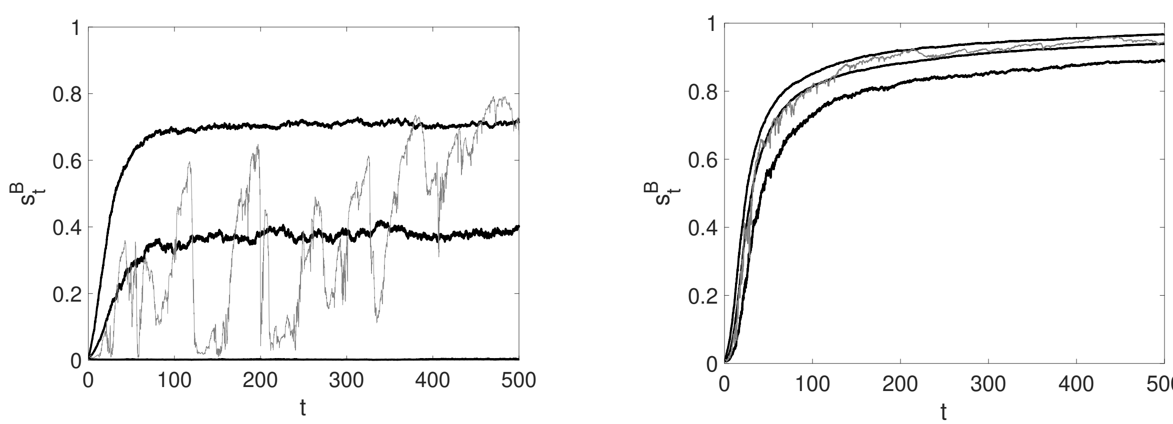

The figure shows the median, 5%, and 95% quantile paths of the consumption share of agent B for 1,000 samples—each consisting of 500 years of simulated data—as well as a sample path (grey line). Agent B believes that $\rho_x^B = 0.95$ and agent A believes that $\rho_x^A = 0.985$. Results are shown for an initial consumption share of $s_0^B = 0.01$. The left panel depicts the case where agent B has the wrong beliefs about the long-run risk process ($\rho_x = 0.985 = \rho_x^A$) and in the right panel agent B has the right beliefs ($\rho_x = 0.95 = \rho_x^B$).

results for $\rho_x = 0.985$ (agent A has the correct beliefs). We observe that the median share increases initially but then remains at about 40 percent. For small s_t^B the influence of the risk-sharing motive is large as there is a large demand for insurance against negative shocks to x_{t+1} from the first group of investors; recall Figure 2. Therefore, the consumption share of agent B increases on average. When the consumption share of agent B becomes larger, the supply (demand) of the insurance against long-run risks increases (decreases) and hence the influence of the risk-sharing motive becomes smaller; see also Figure 4. Furthermore, the influence of the speculation motive is large in this case of large belief differences. As agent

B has the wrong beliefs, this motive on average transfers wealth to investor A. We observe that the two motives are roughly equally strong for a consumption share of investor B of about 40 percent. Furthermore, the 5% and 95% quantile paths and the grey sample path show that there is significant variation in the consumption shares. So the interaction between the speculation and risk-sharing motives induces large changes in the wealth distribution over time.

In sum, the long-run simulations of the various specifications of the heterogeneous-agent economy give us some critical insights. If the belief difference between the two agents is sufficiently small, then agent B—who believes in a smaller persistence—dominates the economy rather quickly. Similarly, a heterogeneous-agent model with a sizable belief difference and in which agent B—the agent with the smaller persistence parameter—has the correct belief also quickly reduces to a model that is essentially a representative-agent model. On the contrary, if agent A—the agent with the higher persistence parameter—has the correct belief then both investors have a significant consumption share in the long run and the risk sharing of the investors will have a large impact on asset prices.

5.2 Robustness of the Results

We complete our analysis with several other robustness checks. In Figure 11a we show the median consumption share of agent B (as in Figure 9) for different degrees of risk aversion $\gamma^h = \{2, 5, 10\}$. The risk-sharing motive should increase with the degree of risk aversion, while the speculation motive decreases (see equation (24)). Hence, the smaller γ is, the less wealth should be transferred to type B investors, who hold the wrong beliefs. And indeed, this is exactly what we observe in Figure 11a. For $\gamma^h = 10$ (yellow line) the influence of the risk-sharing motive is strong. Hence, agent B profits from selling the insurance against long-run risks and rapidly accumulates wealth. For $\gamma^h = 5$ (red line) this effect becomes less severe and agent B's consumption share increases less quickly. For $\gamma^h = 2$ (blue line) the risk-sharing motive has little influence as investor A is less averse to long-run risks and the speculation motive dominates equilibrium outcomes. As $\rho_x = \rho_x^A$, the speculation motive works in favor of agent A (agent B bets on states that have a smaller probability under the true probability measure) and agent A dominates the economy in the long run. If agent B has the correct beliefs, $\rho_x = \rho_x^B$, the speculation motive works in favor of agent B. We show this case in Figure 11b. The blue line shows the consumption shares for $\rho_x = \rho_x^A$ and the red line those for $\rho_x = \rho_x^B$. So, in the absence of the risk-sharing motive, the speculation motive determines equilibrium outcomes.

In Figure 11c we depict the robustness of our findings with regard to the persistence levels of x_t. We show the consumption paths for $\rho_x^B = 0.6$ and $\rho_x^A = 0.5$ instead of 0.975 and 0.985, respectively. Lowering the persistence will—similarly to the decrease in risk aversion—decrease the risk-sharing motive. Risk premia in the economy are only large for ρ_x close to 1 but collapse for smaller ρ_x; see Bansal and Yaron (2004). Hence, even those investors who believe in a larger (but significantly smaller than 1) value for ρ_x have only small hedging demands. Consequently, we observe that in this setup the dynamics of the consumption shares strongly depend on the true value of ρ_x as the speculation motive dominates—that is, the agent with the correct beliefs will dominate the economy.

Figure 11: The Risk-Sharing and Speculation Motives

(a) $\rho_x^A = \rho_x = 0.985$, $\rho_x^B = 0.975$

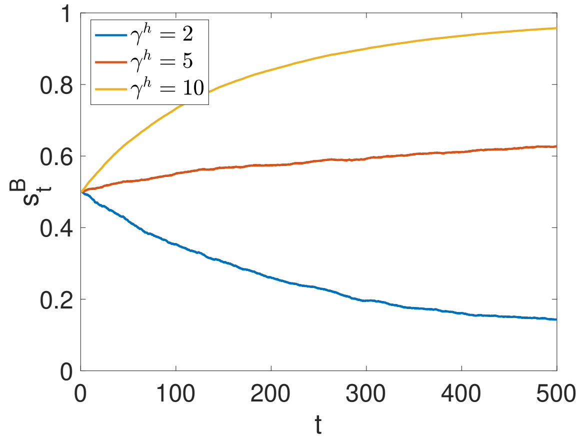

(b) $\gamma^h = 2$, $\rho_x^A = 0.985$, $\rho_x^B = 0.975$

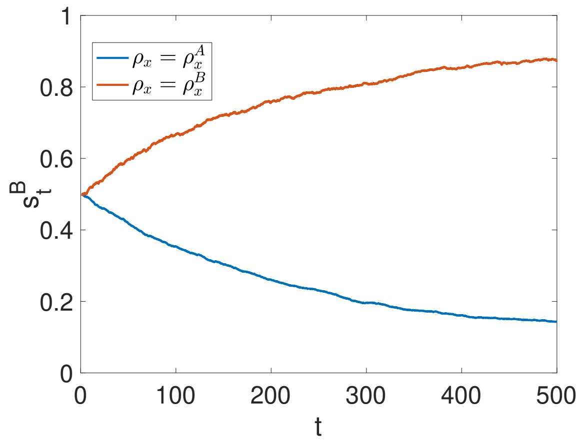

(c) $\gamma^h = 10$, $\rho_x^B = 0.6$, $\rho_x^A = 0.5$

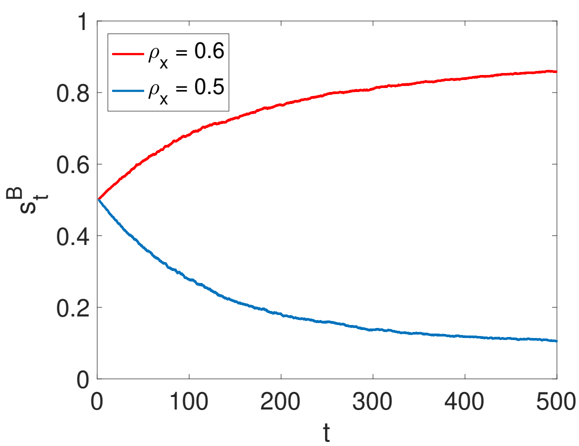

The figure shows the median consumption share of agent B for 1,000 samples each consisting of 500 years of simulated data. Panel (a) shows the time series for different degrees of risk aversion $\gamma^h \in \{2, 5, 10\}$. Agent B believes that $\rho_x^B = 0.975$ and agent A has the correct beliefs with $\rho_x^A = \rho_x = 0.985$. Panel (b) shows the time series for $\gamma^h = 2$, $\rho_x^A = 0.985$, and $\rho_x^B = 0.975$ for the two cases in which either agent A (blue line) or agent B (red line) has the correct beliefs. Panel (c) shows the time series for $\gamma^h = 10$, $\rho_x^A = 0.6$, and $\rho_x^B = 0.5$ for the two cases where either agent A (blue line) or agent B (red line) has the correct beliefs.

6 Conclusion

In this paper, we show that belief differences can solve several asset-pricing puzzles. In particular, we demonstrate that even small belief differences have large effects on prices in models with Epstein–Zin preferences and long-run consumption risks. Solving models with heterogeneous agents and recursive preferences is not a simple task as equilibrium allocations are no longer a function of the exogenous state alone. Therefore, we first derive a recursive formulation of the first-order conditions for equilibrium and present a numerical solution method to solve for the equilibrium. This methodology can be applied to solve a broad class of models featuring multiple agents with recursive utility and continuous or discrete state processes.

We then apply the methodology and analyze the influence of belief heterogeneity in long-run risk asset-pricing models. For this purpose, we take a standard long-run risk model with persistent changes in the mean growth rate of consumption as in Bansal and Yaron (2004, Case I), with the only exception that we use two different persistence levels: one agent believes in a slightly smaller amount of persistence relative to the original paper, while one agent believes in a slightly larger amount (and is correct). This model not only generates a large and significant equity premium, it also addresses many of the empirical deficiencies of the representative-agent model. Notably, it adds significant countercyclical time variation in expected risk premia to the model, consistent with the data reported in Martin (2017). Furthermore, shifts in the wealth distribution increase the volatility of the price–dividend ratio to levels close to the data as the impact of the different agents on asset prices varies over time. The variation in the wealth distribution also helps to address the predictability puzzle pointed out by Beeler and Campbell (2012) as well as the countercyclical variation in return predictability reported by Henkel, Martin, and Nardari (2011) and Dangl and Halling (2012). The endogenous variation in asset prices increases the predictability of returns especially in recessions while simultaneously decreasing the predictability of consumption and dividend growth. Therefore, belief heterogeneity in long-run risk models can explain many empirical patterns on financial markets that seem puzzling from the viewpoint of the representative-agent model. While our analysis in this paper focuses on investor heterogeneity in the long-run risk model, we expect the key channel—namely the strong influence of belief differences on asset prices—to carry over to other asset-pricing models with recursive preferences.

References

Alchian, A. A. (1950): "Uncertainty, Evolution, and Economic Theory," *Journal of Political Economy*, 58, 211–221.

Anderson, E. W. (2005): "The dynamics of risk-sensitive allocations," *Journal of Economic theory*, 125, 93–150.

Anderson, E. W., E. Ghysels, and J. L. Juergens (2005): "Do Heterogeneous Beliefs Matter for Asset Pricing?" *The Review of Financial Studies*, 18, 875–924.

Andrei, D., B. Carlin, and M. Hasler (2019a): "Asset Pricing with Disagreement and Uncertainty about the Length of Business Cycles," *Management Science*, 65, 2900–2923.

Andrei, D., M. Hasler, and A. Jeanneret (2019b): "Asset Pricing with Persistence Risk," *Review of Financial Studies*, 32, 2809–2849.

Bansal, R., D. Kiku, and A. Yaron (2012): "An Empirical Evaluation of the Long-Run Risks Model for Asset Prices," *Critical Finance Review*, 1, 183–221.

——— (2016): "Risks For the Long Run: Estimation with Time Aggregation," *Journal of Monetary Economics*, 82, 52–69.

Bansal, R. and A. Yaron (2004): "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 59, 1481–1509.

Beeler, J. and J. Y. Campbell (2012): "The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment," *Critical Finance Review*, 1, 141–182.

Bhamra, H. S. and R. Uppal (2014): "Asset Prices with Heterogeneity in Preferences and Beliefs," *Review of Financial Studies*, 27, 519–580.

Bidder, R. and I. Dew-Becker (2016): "Long-run risk is the worst-case scenario," *American Economic Review*, 106, 2494–2527.

Billingsley, P. (1999): *Convergence of Probability Measures*, Wiley Series in Probability and Statistics: Probability and Statistics, New York: John Wiley & Sons Inc., second ed.

Blume, L. and D. Easley (2006): "If You're so Smart, why Aren't You Rich? Belief Selection in Complete and Incomplete Markets," *Econometrica*, 74, 929–966.

Borovička, J. (2019): "Survival and Long-Run Dynamics with Heterogeneous Beliefs Under Recursive Preferences," *Journal of Political Economy*, Forthcoming.

Branger, N., I. Dumitrescu, V. Ivanova, and C. Schlag (2011): "Preference Heterogeneity and Survival in Long-Run Risk Models," Working Paper, Goethe University.

Branger, N., P. Konermann, and C. Schlag (2019): "Optimists and Pessimists in (In)Complete Markets," *Journal of Financial and Quantitative Analysis*, Forthcoming.

Buraschi, A., F. Trojani, and A. Vedolin (2014): "Economic Uncertainty, Disagreement, and Credit Markets," *Management Science*, 60, 1281–1296.

Campbell, H. and J. Cochrane (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107, 205–251.

Carlin, B. I., F. A. Longstaff, and K. Matoba (2014): "Disagreement and asset prices," *Journal of Financial Economics*, 114, 226–238.

Chen, H., S. Joslin, and N.-K. Tran (2012): "Rare Disasters and Risk Sharing with Heterogeneous Beliefs," *Review of Financial Studies*, 25, 2189–2224.

Colacito, R. and M. M. Croce (2013): "International Asset Pricing with Recursive Preferences," *The Journal of Finance*, 68, 2651–2686.

Colacito, R., M. M. Croce, S. W. Ho, and P. Howard (2018a): "BKK the EZ Way: International Long-Run Growth News and Capital Flows," *American Economic Review*, Forthcoming.

Colacito, R., M. M. Croce, Y. Liu, and I. Shaliastovich (2018b): "Volatility Risk Pass-Through," Working paper.

Colacito, R., M. M. Croce, and Z. Liu (2018c): "Recursive Allocations and Wealth Distribution with Multiple Goods: Existence, Survivorship, and Dynamics," *Quantitative Economics*, Forthcoming.

Collin-Dufresne, P., M. Johannes, and L. A. Lochstoer (2015): "A Robust Numerical Method for Solving Risk-Sharing Problems with Recursive Preferences," Working Paper.

——— (2016a): "Asset Pricing when 'This Time is Different'," *Review of Financial Studies*, 30, 505–535.

——— (2016b): "Parameter Learning in General Equilibrium: The Asset Pricing Implications," *American Economic Review*, 106, 664–698.

Cvitanić, J., E. Jouini, S. Malamud, and C. Napp (2012): "Financial Markets Equilibrium with Heterogeneous Agents," *Review of Finance*, 16, 285–321.

Dangl, T. and M. Halling (2012): "Predictive regressions with time-varying coefficients," *Journal of Financial Economics*, 106, 157–181.

David, A. (2008): "Heterogeneous Beliefs, Speculation, and the Equity Premium," *Journal of Finance*, 63, 41–83.

Dindo, P. (2019): "Survival in speculative markets," *Journal of Economic Theory*, 181, 1–43.

Duffie, D., P.-Y. Geoffard, and C. Skiadas (1994): "Efficient and equilibrium allocations with stochastic differential utility," *Journal of Mathematical Economics*, 23, 133–146.

Dumas, B., R. Uppal, and T. Wang (2000): "Efficient Intertemporal Allocations with Recursive Utility," *Journal of Economic Theory*, 93, 240–259.

Epstein, L. G., E. Farhi, and T. Strzalecki (2014): "How Much Would You Pay to Resolve Long-Run Risk?" *American Economic Review*, 104, 2680–2697.

Epstein, L. G. and S. E. Zin (1989): "Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937–969.

Flodén, M. (2008): "A Note on the Accuracy of Markov-Chain Approximations to Highly Persistent AR(1) Processes," *Economics Letters*, 99, 516–520.

Friedman, M. (1953): *Essays in Positive Economics*, Chicago: University of Chicago Press.

Gârleanu, N. and S. Panageas (2015): "Young, Old, Conservative, and Bold: The Implications of Heterogeneity and Finite Lives for Asset Pricing," *Journal of Political Economy*, 123, 670–685.

Grammig, J. and E. M. Küchlin (2018): "A two-step indirect inference approach to estimate the long-run risk asset pricing model," *Journal of Econometrics*, 205, 6–33.

Henkel, S. J., J. S. Martin, and F. Nardari (2011): "Time-varying short-horizon predictability," *Journal of Financial Economics*, 99, 560–580.

James, G. M. and A. A. Smith (1998): "Approximate Bias Correction in Econometrics," *Journal of Econometrics*, 85, 205–230.

Ju, N. and J. Miao (2012): "Ambiguity, Learning, and Asset Returns," *Econometrica*, 80, 559–591.

Judd, K. L., F. Kubler, and K. Schmedders (2003): "Asset Trading Volume with Dynamically Complete Markets and Heterogeneous Agents," *Journal of Finance*, 58, 2203–2217.

Kan, R. (1995): "Structure of Pareto optima when agents have stochastic recursive preferences," *Journal of Economic Theory*, 66, 626–631.

Kendall, M. G. (1954): "Note on Bias in the Estimation of Autocorrelation," *Biometrika*, 41, 403–404.

Lucas, R. E. and N. L. Stokey (1984): "Optimal growth with many consumers," *Journal of Economic Theory*, 32, 139–171.

Martin, I. (2017): "What is the Expected Return on the Market?" *The Quarterly Journal of Economics*, 132, 367–433.

Martin, I. and C. Wagner (2019): "What is the Expected Return on a Stock?" *Journal of Finance*, 74, 1887–1929.

Mas-Colell, A. and W. R. Zame (1991): "Equilibrium Theory in Infinite Dimensional Spaces," in *Handbook of Mathematical Economics*, ed. by W. Hildenbrand and H. Sonnenschein, Elsevier, vol. IV, 1835–1898.

Patton, A. J. and A. Timmermann (2010): "Why do forecasters disagree? Lessons from the term structure of cross-sectional dispersion," *Journal of Monetary Economics*, 57, 803–820.

Pohl, W., K. Schmedders, and O. Wilms (2016): "Higher-Order Effects in Asset-Pricing Models with Long-Run Risks," Working paper, University of Zurich.

——— (2018): "Higher-Order Effects in Asset Pricing Models with Long-Run Risks," *Journal of Finance*, 73, 1061–1111.

Roche, H. (2011): "Asset prices in an exchange economy when agents have heterogeneous homothetic recursive preferences and no risk free bond is available," *Journal of Economic Dynamics and Control*, 35, 80–96.

Rouwenhorst, G. (1995): *Asset Pricing Implications of Equilibrium Business Cycle Models*, Cooley, T. (Ed.), Frontiers of Business Cycle Research. Princeton University Press, Princeton, NJ, 294–330.

SANDRONI, A. (2000): "Do Markets Favor Agents Able to Make Accurate Predictions?" *Econometrica*, 68, 1303–1341.

SCHORFHEIDE, F., D. SONG, AND A. YARON (2018): "Identifying Long-Run Risks: A Bayesian Mixed-Frequency Approach," *Econometrica*, 86, 617–654.

SHILLER, R. J. (1981): "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?" *American Economic Review*, 71, 421–436.

TAUCHEN, G. AND R. HUSSEY (1991): "Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models," *Econometrica*, 59, 317–396.

WACHTER, J. A. (2013): "Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?" *Journal of Finance*, 68, 987–1035.

WEIL, P. (1989): "The Equity Premium Puzzle and the Risk-Free Rate Puzzle," *Journal of Monetary Economics*, 24, 401–421.

YAN, H. (2008): "Natural Selection in Financial Markets: Does It Work?" *Management Science*, 54, 1935–1950.

For Online Publication

Appendix

A Additional Literature

The study of agent belief heterogeneity begins with the market selection hypothesis of Alchian (1950) and Friedman (1953). By analogy with natural selection, the market selection hypothesis states that agents with systematically wrong beliefs will eventually be driven out of the market. The influence of agent heterogeneity on market outcomes under the standard assumption of discounted expected utility is well understood, and consistent with market selection. Sandroni (2000) and Blume and Easley (2006) find strong support for this hypothesis under the assumption of time-separable preferences in an economy without growth. Yan (2008) and Cvitanić, Jouini, Malamud, and Napp (2012) analyze the survival of investors in a continuous-time framework where there are not only differences in beliefs but also potentially differences in the utility parameters of the investors. They show that it is always the investor with the lowest survival index[12] who survives in the long run. David (2008) considers a similar model setup, in which both agents have distorted estimates of the mean growth rate of the economy, showing that—as agents with lower risk aversion undertake more aggressive trading strategies—the equity premium increases the lower the risk aversion is. Chen, Joslin, and Tran (2012) analyze how differences in beliefs about the probability of disasters affect asset prices. They show that even if there is only a small fraction of investors who are optimistic about disasters, this fraction sells insurance for the disaster states and so eliminates most of the risk premium associated with disaster risk. Bhamra and Uppal (2014) consider the case of habit utility.

For recursive utility, this qualitative behavior changes fundamentally. However, there has been less research in this area as solving such models is anything but trivial. Lucas and Stokey (1984) observe in the deterministic case that the problem of finding all Pareto-optimal allocations can be made recursive if we allow the weights in the social welfare function to be time varying. This approach is extended by Kan (1995) to the stochastic case with finite state spaces. Anderson (2005) develops an extensive theory for the special case of risk-sensitive preferences, no growth, and finite state spaces, and finds first-order conditions similar to

[12] Yan (2008) shows that the survival index increases with belief distortion, risk aversion, and the subjective time discount rate of the investor.

those we use below. In particular, he shows how to characterize the equilibrium by a single value function instead of one value function for each agent. Collin-Dufresne, Johannes, and Lochstoer (2015) derive similar first-order conditions to ours for recursive utility by equating marginal utilities, but use a different procedure to solve for their allocations. Duffie, Geoffard, and Skiadas (1994) formulate the problem in continuous time, while Dumas, Uppal, and Wang (2000) reformulate it in terms of variational utility. Borovička (2019) uses this formulation to explore the question of the survival of agents with recursive utility in continuous time, and shows that agents with fundamentally wrong beliefs can survive or even dominate. So, inferences about market selection and equilibrium outcomes fundamentally differ under the assumption of general recursive utility compared to the special case of standard time-separable preferences. While Borovička (2019) concentrates on the special case of i.i.d. consumption growth, Branger, Dumitrescu, Ivanova, and Schlag (2011) generalize the results to a model with long-run risks as a state variable.

However, most papers with heterogeneous investors and recursive preferences consider only an i.i.d. process for consumption growth. For example, Gârleanu and Panageas (2015) analyze the influence of heterogeneity in the preference parameters on asset prices in a two-agent OLG economy. Roche (2011) considers a model in which the heterogeneous investors can only invest in a stock but there is no risk-free bond. Hence, as there is no savings trade-off, the impact of recursive preferences on equilibrium outcomes will be quite different.

Exceptions that relax the i.i.d. assumptions include the papers by Branger, Konermann, and Schlag (2019) and Collin-Dufresne, Johannes, and Lochstoer (2016a). Both papers re-examine the influence of belief differences regarding disaster risk with Epstein–Zin instead of CRRA preferences as in Chen, Joslin, and Tran (2012). Branger, Konermann, and Schlag (2019) provide evidence that the influence of investors with more optimistic beliefs about disasters is less profound when the disaster occurs to the growth rate of consumption and show that the risk-sharing mechanism persists even when markets are incomplete. Collin-Dufresne, Johannes, and Lochstoer (2016a) make a similar claim but for a different reason. They show that if the investors can learn about the probability of disasters and if they have recursive preferences, the impact of the optimistic investor on asset prices decreases. Optimists are uncertain about the probability of disasters and hence will provide less insurance to pessimistic investors.

A different strand of the literature, which does not rely on the i.i.d. assumption, is comprised of papers on international asset pricing. This area (advanced by Riccardo Colacito and Mariano Croce in particular) considers models with Epstein–Zin preferences, two investors, and also two goods for which investors have different preferences (home bias). For example, Colacito and Croce (2013) argue that a model with highly correlated international long-run components in output can explain both the low correlation between consumption differentials and the tendency of high interest rate currencies to appreciate. The authors furthermore show that an increase in capital mobility can explain the structural break in the data for the pre- and post-1970 period. Colacito, Croce, and Liu (2018c) provide the theoretical foundation for the multiple good economy by providing results on equilibrium existence and agents' survival; they also compare computational methods for solving the model. Furthermore, Colacito, Croce, Ho, and Howard (2018a) use a model with Epstein–Zin preferences and short- and long-run productivity shocks to study the effects of these shocks on international investment flows.

In a different direction, Epstein, Farhi, and Strzalecki (2014) argue that an Epstein–Zin investor dislikes long-run risk to the extent that he or she would pay a substantial premium to get rid of it. In a model with two agents, the agent who believes that risk is longer term than the other is willing to pay an insurance premium to the other agent to hedge against long-run risk.

B Proofs and Details

In this appendix, we provide proofs for the theoretical results presented in Section 2. Along the way, we derive a system of first-order conditions for Epstein–Zin preferences. This system constitutes the foundation for our numerical solution method (see Appendix C).

B.1 Proofs for Section 2.1

Proof of Theorem 1. Let $\boldsymbol{\lambda} = \{\bar{\lambda}^1, \ldots, \bar{\lambda}^H\}$ be a set of Negishi weights and let $\{\boldsymbol{C}\}_0 = \{\{C^1\}_0, \ldots, \{C^H\}_0\}$ denote a vector of agents' consumption processes. The optimal decision $\{\boldsymbol{C}\}_0^*$ of the social planner in the initial period assigns consumption streams to all individual agents for all periods and possible states. Obviously, the optimal decisions must satisfy the market-clearing condition (1) in all periods and states. For ease of notation we again abbreviate the state dependence; we use C_t^h for $C^h(y^t)$ and $U_{\{t\}}^h$ for $U^h\left(\{C^h\}_t\right)$.

To derive the first-order conditions, we borrow a technique from the calculus of variations.

For any function f_t, we can vary the consumption of two agents by

$$\begin{aligned} C_t^h \to & \quad C_t^h + \epsilon f_t \\ C_t^l \to & \quad C_t^l - \epsilon f_t. \end{aligned} \tag{30}$$

It is sufficient to consider the variation with $l = 1$ and $h \in \mathbb{H}^-$. For an optimal allocation it must be true that

$$\left. \frac{\mathrm{d}SP(\{\boldsymbol{C}\}_0; \boldsymbol{\lambda})}{\mathrm{d}\epsilon} \right|_{\epsilon=0} = 0. \tag{31}$$

This gives us

$$\bar{\lambda}^h \hat{U}_{0,t}^h = \bar{\lambda}^1 \hat{U}_{0,t}^1, \qquad h \in \mathbb{H}^-, \tag{32}$$

where $\hat{U}_{t,t+k}^h$ is defined as

$$\hat{U}_{t,t+k}^h = \left. \frac{\mathrm{d}U^h(C_t^h, \ldots, C_{t+k}^h + \epsilon f_{t+k}, \ldots)}{\mathrm{d}\epsilon} \right|_{\epsilon=0}. \tag{33}$$

Using the expression given in Equation (2), the derivative $\hat{U}_{t,t+k}^h$ satisfies a recursive equation with the initial condition

$$\hat{U}_{t,t}^h = \left. \frac{\mathrm{d}U^h(C_t^h + \epsilon f_t, \ldots)}{\mathrm{d}\epsilon} \right|_{\epsilon=0} = F_1^h \left(C_t^h, R_t[U_{\{t+1\}}^h] \right) \cdot f_t, \tag{34}$$

where $F_k^h \left(C_t^h, R_t^h[U_{\{t+1\}}^h] \right)$ denotes the derivative of $F^h \left(C_t^h, R_t^h[U_{\{t+1\}}^h] \right)$ with respect to its kth argument. The recursive step is given by

$$\begin{aligned} \hat{U}_{t,t+k}^h &= \left. \frac{\mathrm{d}F^h \left(C_t^h, R_t^h \left[U^h(C_{t+1}^h, \ldots C_{t+k}^h + \epsilon f_{t+k}, \ldots) \right] \right)}{\mathrm{d}\epsilon} \right|_{\epsilon=0} \\ &= F_2^h \left(C_t^h, R_t^h[U_{\{t+1\}}^h] \right) \cdot \left. \frac{\mathrm{d}R_t^h \left[U^h(\cdot) \right]}{\mathrm{d}\epsilon} \right|_{\epsilon=0} \\ &= F_2^h \left(C_t^h, R_t^h[U_{\{t+1\}}^h] \right) \cdot \frac{\mathrm{d}G_h^{-1} \left(E_t^h G_h \left[U^h(\cdot) \right] \right)}{\mathrm{d}E_t^h G_h[U^h(\cdot)]} \cdot \left. \frac{\mathrm{d}E_t^h G_h[U^h(\cdot)]}{\mathrm{d}\epsilon} \right|_{\epsilon=0} \\ &= F_2^h \left(C_t^h, R_t^h[U_{\{t+1\}}^h] \right) \cdot \frac{1}{G_h'(G_h^{-1}(E_t^h G_h[U_{\{t+1\}}^h]))} \cdot E_t^h \left(G_h'(U_{\{t+1\}}^h) \cdot \hat{U}_{t+1,t+k}^h \right) \\ &= F_2^h \left(C_t^h, R_t^h[U_{\{t+1\}}^h] \right) \cdot \frac{E_t^h \left(G_h'(U_{\{t+1\}}^h) \cdot \hat{U}_{t+1,t+k}^h \right)}{G_h'(R_t^h[U_{\{t+1\}}^h])}, \end{aligned} \tag{35}$$

where we use $\frac{\partial G^{-1}(x)}{\partial x} = \frac{1}{G'(G^{-1}(x))}$ and abbreviate $U^h(C_{t+1}^h, \ldots C_{t+k}^h + \epsilon f_{t+k}, \ldots)$ by $U^h(\cdot)$. We can recast this recursion into a useful form. For this purpose, we define a second recursion

$U^h_{t,t+k}$ by

$$U^h_{t,t} = F^h_1\left(C^h_t, R^h_t[U^h_{\{t+1\}}]\right) \tag{36}$$

and

$$U^h_{t,t+k} = \Pi^h_{t+1} \cdot U^h_{t+1,t+k}, \tag{37}$$

where

$$\Pi^h_{t+1} = F^h_2\left(C^h_t, R^h_t[U^h_{\{t+1\}}]\right) \cdot \frac{G'_h(U^h_{\{t+1\}})}{G'_h(R^h_t[U^h_{\{t+1\}}])} \frac{\mathrm{dP}^h_{t,t+1}}{\mathrm{dP}_{t,t+1}}. \tag{38}$$

A simple induction shows that

$$\hat{U}^h_{t,t+k} = E_t(U^h_{t,t+k} f_t). \tag{39}$$

Plugging (39) into the optimality condition (32) we obtain

$$E_0\left((\bar{\lambda}^h U^h_{0,t} - \bar{\lambda}^1 U^1_{0,t}) f_t\right) = 0, \qquad h \in \mathbb{H}^-. \tag{40}$$

Under a broad range of regularity conditions, this condition implies that

$$\bar{\lambda}^h U^h_{0,t} = \bar{\lambda}^1 U^1_{0,t}, \qquad h \in \mathbb{H}^-. \tag{41}$$

For example, if $\bar{\lambda}^h U^h_{0,t} - \bar{\lambda}^1 U^1_{0,t}$ has finite variance, then this holds for the Riesz Representation Theorem for L^2 random variables. We can then split Expression (41) into two parts. First define $\lambda^h_0 \equiv \bar{\lambda}^h$ to obtain

$$\frac{\lambda^h_0}{\lambda^1_0} = \frac{U^1_{0,t}}{U^h_{0,t}} = \frac{\Pi^1_1}{\Pi^h_1}\frac{U^1_{1,t}}{U^h_{1,t}} = \frac{\Pi^1_1}{\Pi^h_1}\frac{\lambda^h_1}{\lambda^1_1}, \qquad h \in \mathbb{H}^-,$$

where λ^h_1 denotes the Negishi weight in the social planner's optimal solution in $t = 1$. Generalizing this equation for any period t, we obtain the following dynamics for the optimal weight[13] λ^h_{t+1}:

$$\frac{\lambda^h_{t+1}}{\lambda^1_{t+1}} = \frac{\Pi^h_{t+1}}{\Pi^1_{t+1}}\frac{\lambda^h_t}{\lambda^1_t}, \qquad h \in \mathbb{H}^-. \tag{42}$$

Inserting the initial condition (36) into (41) for $t = 0$ and generalizing it for any social planner's optimal solution at time t yields

$$\lambda^h_t F^h_1\left(C^h_t, R^h_t[U^h_{\{t+1\}}]\right) = \lambda^1_t F^1_1\left(C^1_t, R^1_t[U^1_{\{t+1\}}]\right), \qquad h \in \mathbb{H}^-. \tag{43}$$

[13] Note that we can either solve the model in terms of the ratio $\frac{\lambda^h_t}{\lambda^1_t}$ (this is equal to setting $\lambda^1_t = 1$ for all t as done in Judd, Kubler, and Schmedders (2003)) or we can normalize the weights so that they remain bounded in $(0,1)$. Our solution method uses the latter approach as it obtains better numerical properties.

Equation (43) states the optimality conditions for the individual consumption choices at any time t. This completes the proof of Theorem 1. $\square$

Note that for time-separable utility, $F_1^h\left(C_t^h, R_t^h[U_{\{t+1\}}^h]\right)$ is simply the marginal utility of agent h at time t, and so we obtain the same optimality condition as, for example, Judd, Kubler, and Schmedders (2003) (see Equation (7) on page 2209). In this special case the Negishi weights can be pinned down in the initial period and thereafter remain constant. For general recursive preferences this is not true. The optimal weights vary over time following the law of motion described by Equation (42).

We can use Equations (42) and (43) together with the market-clearing condition (1) to compute the social planner's optimal solution. We therefore define $\boldsymbol{\lambda}_t^- = \{\lambda_t^B, \lambda_t^3, \ldots, \lambda_t^H\}$ and let V^h denote the value function of agent $h \in \mathbb{H}$. We are looking for model solutions of the form $V^h(\boldsymbol{\lambda_t^-}, y^t)$. So the model solution depends on both the exogenous state y^t and the time-varying Negishi weights $\boldsymbol{\lambda}_t^-$. An optimal allocation is then characterized by the following four equations:

- the market-clearing condition (1)

$$\sum_{h=1}^{H} C^h(\boldsymbol{\lambda_t^-}, y^t) = C(y^t); \tag{44}$$

- the value functions (2) of the individual agents

$$V^h(\boldsymbol{\lambda_t^-}, y^t) = F^h\left(C^h(\boldsymbol{\lambda_t^-}, y^t), R_t^h[V^h(\boldsymbol{\lambda_{t+1}^-}, y^{t+1})]\right), \qquad h \in \mathbb{H}; \tag{45}$$

- the optimality conditions (43) for the individual consumption decisions for $h \in \mathbb{H}^-$

$$\lambda_t^h F_1^h\left(C^h(\boldsymbol{\lambda_t^-}, y^t), R_t^h[V^h(\boldsymbol{\lambda_{t+1}^-}, y^{t+1})]\right) = \lambda_t^1 F_1^1\left(C^1(\boldsymbol{\lambda_t^-}, y^t), R_t^1[V^1(\boldsymbol{\lambda_{t+1}^-}, y^{t+1})]\right); \tag{46}$$

- the equations (42) for the dynamics of $\boldsymbol{\lambda_t^-}$

$$\frac{\lambda_{t+1}^h}{\lambda_{t+1}^1} = \frac{\Pi_{t+1}^h}{\Pi_{t+1}^1}\frac{\lambda_t^h}{\lambda_t^1}, \qquad h \in \mathbb{H}^-, \tag{47}$$

with

$$\Pi_{t+1}^h = F_2^h\left(C^h(\boldsymbol{\lambda_t^-}, y^t), R_t^h[V^h(\boldsymbol{\lambda_{t+1}^-}, y^{t+1})]\right) \cdot \frac{G_h'(V^h(\boldsymbol{\lambda_{t+1}^-}, y^{t+1}))}{G_h'(R_t^h[V^h(\boldsymbol{\lambda_{t+1}^-}, y^{t+1})])}\frac{\mathrm{dP}_{t,t+1}^h}{\mathrm{dP}_{t,t+1}}. \tag{48}$$

This concludes the general description of the equilibrium obtained from the social planner's optimization problem.

To prove Theorem 2, we first derive a variant of Lemma 1 in Blume and Easley (2006).

Lemma 1. *Let X_t^i, $i = 1, 2, \ldots, H$, be a family of positive random variables for each $t = 0, 1, 2, \ldots$, such that $A \leq \sum_i X_t^i \leq B$ with $B \in \mathbb{R}_{++}$. Let $f^i : \mathbb{R}_{++} \to \mathbb{R}_{++}$, $i = 1, 2, \ldots, H$, be a family of decreasing functions such that $f^i(x) \to \infty$ as $x \to 0$. If $f^i(X_t^i)/f^j(X_t^j) \to \infty$, then $X_t^i \to 0$ for $t \to \infty$. If $X_t^i \to 0$, then for at least one j, $\limsup_t f^i(X_t^i)/f^j(X_t^j) = \infty$.*

Proof. Since X_t^i is positive, $X_t^i \leq B$ for all i, t. By assumption, $0 < f^j(B) \leq f^j(X_t^j)$. Thus, $f^i(X_t^i)/f^j(X_t^j) \to \infty$ if and only if $f^i(X_t^i) \to \infty$, which happens when $X_t^i \to 0$ as $t \to \infty$.

Conversely, assume $X_t^i \to 0$. Every period, for at least one j, $X_t^j \geq A/H$ (otherwise $\sum_{i=1}^H X_t^i < A$). Since there are only finitely many random variables, for at least one j we have $X_t^j \geq A/H$ infinitely often. Then, by assumption, $f^j(X_t^j) \leq f^j(A/H)$ infinitely often, and so $\limsup f^i(X_t^i)/f^j(X_t^j) = \infty$. □

Proof of Theorem 2. By the first-order condition (5), $\lambda_t^j/\lambda_t^i = F_1^i(C_t^i, R_t^i)/F_1^j(C_t^j, R_t^j)$. Since F^h is additively separable, F_1^h is a function of consumption alone. Let $f^i = F_1^i, f^j = F_1^j$, $A = \underline{C}$, and $B = \overline{C}$, and apply Lemma 1. □

B.2 Proofs for Section 2.2

In this section we provide the specific expressions for V^h, F_1^h, F_2^h, and Π^h when the heterogeneous investors have recursive preferences as in Epstein and Zin (1989) and Weil (1989). The value function for Epstein–Zin (EZ) preferences is given by[14]

$$V_t^h = \left[(1-\delta^h)(C_t^h)^{\rho^h} + \delta^h R_t^h \left(V_{t+1}^h\right)^{\rho^h}\right]^{\frac{1}{\rho^h}} \tag{49}$$

with

$$\begin{aligned} R_t^h \left(V_{t+1}^h\right) &= G_h^{-1}\left(E_t^h\left[G_h(V_{t+1}^h)\right]\right) \\ G_h(V_{t+1}^h) &= \left(V_{t+1}^h\right)^{\alpha^h}. \end{aligned}$$

Recall that the parameter δ^h is the discount factor, $\rho^h = 1 - \frac{1}{\psi^h}$ determines the EIS, ψ^h, and $\alpha^h = 1 - \gamma^h$ determines the relative risk aversion γ^h of agent h. The derivatives of

[14] For ease of notation, we again abbreviate the dependence on the exogenous state y_t and the endogenous state $\underline{\boldsymbol{\lambda}}_{\boldsymbol{t}}^-$. Hence we write V_t^h for $V^h(\boldsymbol{\lambda}_{\boldsymbol{t}}^-, y_t)$ or C_t^h for $C^h(\boldsymbol{\lambda}_{\boldsymbol{t}}^-, y_t)$.

$F^h\left(C_t^h, R_t^h[V_{t+1}^h]\right) = V_t^h$ with respect to its first and second argument are then given by

$$F_{1,t}^h = (1-\delta^h)(C_t^h)^{\rho^h-1}(V_t^h)^{1-\rho^h} \tag{50}$$

and

$$F_{2,t}^h = \delta^h R_t^h \left(V_{t+1}^h\right)^{\rho^h-1}(V_t^h)^{1-\rho^h}. \tag{51}$$

In this paper we focus on growth economies. Therefore, we introduce the following normalization to obtain a stationary formulation of the model. We define the consumption share of agent h by $s_t^h = \frac{C_t^h}{C_t}$ and the normalized value functions, $v_t^h = \frac{V_t^h}{C_t}$. Recall that $\Delta c_{t+1} = c_{t+1} - c_t$ with $c_t = \log(C_t)$. The value function (49) is then given by

$$v_t^h = \left[(1-\delta^h)(s_t^h)^{\rho^h} + \delta^h R_t^h \left(v_{t+1}^h e^{\Delta c_{t+1}}\right)^{\rho^h}\right]^{\frac{1}{\rho^h}}. \tag{52}$$

By inserting (50) into (46) we obtain the optimality condition for the individual consumption decisions

$$\lambda_t^h F_1^h\left(C^h(\boldsymbol{\lambda}_t^-, y^t), R_t^h[V^h(\boldsymbol{\lambda}_{t+1}^-, y^{t+1})]\right) = \lambda_t^1 F_1^1\left(C^1(\boldsymbol{\lambda}_t^-, y^t), R_t^1[V^1(\boldsymbol{\lambda}_{t+1}^-, y^{t+1})]\right),$$

which simplifies to

$$\lambda_t^h(1-\delta^h)(C_t^h)^{\rho^h-1}(V_t^h)^{1-\rho^h} = \lambda_t^1(1-\delta^1)(C_t^1)^{\rho^1-1}(V_t^1)^{1-\rho^1}. \tag{53}$$

Recall the definition of the normalized Negishi weights, $\underline{\lambda}_t^h = \frac{\lambda_t^h}{(v_t^h)^{\rho^h-1}}$. From Equation (53) we obtain

$$\underline{\lambda}_t^h(1-\delta^h)(s_t^h)^{\rho^h-1} = \underline{\lambda}_t^1(1-\delta^1)(s_t^1)^{\rho^1-1}. \tag{54}$$

This equation is the optimality condition for the individual consumption decisions we employ for solving for the model with Epstein–Zin preferences. Inserting the de-trended weight $\underline{\lambda}_t^h$ into the dynamics for the weights (47), we obtain

$$\frac{\lambda_{t+1}^h}{\lambda_{t+1}^1} = \frac{\underline{\lambda}_{t+1}^h(v_{t+1}^h)^{\rho^h-1}}{\underline{\lambda}_{t+1}^1(v_{t+1}^1)^{\rho^1-1}} = \frac{\underline{\lambda}_t^h(v_t^h)^{\rho^h-1}}{\underline{\lambda}_t^1(v_t^1)^{\rho^1-1}}\frac{\Pi_{t+1}^h}{\Pi_{t+1}^1}, \qquad h \in \mathbb{H}^-. \tag{55}$$

Plugging the expressions for Epstein–Zin preferences (49)–(51) into Equation (48), we obtain the following expression for Π^h_{t+1}:

$$\begin{aligned}\Pi^h_{t+1} &= \delta^h R^h_t \left(V^h_{t+1}\right)^{\rho^h-1} (V^h_t)^{1-\rho^h} \frac{\left(V^h_{t+1}\right)^{\alpha^h-1}}{R^h_t \left(V^h_{t+1}\right)^{\alpha^h-1}} \frac{\mathrm{dP}^h_{t,t+1}}{\mathrm{dP}_{t,t+1}} \\ &= \delta^h (V^h_t)^{1-\rho^h} \frac{\left(V^h_{t+1}\right)^{\alpha^h-1}}{R^h_t \left(V^h_{t+1}\right)^{\alpha^h-\rho^h}} \frac{\mathrm{dP}^h_{t,t+1}}{\mathrm{dP}_{t,t+1}}. \end{aligned} \tag{56}$$

Using the normalized value function $v^h_t = \frac{V^h_t}{C_t}$, we have

$$\Pi^h_{t+1} = \delta^h (v^h_t)^{1-\rho^h} \frac{\left(v^h_{t+1} e^{\Delta c_{t+1}}\right)^{\alpha^h-1}}{R^h_t \left(v^h_{t+1} e^{\Delta c_{t+1}}\right)^{\alpha^h-\rho^h}} \frac{\mathrm{dP}^h_{t,t+1}}{\mathrm{dP}_{t,t+1}}. \tag{57}$$

Equation (55) can then be written as

$$\frac{\underline{\lambda}^h_{t+1}}{\underline{\lambda}^1_{t+1}} = \frac{\underline{\lambda}^h_t}{\underline{\lambda}^1_t} \frac{\underline{\Pi}^h_{t+1}}{\underline{\Pi}^1_{t+1}}, \qquad h \in \mathbb{H}^-, \tag{58}$$

where

$$\underline{\Pi}^h_{t+1} = \underbrace{\delta^h e^{\rho^h \Delta c_{t+1}} \frac{\mathrm{dP}^h_{t,t+1}}{\mathrm{dP}_{t,t+1}}}_{\text{CRRA-Term}} \underbrace{\frac{\left(v^h_{t+1} e^{\Delta c_{t+1}}\right)^{\alpha^h-\rho^h}}{R^h_t \left(v^h_{t+1} e^{\Delta c_{t+1}}\right)^{\alpha^h-\rho^h}}}_{\text{New EZ-Term}}. \tag{59}$$

For $\alpha^h = \rho^h$, we obtain the standard term for CRRA preferences; the dynamics of $\underline{\lambda}^h_{t+1}$ only depend on the subjective discount factor, the EIS, and the subjective beliefs of the investors. For Epstein–Zin preferences, we obtain an extra term that reflects the time trade-off. Using the normalization $\sum_{h=1}^H \underline{\lambda}^h_t = 1$, the dynamics for $\underline{\lambda}^h_{t+1}$ are then given by

$$\underline{\lambda}^h_{t+1} = \frac{\underline{\lambda}^h_t \underline{\Pi}^h_{t+1}}{\sum_{h=1}^H \underline{\lambda}^h_t \underline{\Pi}^h_{t+1}}. \tag{60}$$

Hence, for Epstein–Zin preferences we obtain the following system for the first-order conditions (44)–(48):

The market-clearing condition:

$$\sum_{h=1}^H s^h_t = 1. \tag{MC}$$

The optimality condition for the individual consumption decisions:

$$\underline{\lambda}_t^h(1-\delta^h)(s_t^h)^{\rho^h-1} = \underline{\lambda}_t^1(1-\delta^1)(s_t^1)^{\rho^1-1}, \quad h \in \mathbb{H}^-, \tag{CD}$$

with $\sum_{h=1}^H \underline{\lambda}_t^h = 1$.

The value functions of the individual agents:

$$v_t^h = \left[(1-\delta^h)(s_t^h)^{\rho^h} + \delta^h R_t^h \left(v_{t+1}^h e^{\Delta c_{t+1}}\right)^{\rho^h}\right]^{\frac{1}{\rho^h}}, \quad h \in \mathbb{H}. \tag{VF}$$

The equation for the dynamics of $\underline{\lambda}_t^h$:

$$\begin{aligned} \underline{\lambda}_{t+1}^h &= \frac{\underline{\lambda}_t^h \underline{\Pi}_{t+1}^h}{\sum_{h=1}^H \underline{\lambda}_t^h \underline{\Pi}_{t+1}^h} \\ \underline{\Pi}_{t+1}^h &= \delta^h e^{\rho^h \Delta c_{t+1}} \frac{\mathrm{dP}_{t,t+1}^h}{\mathrm{dP}_{t,t+1}} \frac{\left(v_{t+1}^h e^{\Delta c_{t+1}}\right)^{\alpha^h-\rho^h}}{R_t^h \left(v_{t+1}^h e^{\Delta c_{t+1}}\right)^{\alpha^h-\rho^h}}, \quad h \in \mathbb{H}^-. \end{aligned} \tag{Dλ}$$

Note that the conditions (MC, CD, VF, Dλ) are just the equilibrium conditions (11)–(15) stated in Section 2.2. We observe that Equation (CD) and hence the individual consumption decisions s_t^h only depend on time t information and that there is no intertemporal dependence. This feature allows us to first solve for s_t^h given the current state of the economy, and in a second step to solve for the dynamics of the Negishi weights. Hence, we can separate solving the optimality conditions (11)–(15) into two steps in order to reduce the computational complexity. In Appendix C we describe this approach in detail.

Using condition (CD) we can prove Theorem 3. Recall that $\rho^h = 1 - \frac{1}{\psi^h} < 1$ for all possible values of an agent's EIS, $\psi^h > 0$.

Proof of Theorem 3. Condition (CD) implies

$$\frac{\underline{\lambda}_t^j}{\underline{\lambda}_t^i} = \frac{(1-\delta^i)(s_t^i)^{\rho^i-1}}{(1-\delta^j)(s_t^j)^{\rho^j-1}}.$$

Now let $f^i(s) = s^{\rho^i-1}$, $f^j(s) = s^{\rho^j-1}$, and $A = B = 1$, and apply Lemma 1. $\square$

C Solution Method

We describe our solution method for asset-pricing models with heterogeneous agents and recursive preferences.

C.1 Computational Procedure—A Two-Step Approach

For ease of notation the following procedures are described for $H = 2$ agents and a single state variable $y_t \in \mathbb{R}^1$. However, the approach can analogously be extended to the general case of $H > 2$ agents and multiple states. We solve the social planner's problem using a collocation projection. For this we perform the usual transformation from an equilibrium described by the infinite sequences (with a time index t) to the equilibrium described by functions of some state variable(s) x on a state space X. We denote the current exogenous state of the economy by y and the subsequent state in the next period by y' with the state space $Y \in \mathbb{R}$. The term $\underline{\lambda}_2$ denotes the current endogenous state of the Negishi weight and $\underline{\lambda}_2'$ denotes the corresponding state in the subsequent period with $\underline{\Lambda}_2 \in (0,1)$.

We approximate the value functions of the two agents, $v^h(\underline{\lambda}_2, y), h = \{1,2\}$, by two-dimensional cubic splines and we denote the approximated value functions by $\hat{v}^h(\underline{\lambda}_2, y)$. For the collocation projection we have to choose a set of collocation nodes $\{\underline{\lambda}_{2_k}\}_{k=1}^n$ and $\{y_l\}_{l=1}^m$ at which we evaluate $\hat{v}^h(\underline{\lambda}_2, y)$. The individual consumption shares only depend on the endogenous state $\underline{\lambda}_{2_k}$. So in the following we show how to first solve for the individual consumption shares at the collocation nodes $s_k^h = s^h(\underline{\lambda}_{2_k})$, which are then used to solve for the value functions v^h and the dynamics of the endogenous state $\underline{\lambda}_2$.

Step 1: *Computing Optimal Consumption Allocations*

Equation (13) has to hold at each collocation node $\{\underline{\lambda}_{2_k}\}_{k=1}^n$:

$$\underline{\lambda}_{2_k}(1-\delta^2)\left(s_k^2\right)^{\rho^2-1} = (1-\underline{\lambda}_{2_k})(1-\delta^1)\left(s_k^1\right)^{\rho^1-1}.$$

Together with the market-clearing condition (11) we get

$$\underline{\lambda}_{2_k}(1-\delta^2)\left(s_k^2\right)^{\rho^2-1} = (1-\underline{\lambda}_{2_k})(1-\delta^1)\left(1-s_k^2\right)^{\rho^1-1}. \tag{61}$$

So for each node $\{\underline{\lambda}_{2_k}\}_{1=0}^n$ the optimal consumption choice s_k^2 can be computed by solving Equation (61) and s_k^1 is obtained by the market-clearing condition (11).[15] For the special case of $\rho^2 = \rho^1$ we can solve for s^2 as a function of $\underline{\lambda}_2$ analytically, and hence we do not have to solve the system of equations for each node.

[15] Note that in the case of H agents we have to solve a system of $H-1$ equations that pin down the $H-1$ individual consumption choices $s^h \in \mathbb{H}^-$.

Step 2: *Solving for the Value Function and the Dynamics of the Negishi Weights*

Solving for the value function is not as straightforward as it depends on the dynamics of the endogenous state $\underline{\lambda}_2$, which are unknown and follow Equation (15). We compute the expectation over the exogenous state by a Gaussian quadrature with Q quadrature nodes. This implies that the values for y' at which we evaluate v^h are given by the quadrature rule. We denote the corresponding quadrature nodes by $\{y'_{l,g}\}_{l=1,g=1}^{m,Q}$ and the weights by $\{\omega_g\}_{g=1}^{Q}$.[16] We can then solve Equation (15) for a given pair of collocation nodes $\{\underline{\lambda}_{2_k}, y_l\}_{k=1,l=1}^{n,m}$ and the corresponding quadrature nodes $\{y'_{l,g}\}_{l=1,g=1}^{m,Q}$ to compute a vector $\vec{\underline{\lambda}}'_2$ of size $n \times m \times Q$ that consists of the corresponding values $\underline{\lambda}'_{2_{k,l,g}}$ for each node. For each $\underline{\lambda}'_{2_{k,l,g}}$, Equation (15) then reads

$$\begin{aligned} \underline{\lambda}'_{2_{k,l,g}} &= \frac{\underline{\lambda}_{2_k}\underline{\Pi}^2}{(1-\underline{\lambda}_{2_k})\underline{\Pi}^1 + \underline{\lambda}_{2_k}\underline{\Pi}^2} \\ \underline{\Pi}^h &= \delta^h e^{\rho^h \Delta c(y'_{l,g})} \left(\frac{v^h(\underline{\lambda}'_{2_{k,l,g}}, y'_{l,g}) e^{\Delta c(y'_{l,g})}}{R^h\left[v^h(\underline{\lambda}'_2, y')e^{\Delta c(y')} | \underline{\lambda}_{2_k}, y_l\right]} \right)^{\alpha^h - \rho^h} \frac{\mathrm{dP}^h(y'_{l,g}|y_l)}{\mathrm{dP}(y'_{l,g}|y_l)}, \end{aligned} \tag{62}$$

where

$$R^h\left[v^h(\underline{\lambda}'_2, y')e^{\Delta c(y')}|\underline{\lambda}_{2_k}, y_l\right] = G_h^{-1}\left(E\left[G_h\left(v^h(\underline{\lambda}'_2, y')e^{\Delta c(y')}\right)\frac{\mathrm{dP}^h(y')}{\mathrm{dP}(y')}\middle|\underline{\lambda}_{2_k}, y_l\right]\right).$$

Note that $\underline{\lambda}'_{2_{k,l,g}}$ depends on the full distribution of $\underline{\lambda}'_2$ through the expectation operator. By applying the Gaussian quadrature to compute the expectation we get

$$E\left[G_h\left(v^h(\underline{\lambda}'_2, y')e^{\Delta c(y')}\right)\frac{\mathrm{dP}^h(y')}{\mathrm{dP}(y')}\middle|\underline{\lambda}_{2_k}, y_l\right] \approx \sum_{g=1}^{Q} G_h\left(v^h(\underline{\lambda}'_{2_{k,l,g}}, y_{l,g})e^{\Delta c(y_{l,g})}\right)\cdot \omega_g.$$

By computing the expectation with the quadrature rule, we do not need the full distribution of $\underline{\lambda}'_2$; instead, we only have to evaluate v^h at those values $\underline{\lambda}'_{2_{k,l,g}}$ that can be obtained by solving (62) for each pair of collocation nodes $\{\underline{\lambda}_{2_k}, y_l\}_{k=1,l=1}^{n,m}$ and the corresponding quadrature nodes $\{y'_{l,g}\}_{l=1,g=1}^{m,Q}$. So at the end we have a square system of equations with $n \times m \times Q$ unknowns, $\underline{\lambda}'_{2_{k,l,g}}$, and as many equations (62) for each $\{k, l, g\}$.

The value function is in general not known so we have to compute it simultaneously when

[16] Note that the quadrature nodes $\{y'_{l,g}\}_{l=1,g=1}^{m,Q}$ depend on the state today, $\{y_l\}_{l=1}^{m}$.

solving for $\underline{\lambda}'_{2_{k,l,g}}$. Plugging the approximation $\hat{v}^h(\underline{\lambda}_2, y)$ into the value function (12) yields

$$\hat{v}^h(\underline{\lambda}_{2_k}, y_l) = \left[(1-\delta^h) \left(s_k^h\right)^{\rho^h} + \delta^h R^h \left(\hat{v}^h(\underline{\lambda}'_2, y') e^{\Delta c(y')} \middle| \underline{\lambda}_{2_k}, y_l \right) \right]^{\frac{1}{\rho^h}}. \tag{63}$$

The collocation projection conditions require that the equation has to hold at each collocation node $\{\underline{\lambda}_{2_k}, y_l\}_{k=1,l=1}^{n,m}$. So we obtain a square system of equations with $n \times m \times 2$ equations (63) and as many unknowns for the spline interpolation at each collocation node, which we solve simultaneously with the system for $\underline{\lambda}'_{2_{k,l,g}}$ described above.

For all results presented in the paper, we choose an approximation interval for x_t that covers ± 4 (unconditional) standard deviations around the unconditional mean of the process. For $\underline{\lambda}_t^2$ the minimum and maximum values are given by 0 and 1 so the full state space is included. We approximate the value functions using two-dimensional cubic splines with not-a-knot end conditions. We provide the solver with additional information that we can formally derive for the limiting cases. For example, we know that for $\underline{\lambda}_t^2 = 1$ ($\underline{\lambda}_t^2 = 0$) agent 2 (1) consumes everything, so this corresponds to the representative-agent economy populated only by agent 2 (1). Hence, we require that the value function for these cases equals the value function for the corresponding representative-agent economy. We also know that for $\underline{\lambda}_t^2 = 0$ ($\underline{\lambda}_t^2 = 1$) the consumption of agent 2 (1) is 0, and hence the value function is also 0. As the shocks in the model are normally distributed, we compute the expectations over the exogenous states by Gauss–Hermite quadrature using five nodes for the shock in x_{t+1} and three for the shock in Δc_{t+1}.

C.2 Accuracy of the Solution Method

In the following we provide details of the accuracy of the solution method. We report numerical errors in the fixed-point equation (12), which determines the value functions of the two agents. In addition, we report the errors in the equilibrium conditions (15) for the Negishi weights. For the models with Epstein–Zin preferences, we compute these errors on a 200×200 uniform grid for the two states of the model—the exogenous long-run risk state, x_t, and the endogenous Negishi weight, $\underline{\lambda}_t^h$.

As a benchmark for our analysis, we also report numerical errors for a model with CRRA preferences. In the CRRA case, the dynamics of $\underline{\lambda}_t^B$ are exogenous and given in closed form; see equation (15), which shows that for $\rho^h = \alpha^h$, $\underline{\lambda}_{t+1}^B$ does not depend on the value functions. Thus, we only have to approximate the value functions in the fixed-point equation (12) for the exogenous process x_t specified in equation (16) and the exogenous weights $\underline{\lambda}_t^B$ given by

(15) using the probability ratio specified in equation (17). This benchmark case gives us a first indication of the adequateness of the projection approach to precisely approximating the value functions of the agents.

Table 7 reports the root-mean-square error (RMSE) as well as the maximum absolute error (MAE) for different numbers of collocation nodes using a uniform grid with n nodes for the $\underline{\lambda}_t^B$ dimension and m nodes for the x_t dimension. The first panel reports these errors for a CRRA model with $\psi^h = \frac{1}{\gamma^h} = 1.5$ and persistence parameter values $\rho_x^A = 0.985$ and $\rho_x^B = 0.975$. We observe that for $n = 16$ and $m = 8$, errors are already small with a maximum error in the value function of agent A of 5.2e-4. $n = 20$ the MAE can be decreased to 2.8e-5 and for $n = 50$ all errors are smaller then 1.0e-5. So, we observe a high accuracy of the projection approach for the approximation of the agents' value functions.

The second panel of Table 7 reports errors for the same calibration but with Epstein–Zin preferences ($\gamma^h = 10$, $\psi^h = 1.5$) as used in Sections 3 and 5. Now the Negishi weights $\underline{\lambda}_{t+1}^B$ depend on the value functions and, therefore, we must solve for the value functions in equation (12) jointly with the dynamics for the Negishi weights; see equation (15). For the small belief difference, the numerical errors are already small for $n = 16$ and $m = 12$ nodes with a maximum absolute error of 3.1e-5 in the equation for the Negishi weights. For $n = 50$ and $m = 22$ the maximum error in the Negishi weights reduces to 8.6e-6 and the root-mean-square errors are below 5e-7.

The bottom panel of Table 7 reports errors for the main economy with $\rho_x^A = 0.99$ and $\rho_x^B = 0.96$, which we discussed extensively in Section 4. We observe that the numerical errors are large for the approximation with $n = 16$ and $m = 12$ with an RMSE of 0.0016 for both v_t^B and $\underline{\lambda}_t^B$ and maximum absolute errors of 0.0229 and 0.0450, respectively. Figure 12 plots the numerical errors in the value functions. We observe that the errors are especially large for $\underline{\lambda}_t^B$ close to 0. Therefore, we adjust the uniform collocation grid in order to obtain higher accuracy close to the boundary. In particular, we choose the grid such that half the nodes are uniformly distributed between 0 and 0.1 and the other half are uniformly distributed between 0.1 and 1. Figure 13 shows that the errors in v_t^B close to the boundary can be reduced by using this adjusted grid. In Table 7 the results for the adjusted grid are marked with an asterisk (*). We observe that using the same number of nodes, the maximum error in v_t^B can almost be halved to 0.0127 and the maximum error in $\underline{\lambda}_t^B$ can be reduced by more than a factor of 3 to 0.0135. Increasing the number of collocation points reduces the errors further. For $n = 50$ and $m = 22$ the largest RMSEs is 2.3e-4 for v_t^B with a corresponding maximum error of 0.0100. Figure 14 shows that the errors are still large, especially close to the boundary of $\underline{\lambda}_t^B = 0$. By increasing

Table 7: Numerical Errors

		$\boldsymbol{v_t^A}$		$\boldsymbol{v_t^B}$		$\boldsymbol{\underline{\lambda}_t^B}$	
n	m	RMSE	MAE	RMSE	MAE	RMSE	MAE
$\rho_x^A = 0.985, \rho_x^B = 0.975$, CRRA							
16	8	2.2e-5	5.2e-4	9.9e-6	1.7e-4	—	—
20	8	2.9e-4	2.8e-4	5.2e-6	9.0e-5	—	—
50	8	1.9e-6	1.0e-5	1.8e-6	7.7e-6	—	—
$\rho_x^A = 0.985, \rho_x^B = 0.975$, EZ							
16	12	1.4e-7	1.3e-6	7.8e-7	6.5e-6	2.7e-6	3.1e-5
50	22	1.1e-8	1.2e-7	7.7e-8	1.2e-6	4.3e-7	8.6e-6
$\rho_x^A = 0.99, \rho_x^B = 0.96$, EZ							
16	12	4.7e-5	6.6e-4	0.0016	0.0229	0.0016	0.0450
16*	12	9.1e-5	6.3e-4	4.5e-4	0.0127	7.7e-4	0.0135
50*	22	1.1e-6	2.4e-5	2.3e-4	0.0100	3.2e-5	0.0032
80*	22	1.9e-7	4.3e-6	6.5e-6	2.0e-4	4.7e-6	2.6e-4

The table shows root-mean-square errors (RMSEs) as well as maximum absolute errors (MAEs) in the value functions (12) as well as the equilibrium conditions for the Negishi weights (15) for different model calibrations and numbers of collocation nodes. Errors are reported for a state grid of ± 4 standard deviations around the unconditional mean of x_t. For $\underline{\lambda}_t^B$ the full grid between 0 and 1 is used. We denote the number of collocation nodes for the $\underline{\lambda}_t^B$-dimension by n and the number of collocation nodes for the x_t-dimension by m. An asterisk (*) denotes cases where we used an adjusted grid—instead of a uniform grid—to account for the nonlinearity close to the boundary at $\underline{\lambda}_t^B = 0$.

the number of nodes in the $\underline{\lambda}_t^B$ dimension further to 80, the errors can be reduced significantly with root-mean-square errors between 1.9e-7 and 6.5e-6 and maximum errors in the order of 1e-4. Note that errors are computed on a very large grid covering ± 4 unconditional standard deviations of the x_t-process. To verify that the errors are not only numerically small but also small in economic terms, we conduct two exercises. First we ask the hypothetical question by how much one needs to change the consumption share of the respective agent in order that the errors are exactly zero. For the high degree approximation with $n = 80$ we find maximum values for agent 1 of 0.0023 and agent 2 of 0.0038. This implies that the consumption shares need to be adjusted by a maximum of only 0.0038 in order to exactly satisfy the equilibrium condition. As a second test for economic significance, we analyze the influence of the errors on asset prices. We find that increasing the approximation degree does not change the model outcomes. This is demonstrated in Table 8 which shows the annualized asset-pricing moments as reported in Table 2 for the different approximation degrees. While the moments change slightly for the degree 16 compared to the degree 50 approximation, there is hardly a difference between the degree 50 and degree 80 solutions. Hence, the small errors near the boundary of the approximation space do not affect the model outcomes and the quantitative conclusions drawn in the paper.

Table 8: Annualized Asset-Pricing Moments

n	m	$E\left(p_t - d_t\right)$	$\sigma\left(p_t - d_t\right)$	$AC1\left(p_t - d_t\right)$	$E\left(R_t^m - R_t^f\right)$	$E\left(R_t^f\right)$	$\sigma\left(R_t^m\right)$	$\sigma\left(R_t^f\right)$
16	12	3.5264	0.3823	0.8007	5.3815	2.2911	21.712	1.9147
16*	12	3.5304	0.3866	0.8005	5.4512	2.2765	21.896	1.9342
50*	22	3.5304	0.3845	0.8002	5.4203	2.2856	21.806	1.9245
80*	22	3.5305	0.3846	0.8002	5.4204	2.2856	21.807	1.9246

The table shows selected annualized moments as in Table 2 for the two-agent economy for different approximation degrees. Agent A has the correct beliefs with $\rho_x^A = \rho_x = 0.99$; agent B has the belief $\rho_x^B = 0.96$. n denotes the number of collocation nodes for the $\underline{\lambda}_t^B$-dimension and m denotes the number of collocation nodes for the x_t-dimension. An asterisk (*) denotes cases where we used an adjusted grid—instead of a uniform grid—to account for the nonlinearity close to the boundary at $\underline{\lambda}_t^B = 0$.

D Closed-Form Solutions for the CRRA Case

In this section we derive closed-form solutions for the long-run risk model (16) with two investors and CRRA utility. Assume that the two agents A and B have CRRA utility with

Figure 12: Numerical Errors for $\rho_x^A = 0.99, \rho_x^B = 0.96, n = 16$, and $m = 12$ with the Uniform Grid

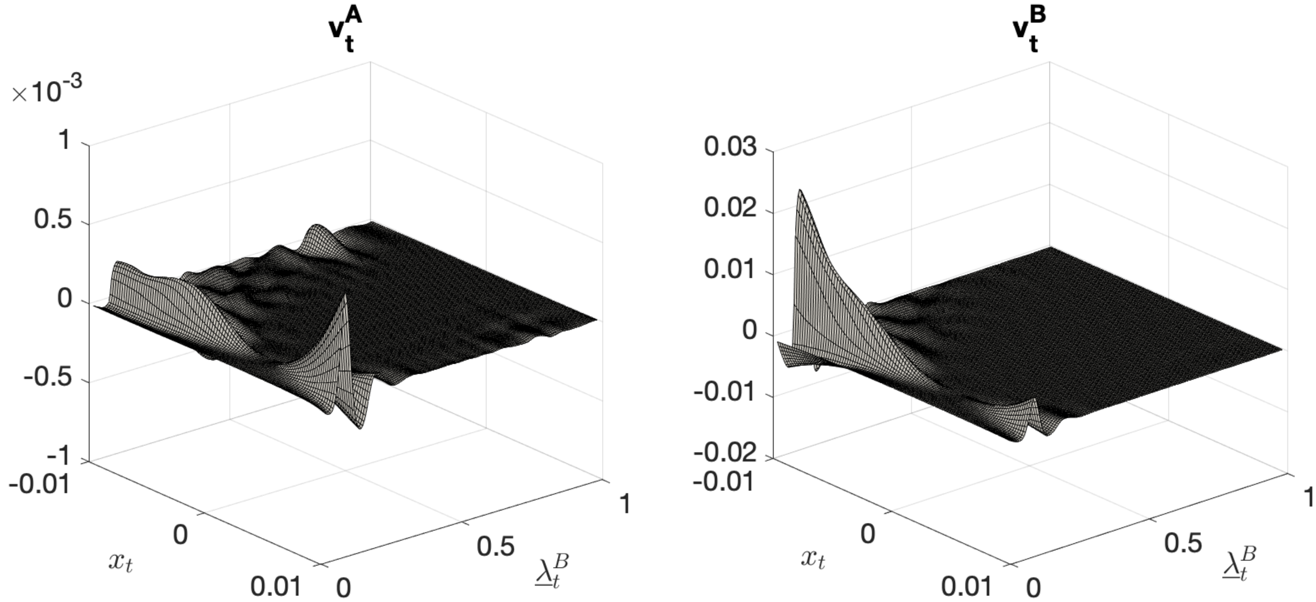

The figure plots numerical errors in the value functions (12) for $\rho_x^A = 0.99$ and $\rho_x^B = 0.96$. A state grid of ± 4 standard deviations around the unconditional mean of the x_t-process is used. For $\underline{\lambda}_t^B$ the full grid between 0 and 1 is used. The projection method uses $n = 16$ collocation nodes for the $\underline{\lambda}_t^B$-dimension and $m = 12$ nodes for the x_t-dimension.

the same degree of risk aversion. The equilibrium conditions (58) and (54) then simplify to

$$\left(\frac{s_{t+1}^A}{s_{t+1}^B}\right)^\gamma = \left(\frac{s_t^A}{s_t^B}\right)^\gamma \frac{\mathrm{dP}_{t,t+1}^A}{\mathrm{dP}_{t,t+1}^B}. \tag{64}$$

Taking logs and using the market-clearing condition yields

$$\log\left(\frac{s_{t+1}^A}{1 - s_{t+1}^A}\right) = \log\left(\frac{s_t^A}{1 - s_t^A}\right) + \frac{1}{\gamma}\log\left(\frac{\mathrm{dP}_{t,t+1}^A}{\mathrm{dP}_{t,t+1}^B}\right). \tag{65}$$

We log-linearize $\log(1 - s_{t+1}^A)$ around $\log(s_{t+1}^A) = \log(s_t^A)$. This step gives us

$$\begin{aligned}
\log\left(1 - e^{\log(s_{t+1}^A)}\right) &\approx \log\left(1 - s_t^A\right) - \frac{s_t^A}{1 - s_t^A}\left(\log\left(s_{t+1}^A\right) - \log\left(s_t^A\right)\right) \\
&= \log\left(1 - s_t^A\right) + \frac{s_t^A}{1 - s_t^A}\log\left(s_t^A\right) - \frac{s_t^A}{1 - s_t^A}a - \frac{s_t^A}{1 - s_t^A}b\eta_{x,t+1}. \quad (66)
\end{aligned}$$

Figure 13: Numerical Errors for $\rho_x^A = 0.99, \rho_x^B = 0.96, n = 16$, and $m = 12$ with the Adjusted Grid

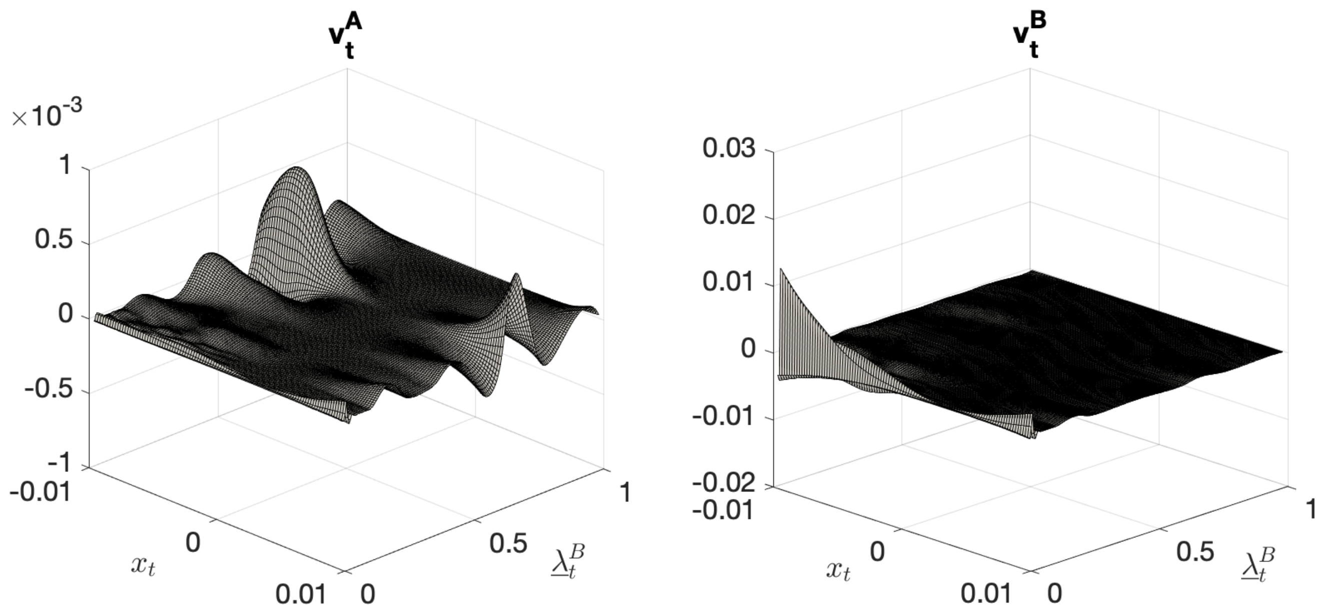

The figure plots numerical errors in the value functions (12) for $\rho_x^A = 0.99$ and $\rho_x^B = 0.96$. A state grid of ± 4 standard deviations around the unconditional mean of the x_t-process is used. For $\underline{\lambda}_t^B$ the full grid between 0 and 1 is used. The projection method uses $n = 16$ collocation nodes for the $\underline{\lambda}_t^B$-dimension and $m = 12$ nodes for the x_t-dimension. Results are shown for the adjusted grid to account for the nonlinearity close to $\underline{\lambda}_t^B = 0$.

Figure 14: Numerical Errors for $\rho_x^A = 0.99, \rho_x^B = 0.96, n = 50$, and $m = 22$

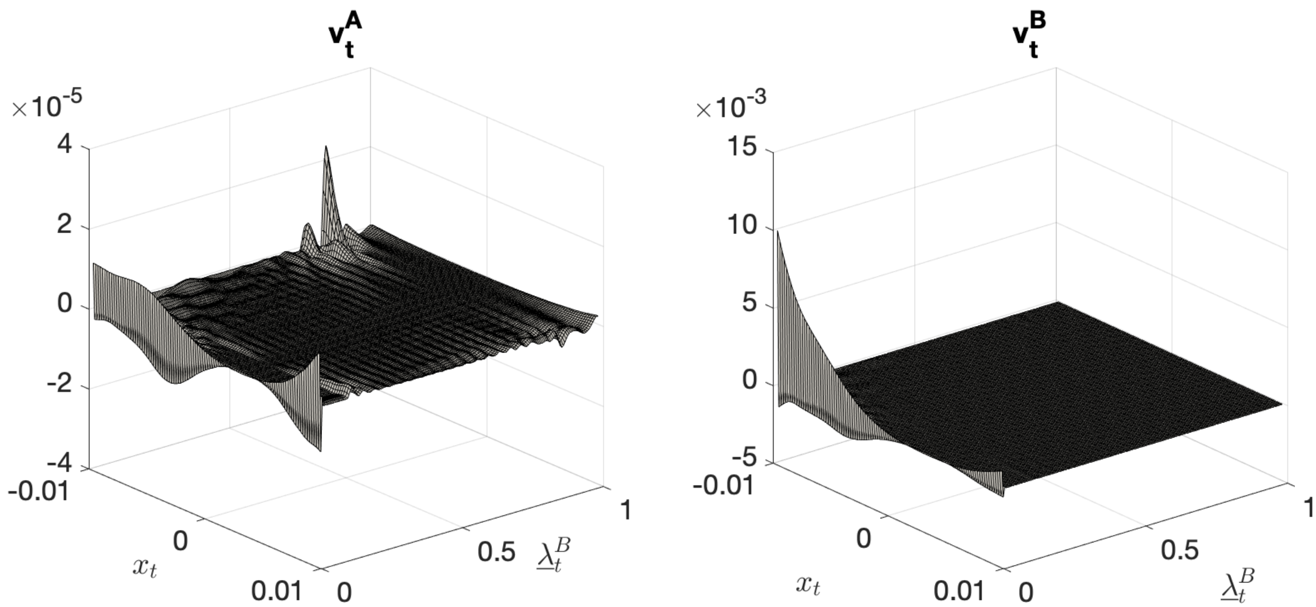

The figure plots numerical errors in the value functions (12) for $\rho_x^A = 0.99$ and $\rho_x^B = 0.96$. A state grid of ± 4 standard deviations around the unconditional mean of the x_t-process is used. For $\underline{\lambda}_t^B$ the full grid between 0 and 1 is used. The projection method used $n = 50$ collocation nodes for the $\underline{\lambda}_t^B$-dimension and $m = 22$ nodes for the x_t-dimension. Results are shown for the adjusted grid to account for the non-linearity close to $\underline{\lambda}_t^B = 0$.

Since $x_{+1} \sim N(\rho_x x_t, \sigma_x^2)$, the probability ratio is given by

$$\begin{aligned}\log\left(\frac{\mathrm{dP}_{t,t+1}^A}{\mathrm{dP}_{t,t+1}^B}\right) &= \log\left(e^{-0.5\frac{(\rho_x x_t + \eta_{x,t+1} - \rho_x^A x_t)^2}{\sigma_x^2} + 0.5\frac{(\rho_x x_t + \eta_{x,t+1} - \rho_x^B x_t)^2}{\sigma_x^2}}\right) \\ &= \frac{x_t^2}{2\sigma_x^2}\left((\rho_x - \rho_x^B)^2 - (\rho_x - \rho_x^A)^2\right) + \frac{x_t}{\sigma_x^2}\left(\rho_x^A - \rho_x^B\right)\eta_{x,t+1}. \qquad (67)\end{aligned}$$

Hence, we find that the consumption share in $t+1$ is a linear function of $\eta_{x,t+1}$,

$$\log\left(s_{t+1}^A\right) = a^{CRRA} + b^{CRRA}\eta_{x,t+1} \qquad (68)$$

and the coefficients are given by

$$\begin{aligned} b^{CRRA} &= \frac{(1 - s_0^A) x_t (\rho_x^A - \rho_x^B)}{\sigma_x^2 \gamma} \\ a^{CRRA} &= \log\left(s_0^A\right) + \frac{(1 - s_t^A) x_t^2}{2\sigma_x^2 \gamma}\left[(\rho_x - \rho_x^B)^2 - (\rho_x - \rho_x^A)^2\right]. \end{aligned}$$

The slope b^{CRRA} determines how the consumption share of investor A changes in response to shocks to x_{t+1}. Assume that $\rho_x^A > \rho_x^B$. The sign of b^{CRRA} depends on the sign of x_t. If

Figure 15: Numerical Errors for $\rho_x^A = 0.99, \rho_x^B = 0.96, n = 80$, and $m = 22$

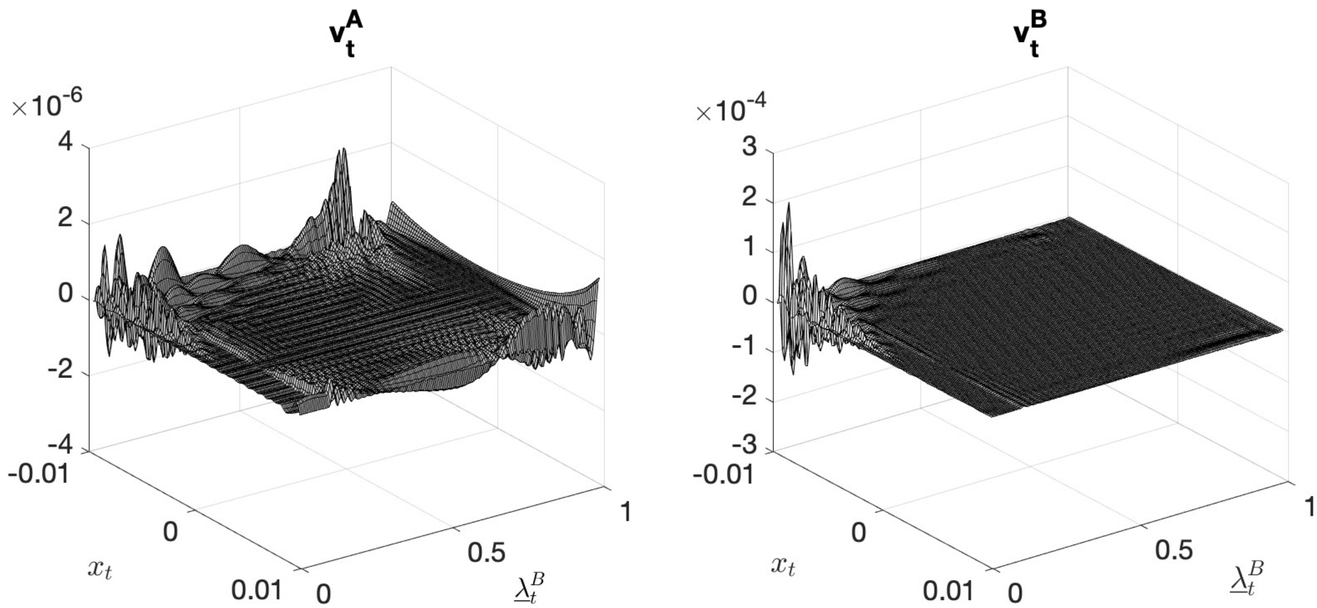

The figure plots numerical errors in the value functions (12) for $\rho_x^A = 0.99$ and $\rho_x^B = 0.96$. A state grid of ± 4 standard deviations around the unconditional mean of the x_t-process is used. For $\underline{\lambda}_t^B$ the full grid between 0 and 1 is used. The projection method uses $n = 80$ collocation nodes for the $\underline{\lambda}_t^B$-dimension and $m = 22$ nodes for the x_t-dimension. Results are shown for the adjusted grid to account for the nonlinearity close to $\underline{\lambda}_t^B = 0$.

x_t is positive (negative), b^{CRRA} is positive (respectively, negative); this sign implies that the larger the shock to x_{t+1}, the larger (smaller) will be $\log s_{t+1}^A$ and, hence, the larger (smaller) the consumption share of agent A. The intuition is that investor B believes in faster mean reversion and hence puts more probability weight on states where x_{t+1} moves toward its long-run mean of 0. So investors bet on states depending on the subjective probabilities they assign to those states. We call this motivation for investments the "speculation motive" of the investors. This motive increases with $|\rho_x^A - \rho_x^B|$ and $|x_t|$ and decreases with the risk aversion γ. The larger the difference in the beliefs, $|\rho_x^A - \rho_x^B|$, the larger is the difference in the probabilities that the investors assign to different states. For $x_t = 0$ investors share the same beliefs but the larger $|x_t|$ is, the more important becomes the difference in the beliefs about the speed of mean reversion. Finally, the more risk averse investors are, the less they are willing to speculate on future outcomes.

We observe that this speculation motive is independent of the true persistence ρ_x. However, the true persistence does influence the average change in the consumption share. Assume that investor A has the correct beliefs, $\rho_x = \rho_x^A$. The average change in the log consumption share is then given by

$$\begin{aligned} E_t\left(\log\left(s_{t+1}^A\right)\right) - \log\left(s_t^A\right) &= a^{CRRA} - \log\left(s_t^A\right) \\ &= \frac{(1-s_t^A)x_t^2}{2\sigma_x^2\gamma}(\rho_x^A - \rho_x^B)^2 \geq 0. \end{aligned} \tag{69}$$

We observe that—independent of the states s_t^A and x_t and whether ρ_x^A is larger or smaller than ρ_x^B—the consumption share of investor A, who has the correct beliefs, will always increase on average. So for CRRA utility, the only thing that matters for the average change in the consumption shares is which investor has the correct beliefs. The speed at which he or she accumulates wealth depends on the risk aversion of the investor. So the more risk averse the investor, the less he or she will be willing to speculate on future outcomes and, hence, the slower will be the wealth accumulation.

E Additional Details for the Model with Time-Varying Persistence

Below we provide the state vectors as well as the Markov transition probabilities for the model with time-varying persistence used in Section 4.5. For the two-state economy, the state vector is given by $\rho_{x,t} = [0.9686, 0.9798]$ and the Markov transition probabilities are given by

$$P = \begin{bmatrix} 0.9963 & 0.0037 \\ 0.0037 & 0.9963 \end{bmatrix}.$$

For the nine-state economy we have

$$\rho_{x,t} = [0.9584, 0.9623, 0.9663, 0.9702, 0.9742, 0.9782, 0.9821, 0.9861, 0.9900]$$

with

$$P = \begin{bmatrix} 0.9708 & 0.0288 & 0.0004 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.0036 & 0.9709 & 0.0252 & 0.0003 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.0072 & 0.9709 & 0.0216 & 0.0002 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.0108 & 0.9710 & 0.0180 & 0.0001 & 0 & 0 & 0 \\ 0 & 0 & 0.0001 & 0.0144 & 0.9710 & 0.0144 & 0.0001 & 0 & 0 \\ 0 & 0 & 0 & 0.0001 & 0.0180 & 0.9710 & 0.0108 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0.0002 & 0.0216 & 0.9709 & 0.0072 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.0003 & 0.0252 & 0.9709 & 0.0036 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0.0004 & 0.0288 & 0.9708 \end{bmatrix}.$$

F Additional Results

This section presents additional results, which have been referenced in the main body of the paper. Figure 16 shows the conditional mean and variance of continuation utilities for the discrete state economy in the good state; see Section 3.2.

In Section 3.1, we examine the distribution of the persistence of forecaster predictions using a panel regression. An alternative approach is to consider the regression coefficients of individual forecasters in the following regression,

$$\Delta c'_{t,i} = A_i + \beta_i \Delta c_{t,i} + \epsilon_{t,i}. \tag{70}$$

Figure 17 shows the distribution of β_i coefficients.

Figure 16: Conditional Mean and Variance of Continuation Utilities: Good State

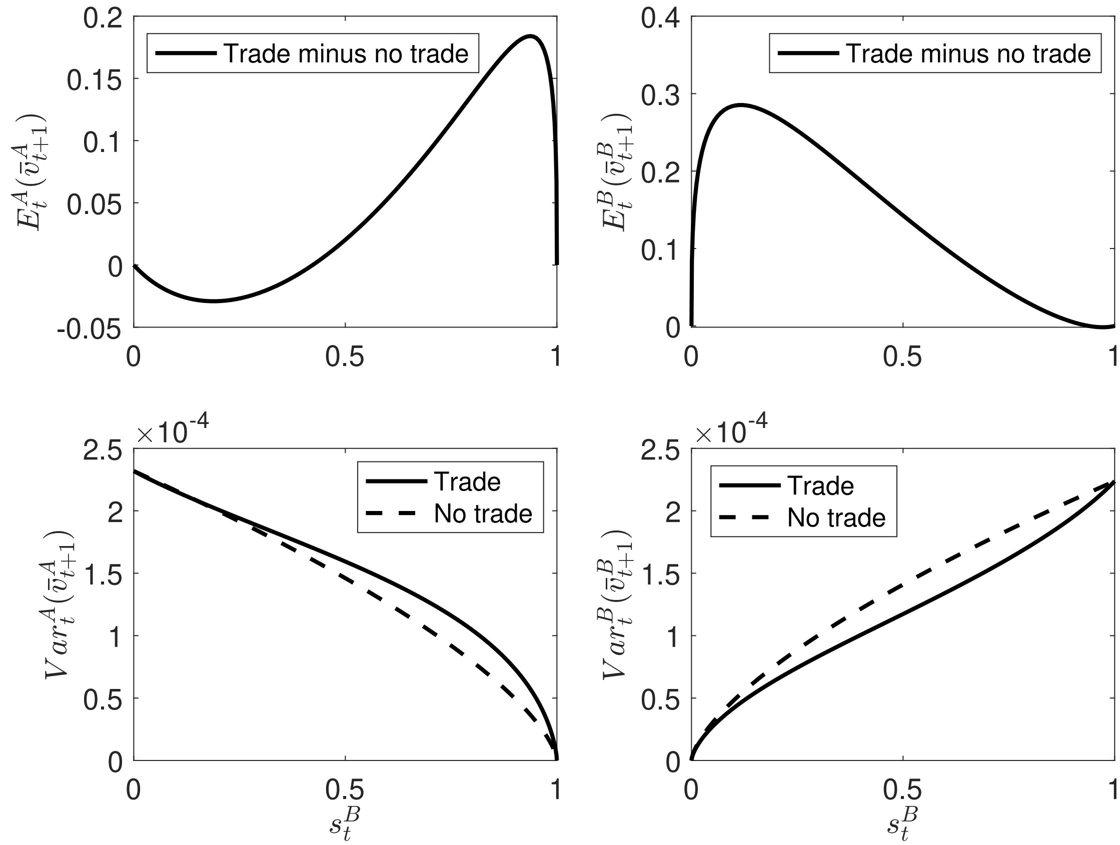

The figure plots the conditional mean and variance of continuation utilities as a function of the consumption share of investor B for the discrete state economy. The top panel shows the difference in expected utility between an economy in which the agents are allowed to trade and a no-trade assumption. The lower panel shows the conditional variance for both the trade and no-trade case. The left panel shows the results for investor A, who believes in the higher persistence and the right panel for investor B, who believes in lower persistence. Results are shown for economy being in the good state.

Figure 17: Individual Forecaster AR(1) Coefficients

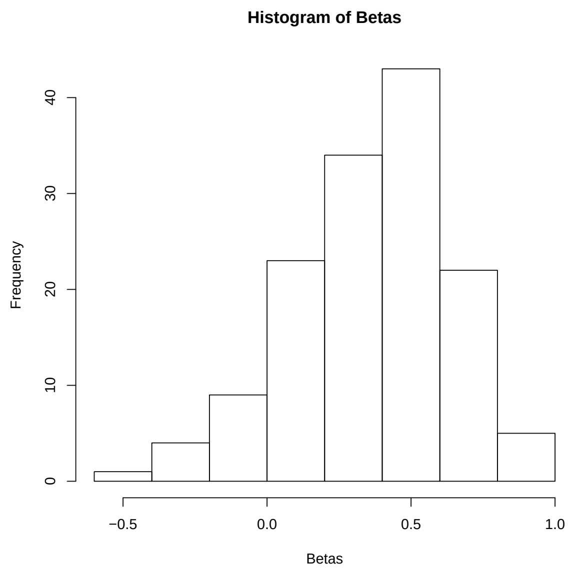

The figure shows a histogram of the AR(1) coefficients (β_i) for real-consumption-growth forecasts, see model (70), for a group of forecasters from the U.S. Survey of Professional Forecasters. Only forecasters with predictions in at least eight surveys are included in the sample. For the presentation of the histogram, the two highest and the two lowest coefficients have been removed.