

Asset Pricing when 'This Time is Different'

Pierre Collin-Dufresne, Michael Johannes and Lars A. Lochstoer*

EPFL

Columbia Business School

September 2015

Abstract

Recent evidence suggests that the young update beliefs about macro outcomes more in response to aggregate shocks than the old. We embed this experiential learning bias in a general equilibrium macro-finance model where agents have recursive preferences and are unsure about the specification of the exogenous aggregate stochastic process. The departure from rational expectations is small in a statistical sense, but generates quantitatively significant increases in risk, as well as substantial and persistent aggregate over- and under-valuation. Consistent with the model, the price-dividend ratio is empirically more sensitive to macro shocks when the proportion of young vs. old is high.

*Pierre Collin-Dufresne is at the Swiss Finance Institute at Ecole Polytechnique Federale de Lausanne. Michael Johannes and Lars A. Lochstoer are at Columbia Business School. We thank Anmol Bhandari, Kent Daniel, Lorenzo Garlappi, Nicolae Garleanu, Paul Tetlock, Colin Ward, and seminar participants at the AFA 2015, Columbia Business School, Duke, NYU, the Macro-Finance Conference 2014, the Minnesota Macro Asset Pricing Conference 2015 (at Carlson SOM), Rochester, Stockholm School of Economics, University of British Columbia, University of Miami, and Yale School of Management for helpful comments. Contact info: Lars A. Lochstoer, 405B Uris Hall, Columbia Business School, Columbia University, 3022 Broadway, New York, NY 10027. E-mail: LL2609@columbia.edu. First draft: November 2013.

1 Introduction

The macro-finance literature typically assumes agents have Rational Expectations and use the entire history of events and Bayes rule to form statistically optimal beliefs. The psychology literature, however, argues that personal experience is more salient and therefore exert a greater influence on agents' decision making than summary information available in historical records.¹ Consistent with the latter view, recent empirical evidence suggests that individual macroeconomic belief formation is in fact subject to age-related experiential learning bias. For instance, Nagel and Malmendier (2013) present direct evidence from survey data that the sensitivity of agents' inflation beliefs to a shock to inflation is decreasing with the age of the agent. In other words, when learning about the economic environment, the young update more in response to shocks than the old, consistent with the notion that the young have more dispersed prior beliefs due to their shorter personal history.²

While intriguing in itself, it is not a priori clear that such cross-sectional evidence on belief biases has first-order relevance for models of aggregate asset prices and macroeconomic dynamics. Consistent with the view that individual biases wash out in the aggregate, Ang, Bekaert, and Wei (2007) document, also using survey data on inflation expectations, that the median inflation forecast outperforms pretty much any other forecast they construct from available macro and asset price data. Thus, the median belief across agents appears to be quite 'rational.' Further, if agents disagree about states that are particularly important for asset prices and marginal utilities, such as a Depression state, there are large gains to trade and an optimistic agent may be willing to provide ample insurance to a 'rational' agent. Indeed, Chen, Joslin, and Tran (2012) show that only a small fraction of optimistic agents are needed in order to eliminate most of the risk premium associated with disaster risk. Thus, allowing for belief heterogeneity can affect the asset pricing performance of standard models, but often in a way that reduces risk and thus makes it harder to fit the stylized facts.

In this paper, we find that generational biases of the form discussed above can have a significant impact on the joint dynamics of macro aggregates, asset prices, and risk premia that helps account for stylized asset pricing moments. We do so by embedding agents with recursive utility that have an experiential learning bias and overlapping generations into otherwise standard macro-finance models. We consider models with and without severe crisis events, in order to assess the effects of the bias on aggregate asset prices in a variety of settings.

¹e.g., Nisbett and Ross (1980), Weber, Boeckenholt, Hilton and Wallace (1993), Hertwig, Barron, Weber, and Erev (2004).

²In other work, Nagel and Malmendier (2011) argue that investors who experienced the Great Depression are more pessimistic about stock returns than (younger) investors who did not. Generational learning bias also present for investor return expectations over the dot-com boom (Vissing-Jorgensen, 2003), for mutual fund managers (Greenwood and Nagel, 2009), in Europe (Amphudia and Ehrmann, 2014).

We calibrate the magnitude of the belief bias to the micro estimates of age effects in macroeconomic expectation formation in Nagel and Malmendier (2013), and assume agents are Bayesian when they update beliefs based on data realized in their lifetime. Together, these restrictions ensure that the average agent's beliefs about macro outcomes such as consumption or GDP growth are very close (in a likelihood ratio sense) to being 'Rational.' The experiential learning bias in the model arises as the Young, when born, are endowed with prior beliefs that are more dispersed than the dying Old's posterior beliefs. One can think of this either as an inability to efficiently process information from before one's lifetime or as some information about the economic environment that cannot be communicated from previous generations to the new Young. Thus, the Young suffer from a 'This Time is Different'-bias in that they treat their birth effectively as a structural break in terms of forming expectations about the future. That is, relative to what the full historical record indicates, the Young will see a sequence of positive shocks as evidence of a higher average growth rate, they will see the occurrence of a severe crisis as a signal that such crisis are more likely to occur again in the future, and so on.

Specifically, agents have Epstein-Zin preferences and are uncertain about the specification of the exogenous aggregate stochastic process. There are two generations alive at each point in time, young and old. Each generation lives for 40 years, so there is a 20 year overlap between generations. When born, agents inherit the mean beliefs about the model specification from their parent generation (who die and are the previously Old), but with a prior variance of beliefs that is higher than the posterior variance of their parent generation's beliefs. We consider the particular cases where agents are unsure about the mean growth rate of the economy or the probability of a severe crisis state. A fully rational, Bayesian agent would eventually learn the true model, but due to the 'this time is different' OLG feature of the model, parameter learning persists indefinitely in this economy.

First, even though agents beliefs' are close to rational, the generational nature of the bias implies that mistakes are highly persistent. For that reason, the aggregate market typically displays over- and undervaluation of the order of $\pm 20\%$ relative to the rational expectations versions of the models we consider, with the 99% interval of misvaluation from about -50% to $+100\%$ of fundamental value. Aggregate beliefs fluctuate around the true values, which leads to long-run excess return predictability.

Second, this excess volatility helps resolve standard asset pricing puzzles when agents have Epstein-Zin preferences. Indeed, as pointed out in Collin-Dufresne, Johannes and Lochstoer (CJL 2015) revisions in beliefs due to learning generate subjective long-run risks that can dramatically amplify macroeconomic shocks and lead to high ex-ante prices of risk even though agents have low risk aversion and consumption growth is actually i.i.d.. As in Bansal and Yaron (2004) these long run risks are priced only if agents have a preference for the early

resolution of uncertainty.

A central prediction of the model is that the aggregate valuations are more sensitive to macroeconomic shocks when the young control more of the total wealth (including human capital) in the economy. Consistent with this, we document in the data that the annual change in the price-dividend ratio is more sensitive to annual contemporaneous GDP growth when the fraction of young vs. old in the economy is high. Further, 10-year changes in the price-dividend ratio are significantly negatively correlated with 10-year changes in the fraction of young vs. old, again consistent with the model where the young perceive more risk as they are more unsure about the specification of the data generating process.

Even though agents in the model are learning only from fundamentals (macroeconomic shocks), past stock market returns can positively impact investors' assessment of future returns in the model. It is well-documented that investors tend to extrapolate from recent past stock returns when forming expectations of future stock returns (see Greenwood and Shleifer (2014) for a survey)—a feature of the data that is hard to match in a model where agents use only fundamental information when forming beliefs (see Barberis, Greenwood, Jin, and Shleifer (2014)).

Further, we show, in the case of learning about the probability of a Depression, that belief uncertainty combined with recursive preferences decreases the impact of optimists on asset prices. In particular, we find that while the risk premium is decreasing in the fraction of optimists, the effect is much smaller than in the case where agents are certain about their beliefs, as is the case in Chen, Joslin, and Tran (2012). In particular, optimists are less willing to provide disaster state hedges as a disaster realization leads to an adverse update in beliefs about the disaster probability. This parameter uncertainty risk makes agents less willing to speculate based on their current average beliefs.

In terms of the dynamics that arise from having heterogeneous agents, the optimal risk-sharing in the Epstein-Zin model tends to exacerbate the impact of biased beliefs on asset prices as the more optimistic (pessimistic) agents hold more (less) stock. A positive (negative) shock is therefore amplified in terms of the wealth-weighted average belief in this model. This endogenous amplification of shocks is much stronger when agents have Epstein-Zin preferences as there is a larger difference in the impact of model risk on utility across the generations when agents are averse to model uncertainty (when they have a preference for early resolution of uncertainty). This means the average difference in portfolio holdings across generations is also large. This is very different from the power utility case where model uncertainty has much less impact on utility.

There are four state-variables in each of the models. Solving the endogenous risk sharing problem is non-trivial when agents have Epstein-Zin preferences. We solve the model using a new robust numerical solution methodology developed by Collin-Dufresne, Johannes, and

Lochstoer (2013b) for solving risk-sharing problems in complete markets when agents have recursive preferences. This numerical method does not rely on approximations to the actual economic problem (e.g., it does not rely on an expansion around a non-stochastic steady-state) and therefore provides an arbitrarily accurate solution (depending of course on the chosen coarseness of grids and quadratures).³

Related literature. There is a large literature on the effects of differences in beliefs on asset prices. Harrison and Kreps (1978) and Scheinkman and Xiong (2003) show how overvaluation can arise when agents have differences in beliefs and there are short sale constraints. Dumas, Kurshev, and Uppal (2009) consider a general equilibrium, complete markets model where two agents with identical power utility preferences disagree about the dynamics of the aggregate endowment. Bhamra and Uppal (2014) consider two agents with heterogeneous beliefs, different risk aversion and “catching up with the Joneses” preferences. Baker, Hollifield, and Osambela (2014) consider a general equilibrium production economy where two power-utility agents have heterogeneous beliefs about the mean productivity growth rate, where the agents agree-to-disagree and do not update their beliefs (static beliefs). The authors show that speculation leads to a counter-cyclical risk premium and that the investment and stock return volatility dynamics are counter-cyclical when agents have high elasticity of intertemporal substitution.

The main contribution of our paper is to consider agents with Epstein-Zin preferences as well as overlapping generations, which together with the experiential learning bias generates the cross-sectional heterogeneity in beliefs. It turns out that allowing for preferences for early resolution dramatically changes the quantitative predictions of the model because, as shown in CJL 2015, learning generates subjective long-run risks for an agent with preferences for early resolution of uncertainty. We illustrate this in the paper by comparing with the standard time separable CRRA utility case. In contrast to CJL 2015 who only consider learning in a rational representative agent model, the small deviation from rationality we introduce here generates long-term (stationary) deviations of prices from fundamentals, which in turn generates excess return predictability. These price deviations are driven by small, but persistent mistakes in cash flow expectations. However, an econometrician with a long sample, who therefore estimates i.i.d. cash flow growth, would attribute all of the variation in the price-dividend ratio to discount rate shocks. Thus, the joint dynamics of consumption and asset prices in the model are, from the perspective of the econometrician, similar to those in the habit formation model of Campbell and Cochrane (1999)—even though from the perspective of the agents within the model, risk and price fluctuations are driven by long-run cash flow risk as

³Accurate solutions do require efficient coding in a fast programming language, such as C++ or Fortran, and extensive use of the multiprocessing capability of high-performance desktops. Such technology is, however, easily available.

in Bansal and Yaron (2004) and CLJ 2015.

A new feature of our model is that agents are not only heterogeneous with respect to their mean beliefs, but also with respect to the confidence they exhibit in their beliefs. This is an important feature when agents have recursive preferences as the level of confidence (the precision of posteriors) determines the magnitude of updates in beliefs, which are priced with these preferences. Bansal and Shaliastovich (2010) present an asset pricing model with confidence risk in a representative agent setting. In our model, the different levels of confidence are strong determinants of the optimal risk sharing.

In contemporaneous work, Ehling, Graniero, and Heyerdahl-Larsen (2013) consider a similar learning bias in an OLG endowment economy framework, but with log utility preferences. These authors present empirical evidence that the expectations of future stock returns are more highly correlated with recent past returns for the young than the old, consistent with the overall evidence given by Malmendier and Nagel (2011, 2013). Nakov and Nuno (2015) consider an OLG economy where agents have power utility preferences and learn about both the mean growth rate in market equity prices and in dividends using a constant gain learning rule. In order to solve the model, the agents face constraints on the minimum and maximum asset exposure such that the stock price equals the reservation price of the marginal stock bidder. In other recent work, Choi and Mertens (2013) solve a model with two sets of infinitely-lived agents with Epstein-Zin preferences, where one set of agents has extrapolative beliefs, in an incomplete markets setting with portfolio constraints. These authors estimate the size of the belief bias by backing it out from standard asset price moments, whereas we calibrate the bias to available micro estimates as given by Malmendier and Nagel (2011, 2013) and solve an OLG model. Both these authors and Dumas, Kurshev, and Uppal (2009) have only one set of agents with biased beliefs, while the generational 'this time is different'-bias leads to multiple agents with biased beliefs. Thus, the nature of the OLG problem we solve has more state variables as we need to keep track of the individual beliefs of multiple generations (the young and the old in our case). Barberis, Greenwood, Jin, and Shleifer (2014) propose a model where there are two sets of CARA utility agents—extrapolators and rational—where the former form beliefs about future asset returns by extrapolating past realized returns, consistent with survey evidence on investor beliefs. In terms of heterogeneous agent models with Epstein-Zin preferences, Garleanu and Panageas (2012) solve an OLG model with Epstein-Zin agents with different preference parameters, while Borovicka (2012) shows long-run wealth dynamics in a two-agent general equilibrium setting where agents have Epstein-Zin preferences and differences in beliefs. Finally, Marcet and Sargent (1989), Sargent (1999), Orphanides and Williams (2005a), and Milani (2007) are prominent examples of the effects of perpetual, non-Bayesian learning in macro economics.

2 The Model

General equilibrium models with parameter learning and heterogenous beliefs are difficult to solve as the state space quickly becomes prohibitively large. For that reason, we focus on settings that are not only simple and tractable, but also quantitatively interesting and which can be easily calibrated to the microeconomic evidence presented in Malmendier and Nagel (2013).

We assume there are two sets of agents alive at any point in time, young and old. A generation lasts for T periods, and each agent lives for $2T$ periods. Thus, there is no uncertainty about life expectancy. All young and old agents currently alive were born at the same time, and agents born at the same time have the same beliefs. These assumptions imply that (a) there are no hedging demands related to uncertain life span, and (b) there is a two-agent representation of the economy. The latter is important in order to minimize the number of state variables. The former is a necessary assumption for the latter to be true given our learning problem, as shown below.

When an old generation dies, the previously young generation becomes the new old generation and a new young generation is born. The old leave their wealth for their offspring. In terms of beliefs, the new young inherit their parent's mean beliefs in a manner that will be made precise below.⁴ The bequest motive is similar to that in a Dynasty model, that is, the parents care as much about their offspring as themselves (with the usual caveat that there is time-discounting in the utility function). Thus, there are two representative agents from each Dynasty, A and B. Figure 1 provides a timeline of events related to the cohorts of each dynasty in the model.

2.1 Aggregate dynamics and cohort belief formation

The agents in the economy are not able to learn the true model specification for aggregate consumption dynamics due to an experiential learning bias. In particular, we assume agents are Bayesian learners with respect to data they personally observe, e.g., aggregate consumption growth realized during their lifetime. However, they downweight data prior to their lifetime in the following way: the young inherit the mean beliefs about the model of consumption growth from their parents (the dying old), but they are endowed with more dispersed initial

⁴The labels 'old' and 'young' in this model refer to the two generations currently alive. A new generation could be born, say, every 20 years, which implies the investors in this economy live for 40 years. When the old 'die' they give life to new 'young,' and so 'death' may be thought of as around age 70 and the new 'young' as around 30 years of age. In other words, the model is stylized in order to capture a 'this time is different'-bias related to personal experience in a quantitatively interesting and tractable setting.

Figure 1 - Model timeline

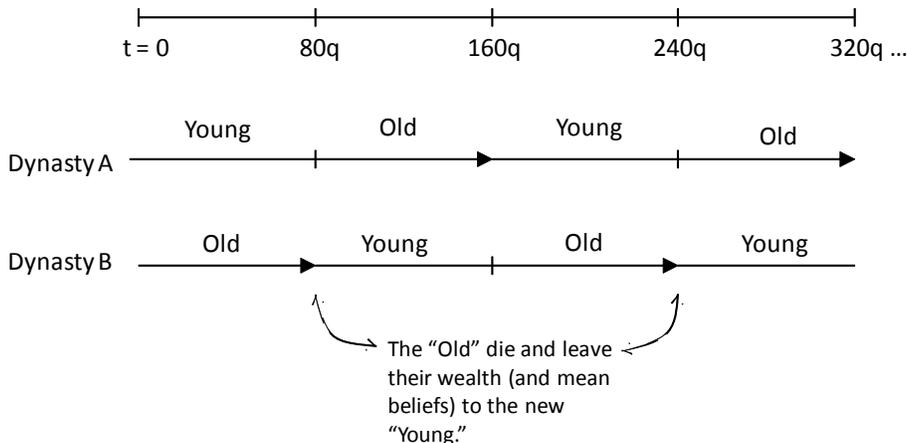


Figure 1: The plot shows the timeline of the model over an 80 year period (the model is an infinite horizon model, so the pattern continues ad infinitum). Model time is in quarters, and a generation lasts for 20 years (80 quarters), while agent's "investing lives" are 40 years. Upon death, represented as an arrowhead in the figure, the Old leave their wealth to their offspring—the new Young. The Young also inherits their parent generations mean beliefs about model parameters, but start their lives with a prior variance of beliefs that is higher than their parents posterior dispersion of beliefs. It is the latter "This Time is Different" bias that makes experiential learning important for belief formation.

beliefs or uncertainty than their parents had at the end of their life. This is the source of the 'This Time is Different'-bias.

We assume aggregate consumption growth is i.i.d., with both standard normal shocks and 'disasters,' a small probability of a large negative consumption drop:

$$\Delta c_{t+1} = \mu + \sigma \varepsilon_{t+1} + d_{t+1}, \quad (1)$$

where $\varepsilon_{t+1} \stackrel{i.i.d.}{\sim} N(0, 1)$, and $d_{t+1} = \underline{d} \ll 0$ with probability p and zero otherwise, similar to the specification in Barro (2006). We calibrate the size of the consumption drop to the U.S. Great Depression experience. We assume σ and \underline{d} are known and that both Δc_{t+1} and d_{t+1} are observed, but that μ and p are unknown.⁵

The time t posterior beliefs of agent i about μ are $N(m_{i,t}, A_{i,t}\sigma^2)$, where beliefs are

⁵In continuous-time, σ would be learned immediately. The constant \underline{d} would be known with certainty after its first realization. Since we consider large jumps, it would be easy for the agent to assess whether $d_{t+1} = 0$ or not. Thus, we lose little by assuming d_{t+1} is directly observed, but gain tractability. It is, however, hard to learn μ and p , which is why we focus on these parameters.

updated according to Bayes rule:

$$m_{i,t+1} = m_{i,t} + \frac{A_{i,t}}{1 + A_{i,t}} (\Delta c_{t+1} - d_{t+1} - m_{i,t}) \quad (2)$$

and

$$A_{i,t+1} = \frac{A_{i,t}}{1 + A_{i,t}}. \quad (3)$$

Agent i 's time t posterior beliefs about p are Beta distributed, with $p \sim \beta(a_{i,t}, A_{i,t}^{-1} - a_{i,t})$. The properties of the Beta distribution and Bayes rule then imply that:

$$E_t^i [p] = a_{i,t} A_{i,t} \quad \text{and} \quad Var_t^i [p] = \frac{A_{i,t}}{1 + A_{i,t}} E_t^i [p] (1 - E_t^i [p]) \quad (4)$$

and

$$a_{i,t+1} = \begin{cases} a_{i,t} + 1 & \text{if } d_{t+1} = \underline{d} \\ a_{i,t} & \text{if } d_{t+1} = 0 \end{cases}, \quad (5)$$

where the updating equation for $A_{i,t}$ is as in Equation (3). Agent i 's mean belief of the likelihood of a disaster event thus evolves as:

$$E_{t+1}^i [p] = E_t^i [p] + (\mathbf{1}_{d_{t+1}=\underline{d}} - E_t^i [p]) \frac{A_{i,t}}{1 + A_{i,t}}. \quad (6)$$

First, note that if agent i were to live forever, $A_{i,\infty} = 0$ and the variance of her subjective beliefs about both μ and p would go to zero. Further, mean beliefs would converge to the true parameter values, $m_{i,\infty} = \mu$, and $E_\infty^i [p] = p$. However, the generational ‘This Time is Different’-bias implies that learning persists indefinitely. In particular, for a time t that corresponds to the death of the current old generation, let the posterior beliefs of the old be the sufficient statistics $m_{old,t}$, $a_{old,t}$ and $A_{old,t}$. The new young are then assumed to be born and consume at time $t + 1$ with prior beliefs $m_{young,t} = m_{old,t}$, $A_{young,t} = k A_{old,t}$, where $k > 1$, and $a_{young,t} = k a_{old,t}$ (i.e., $E_t^{old} [p] = E_t^{young} [p]$). Thus, the mean parameter beliefs are inherited by the young, but the prior dispersion of the beliefs of the young is higher than the posterior dispersion of the old. The constant k determines the amount of the experiential learning bias, and we set $k = A_0 / (A_0^{-1} + 2T)^{-1}$ such that the prior belief dispersion parameter of the young is always $0 < A_0 < 1$, which ensures that the beliefs process is stationary. Finally, we assume that the young and old generations living concurrently do not mutually update, that is, they ‘agree to disagree.’

As should be clear from the preceding discussion, the updating scheme with the ‘This Time is Different’-bias implies that the Young place too much weight on personal experience relative to a full-information, known parameters benchmark case, consistent with the micro

evidence presented by Malmendier and Nagel (2011, 2013). We compare the relation to their evidence in more detail in the calibration section.

2.2 Utility and the bequest motive

We assume agents have Epstein and Zin (1989) recursive preferences. In particular, the value function $V_{i,t}$ of agent i alive at time t who will ‘die’ at time $\tau > t + 1$ is:

$$V_{i,t}^\rho = (1 - \beta) C_{i,t}^\rho + \beta E_t^i [V_{i,t+1}^\alpha]^\rho. \quad (7)$$

Here, $\rho = 1 - 1/\psi$ where ψ is the elasticity of intertemporal substitution (EIS) and $\alpha = 1 - \gamma$, where γ is the risk aversion parameter.

In terms of birth and death, an agent’s last consumption date is τ , and at $\tau + 1$ the agent’s offspring, i' , comes to life and starts consuming immediately. The offspring have different beliefs about the aggregate endowment, as described earlier. We consider a bequest function of the form:

$$B_i(W_{i',\tau+1}) = \phi_{i'}(X_{\tau+1}) W_{i',\tau+1}, \quad (8)$$

where $\phi_i(\cdot)$ will be defined below, X_t is a vector of state variables and $W_{i,t}$ is the agent i ’s wealth at time t . State variables include all agents’ beliefs as well as a measure of the time each class of agent has been alive.

With this bequest function, we have that:

$$V_{i,\tau}^\rho = (1 - \beta) C_{i,\tau}^\rho + \beta E_\tau^i [\phi_{i'}(X_{\tau+1})^\alpha W_{i',\tau+1}^\alpha]^\rho. \quad (9)$$

Substituting in the usual budget constraint, we have:

$$V_{i,\tau}^\rho = (1 - \beta) C_{i,\tau}^\rho + \beta (W_{i,\tau} - C_{i,\tau})^\rho E_\tau^i [\phi_{i'}(X_{\tau+1})^\alpha R_{w_{i,\tau+1}}^\alpha]^\rho. \quad (10)$$

The first order condition over consumption implies that

$$\rho(1 - \beta) C_{i,\tau}^{\rho-1} = \rho\beta (W_{i,\tau} - C_{i,\tau})^{\rho-1} \mu_{i,\tau}, \quad (11)$$

equivalently,

$$\mu_{i,\tau} = \left(\frac{1 - \beta}{\beta}\right)^{1/\rho} \left(\frac{W_{i,\tau}}{C_{i,\tau}} - 1\right)^{(1-\rho)/\rho}, \quad (12)$$

where the certainty equivalent is $\mu_{i,\tau} = E_\tau^i [\phi_{i'}(X_{\tau+1})^\alpha R_{w_{i,\tau+1}}^\alpha]^{1/\alpha}$. Inserting this back into

the value function,

$$\frac{V_{i,\tau}}{W_{i,\tau}} = (1 - \beta)^{1/\rho} \left(\frac{W_{i,\tau}}{C_{i,\tau}} \right)^{1/\rho-1}. \quad (13)$$

The W/C ratio is a function of the state variables X_t . Let:

$$\phi_i(X_t) = (1 - \beta)^{1/\rho} \left(\frac{W_{i,\tau}}{C_{i,\tau}} \right)^{1/\rho-1}. \quad (14)$$

Then, $V_{i,t} = \phi_i(X_t) W_{i,t}$ for each t during the life of agent i . Since i was a general agent, it follows that $B_i(W_{i',\tau+1}) = V_{i',\tau+1}$. In this sense, the bequest function is dynastic, where the agent cares ‘as much’ about their offspring as themselves. Note that the expectation of the offspring’s indirect utility is taken using the parent generation’s beliefs. Thus, each dynasty can be represented as an agent that has the dispersion of beliefs reset every $2T$ periods as in the generational belief transmission explained in Section 2.1.

When there is no model/parameter uncertainty (that is, a full-information or the rational expectations case corresponding to $k = 1$ and $t = \infty$), the model reduces to an infinitely-lived Epstein-Zin representative agent with the same preference parameters as those assumed above (β, γ, ψ) . This agent, together with the maintained assumption of i.i.d. consumption growth, implies that the risk premium, the risk-free rate, the price-dividend ratio, and the price of risk are all constant in this benchmark economy.

2.3 The consumption sharing rule and model solution

We assume markets are complete, so each agent’s intertemporal marginal rates of substitution are equal for each state $(\Delta c_t, d_t)$. We index the two agents in the economy as belonging to Dynasty A or Dynasty B, where as explained above a Dynasty consists of a lineage of parent-child relations. Given Equations (7), (8), and (14), and the assumption of complete markets, we have that the two representative agents’ ratios of marginal utilities are equalized in each state (i.e., the stochastic discount factor is unique and both agents’ IMRS price assets given the respective agent’s subjective beliefs):

$$\begin{aligned} \pi^A(\Delta c_{t+1}, d_{t+1}|X_t) \left(\frac{c_{A,t+1}}{c_{A,t}} \right)^{\rho-1} \left(\frac{v_{A,t+1}}{\mu_{A,t}(v_{A,t+1}C_{t+1}/C_t)} \right)^{\alpha-\rho} = \dots \\ \pi^B(\Delta c_{t+1}, d_{t+1}|X_t) \left(\frac{1 - c_{A,t+1}}{1 - c_{A,t}} \right)^{\rho-1} \left(\frac{v_{B,t+1}}{\mu_{B,t}(v_{B,t+1}C_{t+1}/C_t)} \right)^{\alpha-\rho}. \end{aligned} \quad (15)$$

Here, X_t , which will be defined below, is the vector of state variables for the economy, including the sufficient statistics for each agent’s beliefs. The conditional beliefs about the joint state $(\Delta c_{t+1}, d_{t+1})$ can be further decomposed as $\pi^i(\Delta c_{t+1}, d_{t+1}|X_t) = \pi^i(\Delta c_{t+1} - d_{t+1}|m_{i,t}, A_{i,t}) \pi^i(d_{t+1}|a_{i,t}, A_{i,t})$

given the independence between ε_{t+1} and d_{t+1} and the assumption that agents agree-to-disagree. Further, in Equation (15), we set $c_{i,t} \equiv C_{i,t}/C_t$, $v_{i,t} \equiv V_{i,t}/C_t$ and impose the goods market clearing condition $c_{A,t} + c_{B,t} = 1 \iff C_{A,t} + C_{B,t} = C_t$ for all t .

With recursive preferences the value functions appear in the intertemporal marginal rates of substitution. Thus, unlike in the special case of power utility, Equation (15) does not provide us with an analytical solution for the evolution of the endogenous state variable—the relative consumption (or equivalently, wealth) of agent A . This complicates significantly the model solution. We solve the model using the numerical solution technique given in Collin-Dufresne, Johannes, and Lochstoer (2013b) using a backwards recursion algorithm that solves numerically for the consumption sharing rule starting at a distant terminal date for the economy, \tilde{T} . The solution corresponds to the infinite horizon economy when the transversality condition is satisfied and \tilde{T} is chosen sufficiently far into the future (e.g., 500+ years). The solution technique does not approximate the objective function and thus, with the caveat that it is numerical, provides an exact solution to the model. See the Appendix for further details.

The state variables in this model are $m_{A,t}$, $m_{B,t}$, $a_{A,t}$, $a_{B,t}$, $c_{A,t}$, and t . Time t is a sufficient statistic for $A_{(i=A,B),t}$ as $A_{i,t}$ is deterministic. We note that for general ρ and α , the prior distributions for the mean growth rate μ and the jump probability p must be truncated in order to have existence of equilibrium. The truncation bounds do not affect the updating equations, but $m_{i,t}$ and $a_{i,t}A_{i,t}$ in general no longer equal the conditional mean beliefs of μ and p .

2.4 Model Calibration

2.4.1 The belief process

The belief process of the stationary equilibrium in the model is governed by A_0 —the parameter that controls the severity of the 'This Time is Different'-bias. We calibrate this parameter to be consistent with the micro-evidence documented by Malmendier and Nagel (2013). In particular, Malmendier and Nagel (2013) estimate the sensitivity of the young (at age 30) to updates in beliefs from model learning to be about 2.5% of the size of the macro shock (in their case, quarterly inflation). Towards the end of their life (at age 70), the old have a sensitivity of about 1%. The estimates provided are based on inflation data and survey forecasts using available data in the post-WW2 period, and so these do not correspond to more extreme periods like the Great Depression. We therefore calibrate the value of A_0 to be 0.025, such that the updates in beliefs of the young from a 'regular' quarterly macro shock (ε in our model) is about 2.5% of the size of the macro shock, consistent with the estimate

of Malmendier and Nagel (2013).⁶ This implies that $k = 0.025/(1/0.025 + 160)^{-1} = 5$, when $T = 80$.

Figure 2 plots the weights that are put on lagged consumption data implied by our Bayesian-based learning scheme and the survey-implied estimates of Malmendier and Nagel (2013). The two learning schemes are quite close, though Bayesian learning has longer memory and is more efficient in that our agents more quickly put a lower weight on recent evidence. We choose Bayesian within-generation learning as a parsimonious way of ensuring that learning is consistent across different dimensions of uncertainty. We find this particularly useful since we consider learning both about the quarterly mean growth rate and about rare disasters. Note that Malmendier and Nagel’s estimates imply a zero weight on data from before one is born, which seems inappropriate for learning about rare events. Indeed, it seems more plausible that agents alive today put some weight on the Great Depression when valuing the stock market, even if they were not alive during the thirties. Our within-generation Bayesian agents do not forget past events, they just put lower weights on past events if they did not experience these.

To gauge how irrational our representative agents are we measure how long it would take on average to reject at the 5% level one Dynasty’s agents’ average subjective model (given by Equations 2)–(5)) relative to the true model (given by Equation (1)). We compute the sequential model probabilities over time, averaged across 100,000 simulations, given an initial model probability of 50/50 where the initial mean beliefs of the agent is centered around truth.⁷ Figure 3 gives the average evolution of the model probability of the agents’ model over time. When agents are learning both about μ (upper plot), it takes on average

⁶This calculation is based on the following. With a prior $\mu \sim N(m_0, A_0\sigma^2)$, the subjective consumption dynamics for the next period are:

$$\Delta c_1 = m_0 + \sqrt{A_0 + 1}\sigma\tilde{\varepsilon}_{t+1},$$

and the update in belief can be written:

$$m_1 = m_0 + \frac{A_0}{\sqrt{A_0 + 1}}\sigma\tilde{\varepsilon}_{t+1}.$$

Thus, the sensitivity of the update in mean beliefs to the macro shock when $A_0 = 0.025$ is $\frac{0.025}{\sqrt{0.025+1}} \approx 0.025$. The old then have a posterior sensitivity to shocks of 0.5% of the size of the shock, somewhat lower than that estimated by Malmendier and Nagel. Note that a different learning problem, for instance learning about a persistence parameter, would lead to slower learning relative to the simple learning about the mean case that we consider here. The micro estimates from Malmendier and Nagel do not correspond directly to the learning problem we consider, both since they allow for a non-Bayesian learning scheme and because they consider a different model (not just learning about a mean parameter). A more general learning model leads to many more state variables and is left for future research.

⁷The sequential updating of the probability of the ‘This Time is Different’ learning model versus the true iid model, $p_{M,t+1}$, is:

$$p_{M,t+1} = \frac{L(\Delta c_{t+1}, d_{t+1}|M_{TTiD}, a_t, m_t, A_t) p_{M,t}}{L(\Delta c_{t+1}, d_{t+1}|M_{TTiD}, a_t, m_t, A_t) p_{M,t} + L(\Delta c_{t+1}, d_{t+1}|M_{iid})(1 - p_{M,t})}.$$

**Figure 2 - Belief formation:
Weights on lagged data**

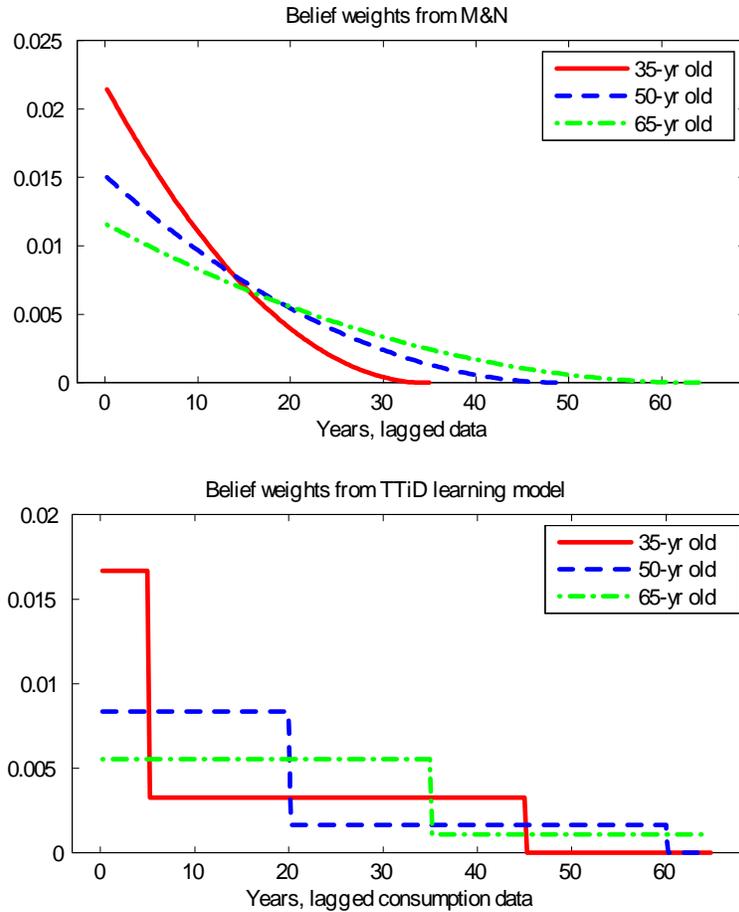


Figure 2: The top plot shows the weights the agent puts on increasingly lagged data when forming beliefs, as estimated by Malmendier and Nagel (2013). The solid line shows the weights corresponding to a 35-year old agent, the dashed line a 50-year old agent, and the dashed-dotted line a 65-year old agent. The weights are in this case zero for observations before the agent was born. The lower plot shows the corresponding weights for the Bayesian agents in the "This Time is Different"-model. The kink corresponds to a generational shift, assuming the agent is born as 'Young' at age 30. The weights for the preceding years are calculated using Bayes rule with the assumed increase in prior variance at each generational shift. The belief-weights are in this case flat within a generation due to Bayesian within-generation learning.

about 400 years before the 'This Time is Different'-bias is detected at the 5% level, while if the agent is learning about p it takes about 350 years (lower plot). Thus, it is very hard to learn, using only time series data, that the belief formation process of a representative agent is not correct. On the other hand, it is immediate to discover this from the the cross-section of agents as the Young update beliefs differently from the Old, even though they observe the

same shock.

Figure 3 - Model Probability: Agent Beliefs vs. Truth

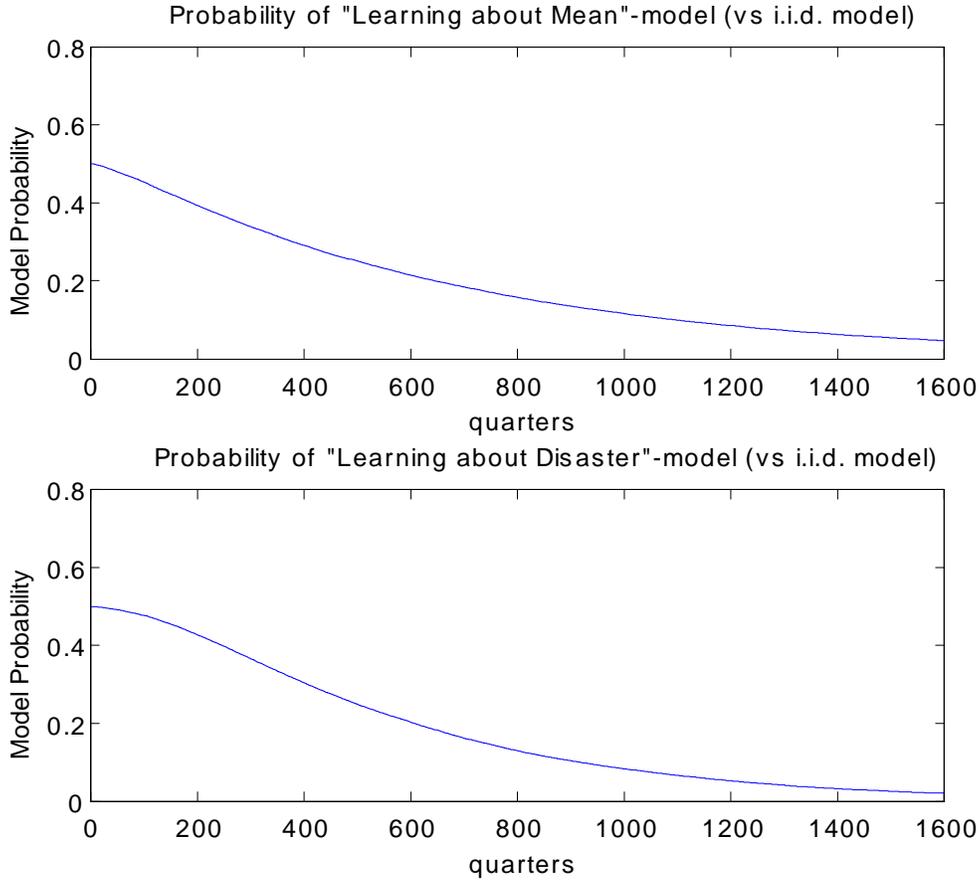


Figure 3: The top plot shows the probability of the subjective consumption dynamics as perceived by each agent, averaged across agents, versus the true model specification. The probabilities are in each case the mean outcomes across 20,000 simulations. The upper plot shows the case of "Learning about the Mean" and the lower plot shows the case of "Learning about a Depression Probability." The initial model probability is set at 50% and the x-axis is the observed sample length in quarters.

2.4.2 Preference and consumption parameters

We assume a generation lasts for 20 years (i.e., $T = 80$ quarters). We separately consider models with learning about the mean parameter μ or the jump probability p .⁸

⁸It is useful to analyze separately both models to understand the different economic implications of learning about frequent vs. rare events. Further, there is a computational advantage since it reduces the number of state variables. The general model (learning about both p and μ) would require about a month to solve with reasonable accuracy given our current computing capabilities.

In the 'Uncertain mean'-calibration, we let the preference parameters be $\gamma = 9$, $\psi = 1.5$ and $\beta = 0.994$, the true mean $\mu = 0.45\%$ and $\sigma = 1.35\%$, while $\underline{d} = 0$, i.e., there are no Depressions in this calibration. In the 'uncertain probability'-calibration, we let $\gamma = 4.5$, $\psi = 1.5$ and $\beta = 0.994$. Thus, the risk-aversion is half of that in the former case, but otherwise the preference parameters are the same. We set the risk aversion parameter in both models so that the Sharpe ratio on the equity claim is similar to that in the data. The true quarterly probability of a Depression, $p = 1.7\%/4$, is set consistent with the estimate used in Barro (2006), while the consumption drop in a Depression, \underline{d} , is -18% . Finally, we let $\mu = 0.53\%$ and $\sigma = 0.8\%$ in this calibration so as to match the mean and volatility of time-averaged consumption growth also in this case.

In the Great Depression, per capita real log consumption dropped by 18% from 1929 through 1933 (using data from the National Income and Product Accounts data from the Bureau of Economic Analysis). Of course, this four-year decline is quite different from a quarterly drop of 18% . However, since agents have Epstein-Zin preferences with $\gamma > 1/\psi$, the risk-pricing is related to the overall drop in consumption, so unlike for the power utility case, this distinction is not as important.⁹ We also consider power utility versions of the economies to assess the role of the EIS, ψ . Table 1 shows all relevant model parameters.

In order to ensure the existence of equilibrium, we truncate the priors for the mean growth rate and the Depression probability for the respective models. In particular, the upper (lower) bound for the prior over μ are 1.35% (-0.45%), while the upper (lower) bound for the prior over p are 0.04 (0.00001). Given that the prior standard deviations of beliefs when born are 0.21% for the mean and 0.01 for the probability, the truncation bounds are quite wide and therefore typically will not strongly affect the update in mean beliefs relative to the untruncated prior cases.

The equity claim is a claim to the exogenous dividend stream

$$\Delta d_{t+1} = \lambda \Delta c_{t+1} + \sigma_d \eta_{t+1}, \tag{16}$$

where $\lambda = 3$ is a leverage parameter (as in e.g., Abel (1999), Bansal and Yaron (2004)). The idiosyncratic shock, η_{t+1} , is normally distributed with mean μ_d and volatility σ_d set to match the mean and volatility of dividend growth in the data.

⁹Modeling true consumption as i.i.d. significantly simplifies the learning problem (in particular, it reduces the number of state variables relative to a more realistic, persistent Depression state).

Table 1 - Parameter values for Exchange Economy

Table 1: The top half of this table gives the preference parameters used in the two calibrations of the model. 'Uncertain Mean' refers to the calibration where log consumption growth is Normally distributed and agents are uncertain about the mean growth rate, while 'Uncertain Probability' refers to the case where log consumption growth also has a 'Depression' shock and where agents are uncertain about the probability of such a shock. The bottom half of the table gives the value for the parameters govern the consumption dynamics and agents beliefs. The numbers correspond to the quarterly frequency of the model calibration.

Preference parameters:		Uncertain Mean	Uncertain Probability
γ	(risk aversion parameter)	9	4.5
ψ	(elasticity of intertemporal substitution)	1.5	1.5
β	(quarterly time discounting)	0.994	0.994
Priors and consumption parameters:		Uncertain Mean	Uncertain Probability
A_0	('This Time is Different'-parameter)	0.025	0.025
T	(length of a generation in quarters)	80	80
σ	(volatility of Normal shocks)	1.35%	0.80%
μ	(true mean in 'normal times')	0.45%	0.53%
\bar{m}	(upper truncation point of μ prior)	1.35%	<i>n/a</i>
\underline{m}	(lower truncation point of μ prior)	-0.45%	<i>n/a</i>
\bar{p}	(upper truncation point of p prior)	<i>n/a</i>	0.04000
\underline{p}	(lower truncation point of p prior)	<i>n/a</i>	0.00001
p	(true probability of Depression)	<i>n/a</i>	0.00425
\underline{d}	(consumption shock in Depression)	<i>n/a</i>	-18%

3 Results

We first describe the dynamic portfolio allocations and implied risk-sharing of the two agents and thereafter focus on the asset pricing implications of the model.

3.1 Portfolio allocation and risk-sharing

As is typical in models with heterogeneous beliefs, we would expect that the more optimistic (pessimistic) agent will tend to hold a larger (smaller) portfolio share in risky assets. Below, we discuss how this equilibrium risk-sharing is affected by some of the novel features of our model, and in particular by (a) agents' differences in *confidence* in, or uncertainty over, their mean beliefs ($A_{A,t}$ vs $A_{B,t}$) and (b) recursive utility and therefore high perceived risk (and benefits from risk-sharing) arising from model uncertainty and learning.

Before we describe the portfolio allocations a couple of definitions are in order. First, total wealth in the economy is the value of the claim to aggregate consumption. Second,

since we solve a discrete-time, complete markets problem where one of the shocks has a continuous support (ε), the complete portfolio choice decisions of agents involve positions in principle in an infinite set of Arrow-Debreu securities. To convey the portfolio decisions of the agents in a simple (first-order) manner, we define the weight implicit in the total wealth portfolio of agent i as the local sensitivity of the return to agent i 's wealth to a small shock to total wealth (as arising from an aggregate ε shock close to zero). In the continuous-time limit for the 'Uncertain mean'-calibration, this local sensitivity is exactly the current portfolio allocation of agent i in the total wealth portfolio (because in this case, markets would be dynamically complete with two assets). When evaluating the 'Uncertain Probability'-calibration, we evaluate changes in relative wealth resulting from whether a Depression shock occurred or not.

3.1.1 'Uncertain Mean'-case

Figure 4 shows how risk-sharing operates in the 'Uncertain mean'-economy. In particular, the change in the relative wealth share of the Young is plotted against realizations of the aggregate shock (log aggregate consumption growth).¹⁰ The current wealth of the two agents is assumed equal and both agents are in the middle age of their respective generations (age 10yr and 30yr). In addition, the current mean beliefs of the Old are assumed to be unbiased, $m_{Old,t} = \mu$.

Two features of the model stand out. First, in the upper left plot, the case where the beliefs of the Young are also unbiased (the solid line) shows that the Old are in fact insuring the Young against bad states even when the mean beliefs coincide. This happens also when agents have power utility preferences ($\psi = 1/9$; see the lower left plot) because the Young perceive the world to be more risky than the Old as they are more uncertain about their mean beliefs about μ than the Old. Second, it is clear in the unbiased case that when $\gamma > 1/\psi$ the Old are insuring the Young to a larger extent than in the power utility case. This is because unlike power utility, recursive utility leads to a preference for early resolution of uncertainty that makes model uncertainty a priced risk. Therefore, the difference in confidence leads the Young to perceive the world as more risky, unconditionally. If the Young are sufficiently optimistic (here, about 2 standard deviations above the mean over a life time), the Young are in fact insuring the Old who are more pessimistic, and vice versa for the case where the Young are pessimistic (about 2-standard deviations below the mean over a life time).

The right-hand plots of Figure 4 show the portfolio weight of the Old versus the Young over the span of a generation (80 quarters). The wealth is held equal across the two agents

¹⁰To be precise, the plot shows the log return on the wealth of the Young minus the log return to total wealth.

Figure 4 - "Uncertain Mean"-case: Risk-sharing and portfolio allocations

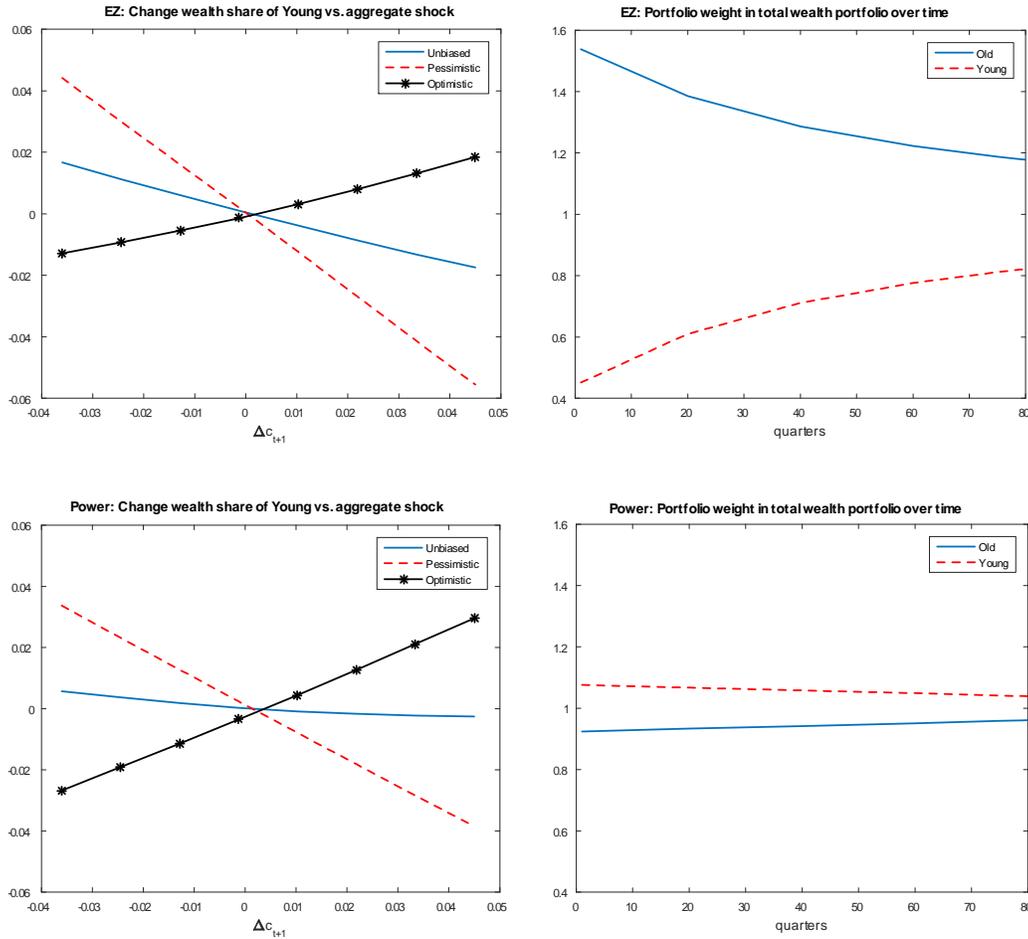


Figure 4: The left plots show the change in the wealth share of the Young for different realizations of the aggregate shock (consumption growth). The current wealth of the agents is set equal, the current age of the Young and the Old are in the middle of their generations (at 10 and 30 years, respectively), and the current beliefs of the Old are unbiased. The solid line shows the change in the relative wealth share when the current beliefs of the Young are also unbiased, whereas the red dashed line shows the case where the Young are pessimistic (the belief of the mean growth rate 2 standard deviations below the true mean), and the dash-dotted line shows the case where the Young are optimistic (2 standard deviations above the true mean)). The right plots show the portfolio allocation of the Young and the Old agent over time (from 1 to 80 quarters), where beliefs are held unbiased and the wealth-share is held equal across agents. The top plots show the result from the 'EZ' case whereas the bottom plots show the result from the 'Power' case.

and the beliefs of both agents are assumed to be unbiased over time. With recursive utility, the Old start with a portfolio allocation of about 1.5 (150%) to the total wealth portfolio, while the Young starts with 0.5. Subsequently both are pulled towards 1 as the difference in the dispersion of beliefs decreases over time. This is because with Bayesian learning, as can be

seen from Equation (3), the variance of beliefs decreases more rapidly when prior uncertainty is high than low. Right before the generational shift, there is still a substantial difference, about 1.2 versus 0.8. This implication of the model is consistent with microevidence on portfolio holdings. In particular, Lucas and Heaton (2000) show that risky assets as a fraction of agents' total wealth (including human capital) is strongly increasing with age.

With power utility preferences, however, the portfolio choice is markedly different. First, portfolio weights barely budge over time and they are quite close, about 1.1 versus 0.9. Second, it is the *Young* who are more exposed to the total wealth fluctuations and thus have a higher portfolio weight. This somewhat counter-intuitive result is due to the fact that total wealth covaries *positively* with marginal utility in this case due to the low level of the elasticity of substitution. For instance, an upward update in the mean belief of the growth rate, due to a positive consumption shock, *lowers* the price/consumption ratio as the wealth effect dominates (as now $\psi = 1/\gamma < 1$). This effect is strong enough to make the return to total wealth positively related to marginal utility. The Young still perceive model risk as higher than the Old (remember, the subjective consumption dynamics in the 'Uncertain mean'-case are $\Delta c_{t+1} = m_{i,t} + \sqrt{1 + A_{i,t}\sigma}\varepsilon_{t+1}$, so this follows since $A_{Young,t} > A_{Old,t}$), but given the negative correlation between total wealth returns and aggregate consumption growth and the resulting negative risk premium on the total wealth portfolio, the Old hedges the Young by holding *less* of the total wealth claim.

The reason the portfolio share does not move much over time (holding beliefs and wealth constant) in the power utility case is because model uncertainty is simply not very important for total welfare when agents are indifferent to the timing of resolution of uncertainty. Thus, while the Young experience more model uncertainty relative to the Old, neither CRRA agent cares very much about it. In the recursive utility case however, the agents experience large utility losses from being faced with model uncertainty. In particular, the amount of long-run risk is proportional to the size of the update in mean beliefs, which from Equation (2) is $A_t\sigma$. In the middle of their respective generations, the relative difference in perception of short run risk ($\sqrt{1 + A_t}\sigma$) between the Young and the Old is 0.3%, while the relative difference in long-run risk ($A_t\sigma$) is 67%. Thus, with recursive utility there is a much bigger difference in perceived risk between the two agents, which is also why both the dynamics of portfolio allocation and asset prices (as we show in the next Section) are much more pronounced in this case.

3.1.2 'Uncertain Probability'-case

In the case where the probability of a Depression is not known to agents, the differences in portfolio allocation can be illustrated in terms of agents' different wealth exposure to the Depression event. Therefore, we plot the log change in the relative wealth of the Young

Figure 5 - "Uncertain Probability"-case: Risk-sharing

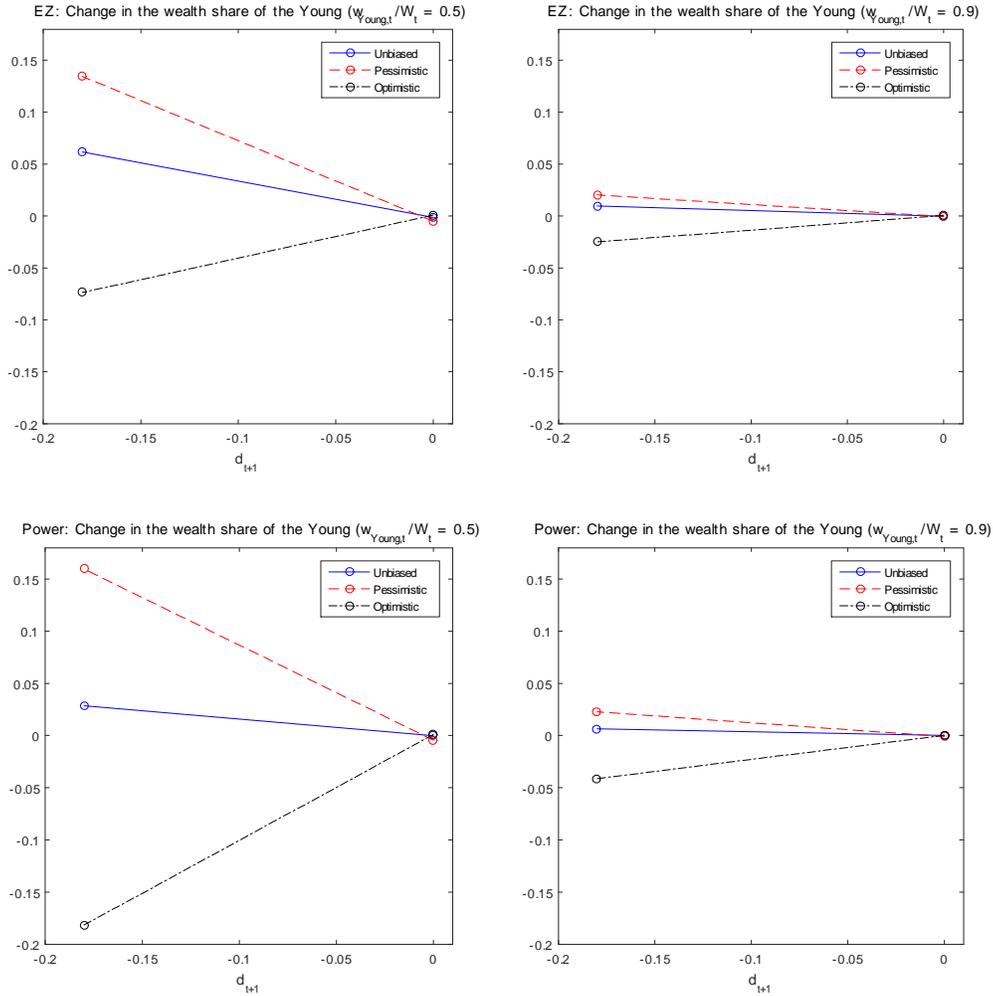


Figure 5: The left plots show the change in the wealth share of the Young for different realizations of the Depression shock (0% or -18%). The current wealth of the agents is set equal, the current age of the Young and the Old are in the middle of their generations (at 10 and 30 years, respectively), and the current beliefs of the Old are unbiased. The solid line shows the change in the relative wealth share when the current beliefs of the Young are also unbiased, whereas the red dashed line shows the case where the Young are pessimistic (the belief of the mean growth rate 2 standard deviations below the true mean), and the dash-dotted line shows the case where the Young are optimistic (2 standard deviations above the true mean)). The right plots show the same but when Young holds 90% of the wealth in the economy. The top plots show the result from the 'EZ' case whereas the bottom plots show the result from the 'Power' case.

(starting from a 50% wealth share) in the relevant cases $d_{t+1} = 0$ and $d_{t+1} = -18\%$. The left plots in Figure 5 show the case where the agents have the same level of wealth before the shock for the EZ calibration (top plot) and the power utility calibration (lower plot), where the Old

have unbiased beliefs ($E_t^{Old}[p] = p$). When the Young also are unbiased, the Old still provide insurance against the Depression event and, as in the previous case, more so with Epstein-Zin utility than with power utility. Also as in the previous case, when the Young are optimistic ($E_t^{Young}[p] = 0.0001 < p = 0.00425$) they provide insurance to the unbiased Old, while the reverse holds when the Young are pessimistic ($E_t^{Young}[p] = 0.02 > p = 0.00425$). Notably, with Epstein-Zin utility, the trade between optimists and pessimists is substantially tempered relative to the power utility case as the optimists now also face parameter uncertainty risk. We revisit this finding in detail in the section discussing the asset pricing implications of wealth dynamics and risk sharing.

The right-hand plots show the same for the case when the wealth share of the Young is 90%. As might be expected, the Old provide less insurance to the Young in this case.

3.2 Asset pricing implications

Parameter learning induces subjective long-run consumption risks because the conditional posterior distribution of future consumption growth varies in a very persistent manner as agents' update their beliefs. With Epstein-Zin preferences and a preference for early resolution of uncertainty ($\gamma > 1/\psi$) agents are averse to long-run risks (see Bansal and Yaron (2004)). As a result learning can be a tremendous amplifier of macro shocks' impact on marginal utility with such preferences and indeed, help resolve asset pricing puzzles even in a traditional representative agent asset pricing model (CJL 2015).

The same amplification mechanism is at work here, but, importantly, there is (a) speculation and risk-sharing across generations related to the model uncertainty and (b) learning persists indefinitely. The former arises as agents that differ in their assessment of probabilities of future states will trade with each other to take advantage of what they perceive to be the erroneous beliefs of other agents. That is, bad states that are perceived as less likely from the perspective of agent A relative to that of agent B are states for which agent B buys insurance from agent A and vice versa, thus making each agent better off given their beliefs. These effects serve to decrease the risk premium and undo some of the asset pricing effects of the long-run risks that arise endogenously from parameter learning. However, the latter element of the 'this time is different'-bias works in the opposite direction in that the magnitude of belief updates from parameter learning remains high indefinitely and so agents are permanently faced with a substantial degree of long-run risk induced by model uncertainty. Further, because learning persists indefinitely, we do not expect to observe any 'drifts' in risk-premia in long time series.

Table 2 - Unconditional Moments: Learning about the Mean

Table 2: This table gives average sample moments from 1,000 simulations of 328 quarters from the "Uncertain mean"-calibration. The columns labelled "EZ" correspond to the calibrations given in Table 1. The "Power" columns have the EIS parameter, ψ , set such that agents have power utility. the column labelled "Known mean" corresponds to moments from the benchmark case where agents know the true mean growth rate of the economy and therefore are not subject to the 'This Time is Different'-bias. The column "One agent" corresponds to the case where there is only one dynasty in the economy and therefore no effects of risk-sharing across agents. Lower case of variables denote the natural log of their upper case counterparts. SR_T denotes the Sharpe ratio, and the T subscript denotes sample mean. The 'data' column gives the empirical moments from U.S. data from 1929 to 2011.

	Data	'This Time is Different'		Known mean	One agent
		<i>EZ</i> : $\gamma = 9$ $\psi = 1.5$	<i>Power</i> : $\gamma=9$ $\psi = 1/9$	<i>EZ</i> : $\gamma = 9$ $\psi = 1.5$	<i>EZ</i> : $\gamma = 9$ $\psi = 1.5$
	1929 – 2011	$\beta = 0.994$	$\beta = 0.994$	$\beta = 0.994$	$\beta = 0.994$
$E_T [r_m - r_f]$	5.0	5.1	0.18	1.5	5.7
$\sigma_T [r_m - r_f]$	20.2	17.6	10.5	12.7	16.7
$SR_T [R_m - R_f]$	0.35	0.37	0.07	0.12	0.42
$\sigma_T (pd)$	0.44	0.20	0.16	0	0.18
$\rho_T (pd)$	0.88	0.91	0.88	<i>n/a</i>	0.91
$E_T [r_f]$	0.7	2.2	15.4	3.4	2.1
$\sigma_T [\ln M]$	-	0.55	0.22	0.24	0.55
$\gamma \times \sigma_T [\Delta c]$	-	0.24	0.24	0.24	0.24
$E_T [\Delta c]$	1.8%	1.8%	1.8%	1.8%	1.8%
$\sigma_T [\Delta c^{TA}]$	2.2%	2.2%	2.2%	2.2%	2.2%

3.2.1 Unconditional Moments

Table 2 shows the unconditional moments from 1,000 simulations of length 328 quarters (as in the data sample) for the "Uncertain mean"-case. The simulations have a 2000 period burn-in period to avoid the effects of initial conditions. The table gives standard moments from four versions of the model—two cases where agents suffer from the 'This Time is Different'-bias, with Epstein-Zin and power utility preferences, respectively, the benchmark case where agents know the mean parameter, and the one-agent case where there is only one dynasty and therefore no effects of risk sharing.¹¹ Table 3 shows the corresponding results for the

¹¹For the power utility specifications, it was necessary to tighten the truncation bounds somewhat to ensure finite utility. Since consumption growth is i.i.d., Sharpe ratios and risk premiums are the same in the power utility and Epstein-Zin economies in the known parameter case and we therefore only present for the EZ known parameter case.

Table 3 - Unconditional Moments: Learning about Disaster Probability

Table 3: This table gives average sample moments from 1,000 simulations of 328 quarters from the "Uncertain probability"-calibration. The columns labelled "EZ" correspond to the calibrations given in Table 1. The "Power" columns have the EIS parameter, ψ , set such that agents have power utility. the column labelled "Known prob." corresponds to moments from the benchmark case where agents know the true probability of a Depression event and therefore are not subject to the 'This Time is Different'-bias. The column "One agent" corresponds to the case where there is only one dynasty in the economy and therefore no effects of risk-sharing across agents. Lower case of variables denote the natural log of their upper case counterparts. SR_T denotes the Sharpe ratio, and the T subscript denotes sample mean. The 'data' column gives the empirical moments from U.S. data from 1929 to 2011.

	Data	'This Time is Different'		Known prob.	One agent
		<i>EZ</i> : $\gamma = 4.5$ $\psi = 1.5$ $\beta = 0.994$	<i>Power</i> : $\gamma=4.5$ $\psi = 1/4.5$ $\beta = 0.994$	<i>EZ</i> : $\gamma = 4.5$ $\psi = 1.5$ $\beta = 0.994$	<i>EZ</i> : $\gamma = 4.5$ $\psi = 1.5$ $\beta = 0.994$
	1929 – 2011				
$E_T [r_m - r_f]$	5.0	4.6	0.26	0.35	4.7
$\sigma_T [r_m - r_f]$	20.2	16.8	13.2	13.1	16.2
$SR_T [R_m - R_f]$	0.35	0.42	0.09	0.10	0.43
$\sigma_T (pd)$	0.44	0.31	0.04	0	0.31
$\rho_T (pd)$	0.88	0.92	0.90	<i>n/a</i>	0.93
$E_T [r_f]$	0.7	1.9	9.6	3.2	1.8
$\sigma_T [\ln M]$	-	0.68	0.20	0.12	0.68
$\gamma \times \sigma_T [\Delta c]$	-	0.12	0.12	0.12	0.12
$E_T [\Delta c]$	1.8%	1.8%	1.8%	1.8%	1.8%
$\sigma_T [\Delta c^{TA}]$	2.2%	2.2%	2.2%	2.2%	2.2%

"Uncertain probability"-case.

Introducing the 'This Time is Different'-bias has strong implications for asset pricing, when agents have a preference for early resolution of uncertainty. Indeed, the equity risk premium increases by a factor more than 3, while the Sharpe ratio and price of risk increases by a factor more than 2, relative to the known parameters cases, matching the data well. Instead, when agents have power utility, the risk premium and Sharpe ratio on equity *decreases*. This is because a 'positive' update in beliefs decreases the price-dividend ratio and therefore the absolute value of the covariance between the pricing kernel and the equity return.

In addition, the heterogeneity in beliefs generated by the 'This time is different'-bias, further reduces unconditional Sharpe ratios and risk premiums. For the power utility case with the 'This Time is Different'-bias and two agents, the unconditional price of risk, measured as the mean conditional volatility of the log pricing kernel, drops slightly from 0.24 in the benchmark i.i.d. case to 0.22. We note that with power utility and only one agent, the price

of risk would have actually increased slightly because of the small increase in risk due to learning. Thus, the drop in the power utility case is driven by the presence of optimists who improve risk-sharing in the market which decreases the required average risk compensation.

Instead, in the Epstein-Zin case with learning, the price of risk increases to 0.68. Thus, the model uncertainty risk channel dominates with these preferences. This is underscored by noting that the unconditional sample moments corresponding to the 'one EZ agent'-cases in these tables are similar to that of the two-EZ-agent economies. Thus, in contrast to the findings of Chen, Joslin and Tran (2012), risk-sharing does not strongly affect the unconditional sample moments in an economy with EZ preferences and learning — a feature we discuss more below.

3.2.2 Wealth dynamics and risk pricing

With heterogenous beliefs the relative wealth of agents in the economy becomes an additional state variable. In our setting, this endogenous state variable affects asset prices in two distinct ways. First, consistent with previous literature, if there are optimists (pessimists) in the market, risk premiums and Sharpe ratios are lower (higher) under the objective measure. Second, and particular to the belief heterogeneity in the generational 'This Time is Different' model, the heterogeneity in agents' confidence in their beliefs matters. As discussed earlier, this uncertainty has particularly strong effects on asset prices when agents have Epstein-Zin utility. Thus, even if agents' mean beliefs coincide, an increase in, say, the wealth of the Young affects asset prices as the wealth-weighted beliefs are now more uncertain.

Figure 6 shows the annualized equity risk premium (under the objective measure) for the 'Unknown probability'-case plotted against the consumption share of the Young—i.e., the agent that perceives more model uncertainty. As before, the agents are in the middle of their respective generations in terms of their age. The dashed blue lines correspond to the case where the agents' mean beliefs about the Depression probability coincide (and are set approximately equal to the true value). The bottom right plot shows the main case where beliefs are uncertain (due to the 'This Time is Different'-bias) and agents have Epstein-Zin preferences. Here the equity risk premium increases from 4% when all the wealth are in the hands of the Old to over 10% when the Young hold all the wealth. The consumption share varies between 0.3 and 0.7 in simulations, so in practice the range of variation in the risk premium resulting from variation in the relative wealth of agents is from 5% to 8%. The bottom left plot shows the same economy but where agents have power utility preferences. Note that the blue line here barely budges, as in this case the model uncertainty has a very low risk price. That is, the aforementioned source of wealth-dynamics-induced variation in the price of risk is absent when agents have power utility.

Another source of variation in the risk premium associated with relative wealth shocks

Figure 6 - Wealth dynamics and the risk premium

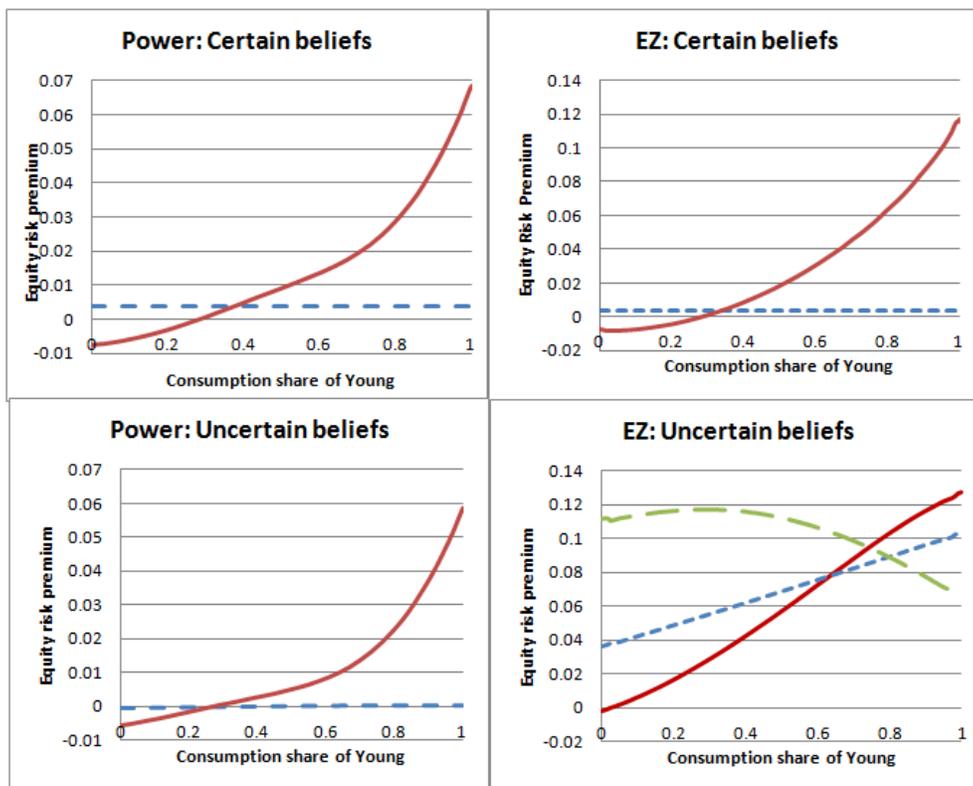


Figure 6: The figure shows the conditional annualized equity risk premium versus the current relative consumption of the Young, for the case with an uncertain Depression probability. Both the Old and the Young are assumed to be in the middle of their respective generations. "Power" refers to models where the agents have power utility, whereas "EZ" refers to models where agents have Epstein-Zin utility with $\gamma > 1/\psi$. The dashed lines correspond to the case where the mean beliefs of the agents coincide. The solid line corresponds to the case where the Young have pessimistic mean beliefs, while the Old have optimistic mean beliefs. The green line with the long dashes in the bottom right graph corresponds to the case where instead the Young are optimists and the Old are pessimists.

is, as mentioned, the difference in mean beliefs. The solid red lines show the case where the Young agent is pessimistic ($E_t^{Young}[p] = 0.02$) and the Old agent is optimistic ($E_t^{Old}[p] = 0.001$). For both the power utility and the EZ cases, the risk premium is increasing the more wealth are given to the pessimistic agent (in this example, the Young). However, note that the sensitivity of the risk premium to the amount of consumption of the optimistic agent is much higher for the power utility case when the wealth of the optimistic agent is low. In particular, the risk premium decreases from 6% to close to zero when the relative consumption share of optimists goes from 0 to 0.5. This echoes the finding of Chen, Joslin, and Tran (2012), who document that in terms of the pricing of disaster risk, the risk premium decreases precipitously when (a small mass, e.g., 10-20% of total wealth) of optimists are introduced in

the model. Thus, the disaster model does not seem robust to a reasonable amount of belief heterogeneity. In the EZ case, however, the same change in the consumption share of the optimists (the Old in this example), decreases the risk premium by only a fraction of 2, from about 12% to about 6%.

In other words, the strong nonlinearity found in the power utility case is not present when agents have EZ preferences and face model uncertainty. Note that in both cases agents are learning, subject to the 'This Time is Different'-bias. In fact, the two top plots shows the same graphs for the case where agents do not learn and are perfectly certain about their (different) beliefs, as in the model of Chen, Joslin, and Tran (2012). Here, each agent remains an optimist or pessimist forever, with no updating of beliefs. Since there are no model-uncertainty-induced long-run risks in this case, both the EZ and the power utility cases exhibit the strong nonlinearity and fragility of the disaster model with respect to belief heterogeneity. In sum, allowing for uncertain beliefs, learning, and Epstein-Zin preferences renders the disaster model more robust along this dimension.

These two sources of excess return predictability, as well as the common movements in beliefs, also generate a positive relation (for the EZ model) between the dividend yield and future excess returns, as in the data. We discuss the model's implications for return predictability in more detail later in the paper.

3.2.3 Over- and undervaluation vs beliefs

The model features periods of over- and undervaluation, as agents at times become either too optimistic (after a sequence of positive Normal shocks or a lack of Depression shocks) or too pessimistic (after a sequence of negative Normal shocks or a recent Depression shock) relative to the true model. Given that agents' beliefs are close to rational—agents update as Bayesian during their lifetime and therefore have mean beliefs close to the truth—one may think the asset misvaluation must be quite small. This is not the case. The reason is that while the belief updates are quite 'small,' they are very persistent. Thus, they affect valuations strongly. In this section, we analyze the dynamics of mispricing and the implications for the conditional equity risk premium and risk prices.

'Uncertain mean'-case The left half of Panel A of Table 4 gives the population dynamics of the log price-dividend ratio. As mentioned earlier, the time-variation in agents beliefs about the mean parameter, μ , cause substantial variation in the price level. This variation is also highly persistent with an annual population autocorrelation of 0.96. The right half of Panel A shows that the persistence of the price-dividend ratio is inherited from the persistence

**Table 4 - Dynamics of Mispricing:
Learning about the mean**

Table 4: This table gives population moments related to the asset pricing impact of agents' boundedly rational learning. The left half of Panel A shows the annual volatility ($\sigma[\cdot]$) and autocorrelation ($\rho[\cdot]$) of the log price-dividend ratio, as well as the average amount of equity market mispricing as a fraction of total market value. The latter is measured as the average absolute value of deviations in the price-dividend ratio away from the unconditional mean normalized by the unconditional mean. The right half of Panel A shows the annualized volatility and autocorrelation of beliefs about the growth rate averaged across the two agents. Panel B shows similar moments for the equity risk premium and the effective risk aversion, $\tilde{\gamma}_t$. The latter is the conditional risk aversion an econometrician with an infinite sample would assign to the representative agent, as explain in the main text.

Panel A:		Dynamic Behavior of:	
The Price Level		Beliefs	
$\sigma [\ln P/D]$	0.25	$\sigma [\bar{m}_i]$	0.1%
$\rho [\ln P/D]$	0.96	$\rho [\bar{m}_i]$	0.96
$E \left[\frac{\text{abs}(P/D - E[P/D])}{E[P/D]} \right]$	20.2%		

Panel B:		Dynamic Behavior of:	
The Risk Premium		Effective Risk Aversion	
$E [E_t [r_{m,t+1}^e]]$	5.1%	$E [\tilde{\gamma}_t]$	20.4
$\sigma [E_t [r_{m,t+1}^e]]$	3.4%	$\sigma [\tilde{\gamma}_t]$	7.3
$\rho [E_t [r_{m,t+1}^e]]$	0.84	$\rho [\tilde{\gamma}_t]$	0.89

in the average mean belief of the two agents, which is also 0.96 due to the generational nature of the learning bias.

The standard deviation of beliefs about the quarterly growth rate is small (0.1%), but the mispricings these belief fluctuations induce are substantial. In particular, the mispricing of the market is typically $\pm 20\%$, as measured by the average absolute deviation of the price-dividend ratio from its unconditional average. Thus, in a long sample (say, 100 years) one can expect occurrences of market levels from 0.5 to 2 times the average price level, implying substantial variation in the long-horizon risk premium under the objective measure.

The left half of Panel B of Table 4 gives the dynamics of the equity premium: its annualized volatility is 3.4% (i.e., the two standard error bounds of annualized quarterly expected returns are -1.7% and 11.9% , with an annual autocorrelation of 0.84). The dividend yield and the risk premium are counter-cyclical as investors tend to be pessimistic (optimistic) in bad

(good) times, when recent shocks to economic outcomes have been negative (positive).

One implication of the model presented here with an experiential learning bias, is that an econometrician who uses a long sample to estimate (the dynamics of) risk aversion jointly with consumption dynamics and assumes Rational Expectations will conclude that investors exhibit high and time-varying risk aversion. In particular, with a long sample the econometrician will conclude that consumption growth is i.i.d. and estimate the mean and volatility parameters to be equal to their true values. The (conditional) price of risk in the i.i.d. economy, when agents are assumed to know the true model, is approximately $\gamma \times \sigma$ —risk aversion times the quantity of risk. Define the estimated conditional relative risk aversion as:

$$\tilde{\gamma}_t \equiv \sigma_t (\ln M_{t+1}) / \sigma. \tag{17}$$

Note that the conditional volatility of the pricing kernel, $\sigma_t (\ln M_{t+1})$, is approximately the maximum conditional Sharpe ratio in the 'This Time is Different'-economy, under the objective measure, and that σ is the objective volatility of consumption growth.

The right half of Panel B in Table 4 shows that the average level of the 'estimated conditional risk aversion' is about 20, more than twice as high as the actual risk aversion of the representative agent (recall, $\gamma = 9$ in the 'Uncertain mean'-model). Further, the estimated conditional risk aversion, $\tilde{\gamma}_t$, has a volatility of about 7 and an annual autocorrelation of 0.89. Thus, an econometrician would find strong, counter-cyclical time-variation in the pricing of macro risks due to the belief fluctuations (it is counter-cyclical as beliefs about μ are procyclical). In fact, the dynamics of this 'estimated risk aversion' is reminiscent of an external habit formation model, such as Campbell and Cochrane (1999), and due to the small but persistent mistakes agents make in their belief formation as well as the fact that the model uncertainty is priced when agents have Epstein-Zin preferences.

Cash Flow vs. Discount Rate Shocks. In his presidential address, Cochrane (2011) argues that historically all variation in price-dividend ratios correspond to variation in discount rates and none to variation in expected cash flows. In the models we consider here, however, the variation in the price-dividend ratios is instead mainly due to variation in agents' expectations of future cash flows as agents update beliefs about the unknown parameters.

Under the objective measure, however, dividend growth is i.i.d., and so there is in fact *no* variation in expected cash flows under this measure. Thus, historical analysis of the relationship between the price-dividend ratio and future returns will attribute all variation in the price-dividend ratio to discount rate variation. In this sense, this paper highlights how even statistically 'small' departures from Rational Expectations can lead to a radically different view of the fundamental drivers of price variation.

Figure 3.2.3 - Predictability Regressions

Table 5: This table gives average results from aggregate log excess return and log dividend growth forecasting regressions in the data and from 1,000 simulations of 328 quarters from the "Uncertain mean"-calibration. The table shows results for forecasting horizons of 1, 3, and 5 years, where the dependent variable is future log excess market returns in Panel A and log dividend growth in Panel B. The independent variable is the appropriately lagged aggregate log dividend yield. 't-stat' denotes Newey-West corrected t-statistics for the coefficient on the dividend yield, while R_{adj}^2 denotes the adjusted R-squared of the regression. The number of lags in the Newey-West procedure is 1.5 times the forecasting horizon (e.g., 18 for a 3-year horizon with quarterly overlapping observations). The data columns correspond to corresponding regressions presented in Beeler and Campbell (2012).

Panel A: Returns					
Forecasting horizon ($\sum_{j=1}^s r_{t+j}^e$)	Data		Model		
	coeff. (t-stat)	R_{adj}^2	coeff. (t-stat)	R_{adj}^2	%-tile of data R_{adj}^2
1 year	-0.09* (-1.8)	0.04	-0.21* (-1.8)	0.03	0.70
3 years	-0.26*** (-3.2)	0.17	-0.58** (-2.2)	0.09	0.87
5 years	-0.41*** (-3.8)	0.26	-0.92*** (-2.7)	0.15	0.87

Panel B: Dividends					
Forecasting horizon ($\sum_{j=1}^s \Delta d_{t+j}$)	Data		Model		
	coeff. (t-stat)	R_{adj}^2	coeff. (t-stat)	R_{adj}^2	%-tile of data R_{adj}^2
1 year	0.07** (2.0)	0.09	0.02 (0.3)	0.01	0.99
3 years	0.11 (1.3)	0.06	-0.12 (-0.59)	0.03	0.81
5 years	0.09 (1.2)	0.04	-0.25 (-0.90)	0.06	0.51

Table 3.2.3 shows how standard forecasting regressions perform within the 'Uncertain mean'-model versus the data. We do not show the intercepts from the regressions for brevity. Panel A shows results from log excess return forecasting regressions, while Panel B shows results for log real dividend growth forecasting regressions. In all cases, the lagged log Price-Dividend ratio is the predictor variable. The data columns, which are taken from Beeler and Campbell (2012) and cover the long U.S. sample, show that the price-dividend ratio is a significant predictor of future excess returns, especially for long forecasting horizons (>1 year). The price-dividend ratio only predicts dividend growth at the annual forecasting horizon, however, and not over long horizons.

The right half of Table 3.2.3 shows the corresponding regressions using simulated data from the model. In particular, we show the mean coefficient, t-statistics (Newey-West, with

lags set to 1.5 time the forecasting horizon), and R_{adj}^2 's from 1,000 simulated samples of the same length as the data sample. As in the data, the price-dividend ratio is a significant predictor of future excess returns over long forecasting horizons within the model. The mean R_{adj}^2 's are not as high as in the data, as might be expected as the model for simplicity ignores business cycle risks and focuses only on the generational learning bias. The rightmost column shows the percentile that the R_{adj}^2 found in the data corresponds to within the model. Over the 3- and 5-year forecasting horizon, the data R^2 correspond to the 87th-percentile of that from model simulations. Panel B shows that dividend growth is unpredictable within the model, as is natural given the i.i.d. dividend growth assumption. In sum, the model features substantial variation in the market price-dividend ratio, which is all attributed to discount rate variation when running standard forecasting regressions.

Table 6 - Beliefs vs. Returns

Table 6: The table gives the correlation between the conditional expected return on the dividend claim with lagged returns on the dividend claim. The column denoted "Average subjective" gives the average conditional subjective expected return of the agents in the economy, averaged across agents. The column denoted "Objective" gives the conditional expected return using the true model parameters (i.e., under the objective, P-measure). The conditional expected return is always the next quarters' expected return, while the lagged return is measured over different backward-looking horizons (s).

Lagged return horizon ($\sum_{j=0}^s r_{t-j}$)	Correlation between conditional expected returns and lagged returns	
	Average subjective	Objective
1 quarter ($s = 1$)	0.07	-0.08
1 year ($s = 4$)	0.13	-0.15
3 years ($s = 12$)	0.22	-0.26
5 years ($s = 20$)	0.30	-0.32

Beliefs vs. Lagged Returns. As discussed in Greenwood and Shleifer (2014), survey evidence shows a positive relation between lagged stock returns and (some) investors' forward-looking expected returns. Table 6 gives the correlation between conditional expected returns and lagged returns on the dividend claim in the 'Uncertain mean' model. In particular, the column "Average subjective" gives the correlation between next quarter's expected subjective

returns, averaged across agents of all combination of ages, with increasing backward-looking windows of realized returns. As the table shows, this positive relation is present in the “This Time is Different”-model, both in gross returns and in excess returns. This is notable, as in the model investors are learning about fundamentals and do not use returns directly in their belief formation (see Barberis, Greenwood, Lin, and Shleifer (2014) for an analysis of an economy where investors explicitly extrapolate from lagged stock returns). Thus, the relation between returns and beliefs is fully endogenous to the model. To understand this outcome, consider a sequence of positive shocks which yields positive stock returns and increases investors’ mean beliefs about the growth rate of the economy. This would by itself not change risk premiums. However, since growth rates are uncertain, the higher long-run growth rate makes the dividend claim more sensitive to future growth rate shocks. For intuition, consider the Gordon growth formula $P/D = 1/(r - \mu)$. The price-dividend ratio is more sensitive to shocks to growth rates when growth rates are high. Since shocks to growth rates are priced risks, the increased covariance of returns with shocks to growth rates leads to a higher subjective risk premium.

In reality, though, investors update ‘too much’ and so on average high growth rate expectations are too optimistic and prices in fact mean-revert. This is shown in column “Objective” in Table 6, which gives the correlation between lagged stock returns and conditional expected returns using the true probabilities to calculate the expectation. This correlation is negative, as one would expect with mean-reverting price-dividend ratios and given the true, constant growth rate in the economy.

‘Uncertain probability’-case For the "Uncertainty probability"-model, where agents learn about the probability of a Depression shock, the implications for over- and undervaluation are in many ways similar to the "Uncertain mean"-case. The top left plot of Figure 7 shows the mean beliefs about the (annualized) probability of a Depression for the two Dynasties, where both agents start at $t = 1$ with unbiased beliefs. In this simulation, the Depression shock occurs at time $t = 41$, and there are no Depressions thereafter until $t = 160$ (i.e., the time period considered covers the full life of an agent). The agent from Dynasty A was again a new Young at $t = 1$, while the agent from Dynasty B was a new Old. Thus, at the time of the shock, they are both in the middle of their respective generations. The Young updates more quickly in the direction of a low Depression probability as long as no Depression shock occurs. Once the shock hits, it is again the Young that update more and their mean beliefs flip from being the optimists in the market to being the pessimists. At $t = 80$, the Old die (from Dynasty B) and the previous Young become the new Old (from Dynasty A), which is why the mean belief of Dynasty B has a kink at this point in time (the new Young starts updating more quickly).

Figure 7 - "Uncertain Probability": Simulated paths

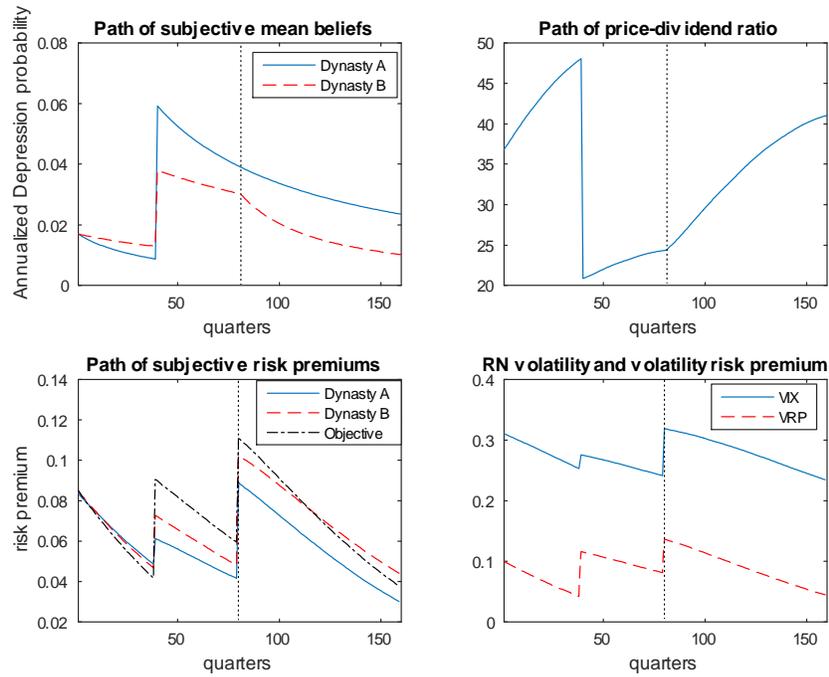


Figure 7: The figure shows selected time series statistics from the 'Uncertain probability'-model. The relevant shocks are Depression realizations. We simulate a path for one full life-time (160 quarters) and let the Depression occur only in the 41st quarter in order to illustrate the model dynamics. The Young generation is assumed to come alive in the first quarter (Dynasty A), while the initial Old generation (Dynasty B) becomes the new Young in the 81st quarter (the dotted, vertical line in each plot). The subjective mean beliefs in the top left plot refer to the mean belief about the probability of a Depression event. The risk premiums and volatilities all refer to the aggregate dividend claim.

At $t > 80$, there is only one of the two generations in the market that has experienced the Depression personally (Dynasty A). This agent assigns a higher probability to the Depression state than the Young who have not experienced the Depression, as documented empirically in Malmendier and Nagel (2011). The upper right plot of Figure 7 shows the model-implied price-dividend ratio. First, the fluctuations are large, from about 25 to about 45. Further, it takes about 120 quarters (30 years) after the Depression shock for the price-dividend ratio to reach its level when beliefs are unbiased, so the effects of the shock are long-lasting. Again, the benchmark known parameters i.i.d. consumption growth model has a constant price-dividend ratio, so these fluctuations reflect misvaluation, as well as the priced model uncertainty.

The bottom left plot of Figure 7 shows three versions of the conditional risk premium—the objective risk premium, the risk premium using agent A's beliefs, and the risk premium using agent B's beliefs. First, note that the risk premium falls in all three cases as long as a

Depression does not occur. Once a Depression occurs, the risk premium goes up in all three cases as agents update the likelihood of a severe consumption drop. However, the conditional risk premium goes up the least for the Young, who become the pessimists in the market upon experiencing the shock as they update their beliefs the most. Consider the time $t = 100$ in the plot. This is after the Depression shock and after a new generation was born (at $t = 81$). Thus, here we can consider the 'Depression babies' effect on expected market returns. The agent (Dynasty A) who experienced the Depression event and is still alive expects a risk premium that is about 1.5% points below that of the current Young (Dynasty B) who did not experience the Depression. Thus, consistent with Malmendier and Nagel's findings, the generation that experienced the Depression have a relatively lower allocation to stocks given their relatively low expectation of future excess stock market returns.

The bottom right plot in Figure 7 shows the path of the VIX (here, risk-neutral annualized quarterly conditional equity return volatility) and the variance risk premium (VRP: the 'VIX' minus the actual (objective measure) conditional annualized quarterly equity return volatility). The Depression event at $t = 41$ is associated with an increase in the VIX as the subjective beliefs about the likelihood of a disaster increases. Interestingly, the variance risk premium increases more. This is due to the fact that the Depression shock causes a wealth-transfer from the Old to the Young. Thus, the wealth-weighted perceived model uncertainty increases and claims that pay off in the Depression state (such as a variance swap) become more valuable hedges, leading to a higher variance risk premium.

3.3 The Empirical Relation between the Price-Dividend Ratio and the fraction of Young vs. Old

A robust implication of the model is that asset prices are more sensitive to macro shocks when the Young control more wealth in the economy, since these agents update beliefs more strongly in response to macro shocks. Further, the price level should on average be lower in this case, as the young perceive more model risk and therefore on average require higher returns.

In lieu of data on the aggregate wealth (including human capital) of the old versus the young, we here use demographic data to proxy for this ratio. In particular, we use Census data from 1900 to 2013 to calculate the log ratio of the number of people in the 25-44 year age bracket versus the 45+ age bracket. We obtain the annual real price-dividend ratio from the Shiller data and annual real, per capita GDP growth from the NIPA tables.

First, as a description of the raw data, we relate the level of the ratio of young to old to the aggregate price-dividend ratio. Panel A of Table 7 shows a regression of the log price-dividend ratio (the *pd*-ratio) on the log ratio of young to old (the *yo*-ratio from here on).

Table 7

The Empirical Relation between the P/D-ratio and Fraction of Young vs. Old

Table 7: The table shows various regressions relating the log price-dividend ratio (pd) to the log ratio of the population of Young (25-44 years) to Old (45+ years), yo . The sample is annual from 1900 to 2013. Δ denotes an annual difference and Δ^{10} denotes a 10-year difference. In Panels A and B, t-statistics are computed using Newey-West standard errors with 20 lags. In Panel C, t-statistics are computed using White standard errors. The regression with header "Time-Trend" uses the deviations from the full-sample time trend of yo , "40-year difference" uses the difference between yo_t and yo_{t-40} , while "40-year moving average" uses the difference between yo_t and the lagged 40-year moving average of yo . * denotes significance at the 10%-level, ** denotes significance at the 5%-level, *** denotes significance at the 1%-level.

Panel A:					
Regression: $pd_t = \alpha + \beta \times yo_t + \varepsilon_t$					
	α		β		R_{adj}^2
Coefficient	3.58***		-1.08***		41%
(t-stat)	(20.9)		(-3.13)		
Panel B:					
Regression: $\Delta^{10}pd_t = \alpha + \beta \times \Delta^{10}yo_t + \varepsilon_t$					
	α		β		R_{adj}^2
Coefficient	-0.002		-1.02**		7.3%
(t-stat)	(-0.02)		(-2.05)		
Panel C:					
Regression: $\Delta pd_{t+1} = \alpha_0 + \alpha_1 yo_t^{\text{detrended}} + \beta_0 \Delta gdp_{t+1} + \beta_1 \Delta gdp_{t+1} \times yo_t^{\text{detrended}} + \varepsilon_{t+1}$					
	α_0	α_1	β_0	β_1	R_{adj}^2
<i>Time-Trend:</i>					
Coefficient	-0.04*	-0.40	0.82***	8.30**	14%
(t-stat)	(-1.93)	(-1.34)	(3.26)	(2.06)	
<i>40-yr difference:</i>					
Coefficient	-0.07	-0.21	3.19**	8.32**	5.9%
(t-stat)	(-1.38)	(-1.51)	(2.18)	(2.21)	
<i>40-yr moving average:</i>					
Coefficient	-0.13***	-0.47**	2.46***	9.49***	18%
(t-stat)	(-3.41)	(2.04)	(3.93)	(2.75)	

Over the sample, the pd -ratio is trending up, while the yo -ratio is trending down, yielding a strong negative relation (an R_{adj}^2 of 41%) that admittedly may be spurious given the high persistence of the series. Given the high persistence, we also look at 10-year changes in the pd -ratio versus 10-year changes in the yo -ratio, in annual overlapping observations regressions

in Panel B. Here, the R_{adj}^2 is 7% and the relation is negative and significant at the 5%-level using Newey-West standard errors and 20 lags, again suggesting the relation between the amount of wealth controlled by the young relative to the old is indeed a factor in determining the level of asset prices.

In Panel C of Table 7 we test whether asset prices are indeed more sensitive to macro shocks when the yo -ratio is high. Here, we can run annual regressions of changes in the pd -ratio, so these regressions have more power and should be better behaved. In particular, Panel C shows results from the regression:

$$\Delta pd_{t+1} = \alpha_0 + \alpha_1 yo_t^{\text{detrended}} + \beta_0 \Delta gdp_{t+1} + \beta_1 yo_t^{\text{detrended}} \times \Delta gdp_{t+1} + \varepsilon_{t+1}, \quad (18)$$

where $yo_t^{\text{detrended}}$ are detrended versions of the yo -ratio. We use three different detrending methods. The first is simply to remove the full-sample time-trend from the raw yo -ratio. The other two are motivated by the model, and we normalize the ratio at a generational frequency. In particular, the second version of $yo_t^{\text{detrended}}$ is simply the difference between the current yo_t -ratio and that 40 years ago, while the third is the difference between the current yo_t -ratio and the current 40-year lagged moving average of the yo_t -ratio.

All three versions yield positive and significant β_0 - and β_1 -coefficients. Thus, the conditional sensitivity of the pd -ratio to GDP shocks, $\beta_t \equiv \beta_0 + \beta_1 yo_t^{\text{detrended}}$ is indeed higher when the fraction of young vs. old is high, as predicted by the model.

4 Conclusion

We have proposed a relatively simple, but quantitatively realistic model that incorporates the 'this time is different'-bias across generations documented by Malmendier and Nagel (2011, 2013). Importantly, consistent with Ang, Bekaert, and Wei (2007) agents' beliefs are very close to truth in a likelihood ratio sense. Thus, the calibration of the bias is not excessive.

In this framework, model-uncertainty persists indefinitely and is an added risk factor due to the combination of learning and a preference for early resolution of uncertainty. This learning-induced long-run risk helps the model match standard asset pricing moments despite the heterogeneity in beliefs. This is because when agents have recursive preferences and are uncertain about their beliefs, the amount of speculative activity stemming from disagreement in mean beliefs is tempered. Agents that are optimistic about the Depression probability refrain from hedging pessimistic agents against this state as they know they will adversely update their beliefs if the state occurs, which will significantly lower their continuation utility.

Despite the small departure (in a statistical sense) from the rational expectations assumption, the conditional asset pricing implications that arise from the experiential learning

bias are large. In particular, the model generates persistent periods of significant over- and underpricing. These price fluctuations are related to the small, persistent mistakes in cash flow expectations on the part of the agents in the economy. An econometrician would, however, attribute almost all of the variation in the price-dividend ratio to discount rate shocks and time-varying risk prices, consistent with the data (see, e.g., Campbell (1991)). That is, the joint behavior of consumption and asset prices is reminiscent of external habit formation models such as Campbell and Cochrane (1999), where risk prices are time-varying and cash flows are i.i.d., even though at the individual level it is time varying expected cash flow risk and a preference for the early resolution of uncertainty that drive consumption and portfolio choices.

References

- Ang, A., G. Bekaert, and M. Wei (2007), "Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better?," *Journal of Monetary Economics* 54, 1163-1212.
- Backus, D., B. Routledge, and S. Zin (2009), "Who Holds Risky Assets?," working paper Stern School of Business.
- Baker, S., B. Hollifield, and E. Osambela (2014), "Disagreement, Speculation, and Aggregate Investment," Carnegie Mellon University working paper.
- Bansal, R. and I. Shaliastovich (2010), "Confidence Risk and Asset Prices," *American Economic Review*, papers and proceedings, 100(2): 537—41.
- Bansal, R. and A. Yaron (2004), "Risks for the Long-Run: A Potential Resolution of Asset Pricing Puzzles", *Journal of Finance* 59(4), 1481 - 1509
- Barberis, N., R. Greenwood, L. Jin, and A. Shleifer (2014), "X-CAPM: An extrapolative capital asset pricing model," forthcoming *Journal of Financial Economics*.
- Bhamra, H. and R. Uppal (2014), "Asset Prices with Heterogeneity in Preferences and Beliefs," *The Review of Financial Studies* 27, 519-580.
- Bloom, N. (2009), "The impact of uncertainty shocks," *Econometrica* 77, 623 - 685.
- Borovicka, J. (2013), "Survival and long-run dynamics with heterogeneous beliefs under recursive preferences", NYU working paper.
- Choi, J. and T. Mertens (2013), "Extrapolative Expectations and the Equity Premium," NYU and Yale working paper.
- Campbell, J. (1991), "A Variance Decomposition for Stock Returns," *The Economic Journal*, 101, 157-179.
- Campbell, J. and J. Cochrane (1999), "By Force of Habit: A Consumption-based Explanation for Aggregate Stock Market Behavior," *Journal of Political Economy*, 107, 205-251
- Cochrane, J. (1991), "Production-Based Asset Pricing and the Link Between Stock Returns and Economic Fluctuations." *Journal of Finance* 46, 209-237.
- Collin-Dufresne, P., M. Johannes, and L. A. Lochstoer (2013a), "Parameter Learning in General Equilibrium: The Asset Pricing Implications", Columbia University working paper.

- Collin-Dufresne, P., M. Johannes, and L. A. Lochstoer (2013b), "A Robust Numerical Procedure for Solving Risk-Sharing Problems with Recursive Preferences", Columbia University working paper.
- Detemple, J. (1986), "Asset pricing in a production economy with incomplete information." *Journal of Finance*, 41, 383–390.
- Doob, J. L. (1949), "Application of the theory of martingales." Coll. Int. du C. N. R. S. Paris, 23 - 27.
- Dumas, B., A. Kurshev, and R. Uppal (2009), "Equilibrium Portfolio Strategies in the Presence of Sentiment Risk and Excess Volatility," *Journal of Finance*, 64, 579 - 629.
- Ehling, P., A. Graniero, and C. Heyerdahl-Larsen (2013), "Asset Prices and Portfolio Choice with Learning from Experience," London Business School working paper.
- Epstein, L. and S. Zin (1989), "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework", *Econometrica* 57, 937 - 969
- Gabaix, X. (2012), "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance," *Quarterly Journal of Economics* 127, 645 - 700.
- Garleanu, N. and S. Panageas (2012), "Young, Old, Conservative and Bold. The implication of finite lives and heterogeneity for asset pricing." UC Berkeley and Chicago Booth working paper.
- Gennotte, G. (1986), "Optimal Portfolio Choice under Incomplete Information," *Journal of Finance*, 61, 733-749.
- Geweke, J. (2001), "A note on some limitations of CRRA utility," *Economics Letters* 71, 341 - 345.
- Goyal A. and I. Welch, "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction," July 2008, *Review of Financial Studies* 21(4), 1455-1508.
- Greenwood, R. and A. Shleifer (2014), "Expectations of returns and expected returns," *Review of Financial Studies* 27, 714-746.
- Hall, R. E (1988), "Intertemporal Substitution in Consumption," *Journal of Political Economy* 96, 339 – 357.
- Harrison J. and D. Kreps (1978), "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations," *Quarterly Journal of Economics*, 92, 323 - 336.

- Hansen, L. (2007), "Beliefs, Doubts and Learning: Valuing Macroeconomic Risk," Richard T. Ely Lecture, *The American Economic Review* 97, No. 2., 1 - 30
- Hansen, L. and T. Sargent (2001), "Robust control and model uncertainty," *American Economic Review* 91, 60-66.
- Heaton, J. and D. Lucas (2000), "Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk," *Journal of Finance* 55, 1163-1198.
- Hirshleifer, D. and J. Yu (2013), "Asset Pricing in Production Economics with Extrapolative Expectations", UC Irvine and University of Minnesota working paper.
- Jermann, U. (1998), "Asset Pricing in Production Economies," *Journal of Monetary Economics* 41, 257 - 275.
- Johannes, M., L. A. Lochstoer, and Y. Mou (2013), "Learning about Consumption Dynamics," forthcoming at the *Journal of Finance*.
- Johnson, T. (2007), "Optimal Learning and New Technology Bubbles," *Journal of Monetary Economics*, 87, 2486 – 2511.
- Kan, R. (1995), "Structure of Pareto Optima When Agents Have Stochastic Recursive Preferences," *Journal of Economic Theory* 66, 626 - 631.
- Kaltenbrunner, G. and L. A. Lochstoer (2010), "Long-Run Risk through Consumption Smoothing," *Review of Financial Studies* 23, 3141 - 3189.
- Kreps, D. and E. Porteus (1978), "Temporal Resolution of Uncertainty and Dynamic Choice Theory", *Econometrica* 46, 185 – 200.
- Lewellen, J. and J. Shanken (2002), "Learning, Asset-Pricing Tests, and Market Efficiency," *Journal of Finance* (57(3)), 1113 - 1145.
- Lucas, R. (1978), "Asset Prices in an Exchange Economy," *Econometrica* 46, 1429-1446
- Lucas, R. and N. Stokey (1984), "Optimal Growth with Many Consumers," *Journal of Economic Theory* 32, 139 - 171.
- Nakov, A. and G. Nuno (2015), "Learning from Experience in the Stock Market," *Journal of Economic Dynamics & Control* 52, 224 - 239.
- Pastor, L. (2000), "Portfolio Selection and Asset Pricing Models," *Journal of Finance* 55, 179 - 223

Pastor, L. and P. Veronesi (2006), "Was there a NASDAQ bubble in the late 1990s?," *Journal of Financial Economics*, 81, 61 – 100.

Pastor, L. and P. Veronesi (2009), "Learning in Financial Markets," *Annual Review of Financial Economics*.

Timmermann, A. (1993), "How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices." *Quarterly Journal of Economics*, 1993, 108, 1135-1145.

Veronesi, P. (2000), "How does information quality affect stock returns?" *Journal of Finance*, 55.

Working, H. (1960), "Note on the Correlation of First Differences of Averages in a Random Chain," *Econometrica* 28(4), 916 - 918.

5 Appendix – Model Solution

We here briefly describe how we solve the model. Denote aggregate consumption C_t . There are two agents with Epstein-Zin preferences and different beliefs about the exogenous aggregate consumption dynamics. The resource constraint is:

$$C_t = C_{A,t} + C_{B,t}. \quad (19)$$

The preferences of agents A and B are given by:

$$V_{A,t} = V_A(C_{A,t}, V_{A,t+1}) = \left[(1 - \beta) C_{A,t}^\rho + \beta E_t^A (V_{A,t+1}^\alpha)^{\rho/\alpha} \right]^{1/\rho}, \quad (20)$$

$$V_{B,t} = V_B(C_{B,t}, V_{B,t+1}) = \left[(1 - \beta) C_{B,t}^\rho + \beta E_t^B (V_{B,t+1}^\alpha)^{\rho/\alpha} \right]^{1/\rho}, \quad (21)$$

where $E_t^i[\cdot]$ denotes an expectation taken with respect to agent i 's beliefs.

With complete markets, the Pareto problem can be written:

$$\max_{\{C_{A,t}, C_{B,t}\}_{t=0}^{\infty}} \lambda V_{A,0} + (1 - \lambda) V_{B,0} \quad \text{s.t.} \quad C_{A,t} + C_{B,t} = C_t \text{ for all states and time.} \quad (22)$$

Even though the individual utility functions are recursive, the social planner function is not recursive. However, there exists a recursive formulation (see Lucas and Stokey (1984), Kan (1995), Backus, Routledge and Zin (2009)):

$$\begin{aligned} J(C_t, V_{B,t}) &= \max_{C_{A,t}, V_{B,t+1}} \left[(1 - \beta) C_{A,t}^\rho + \beta E_t^A [J(C_{t+1}, V_{B,t+1})^\alpha]^{\rho/\alpha} \right]^{1/\rho} \\ \text{s.t. } V_{B,t} &\geq V_B(C_t - C_{A,t}, V_{B,t+1}). \end{aligned} \quad (23)$$

Note that the values $V_{B,t+1}$ we are maximizing over in this problem are for all possible states of nature that can occur at $t + 1$. Thus, we are solving for the consumption-share of agent A and promised utility for agent B for each possible state over the next period. Since preferences are monotonic, the utility-promise constraint will bind and with optimized values we have $V_{A,t} = J(C_t, V_{B,t})$ and $V_{B,t} = V_B(C_t - C_{A,t}, V_{B,t+1})$.

The first order and envelope conditions for the maximization problem imply that for each state, the marginal intertemporal rates of substitution of the two agents must be equal:

$$\pi_{A,t}(\omega_{t+1})\beta\left(\frac{C_{A,t+1}}{C_{A,t}}\right)^{\rho-1}\left(\frac{V_{A,t+1}}{\mu_{A,t}[V_{A,t+1}]}\right)^{\alpha-\rho} = \pi_{B,t}(\omega_{t+1})\beta\left(\frac{C_{B,t+1}}{C_{B,t}}\right)^{\rho-1}\left(\frac{V_{B,t+1}}{\mu_{B,t}[V_{B,t+1}]}\right)^{\alpha-\rho}, \quad (24)$$

where $\pi_{i,t}(\omega_{t+1})$ is agent i 's conditional probability assessment of state ω_{t+1} being realized next period, determined by agent's current beliefs as summarized in the posteriors from the learning problem. Equation (24) is of course a familiar requirement for equilibrium in a frictionless complete markets economy and, as usual, the agents' probability measures must be equivalent measures for this condition to hold.

The problem with solving the recursion in Equation (23) is that the evolution equation for the endogenous state variable (relative wealth, or relative consumption, of the two agents) is not known. In particular, for power utility preferences, where $\alpha = \rho$, Equation (24) provides analytically the evolution equation for the relative consumption of the two agents as a function of the aggregate state. With $\alpha \neq \rho$, however, this is no longer the case, as the value functions of the agents appear and since these value functions are unknown (they are, in fact, what we are trying to solve for). One can start with a guess for the value functions as a function of the aggregate state variables and then try to apply Equation (24) to solve for the consumption sharing rule, but this is highly unstable as one of the state variables is the relative consumption share. One typically needs to effectively guess the equilibrium value functions as the initial value functions in order for the recursion to be well-behaved. In other words, while the recursion in Equation (23) technically provides a value function iteration solution to the risk-sharing problem, it is very hard to implement in practice.

We instead solve the model using the approach outlined in Collin-Dufresne, Johannes, and Lochstoer (2013b). Here, we suggest a new numerical approach to solving these types of problems. The approach relies on two steps. Step 1: solve, in a backwards recursion, the risk-sharing problem in a T -period economy. Step 2: increase T until value functions of both agents no longer change (i.e., has converged according to a convergence criteria) and verify, using the recursion in Equation (23) that the solution indeed corresponds to the solution to the infinite horizon problem. The latter is done by iterating on the recursion given in Equation (23) using both the value functions and the evolution dynamics for the endogenous state variable as obtained in the backwards recursion solution to the T -period problem where

T is large. Note that we do not need the economy to be stationary. I.e., there could be degenerate wealth dynamics. This should be clear from the following discussion, where we outline the approach in detail.

It is convenient to solve a normalized version of this model, where all variables are divided by aggregate consumption. Let lower case of variables denote the normalized counterpart. Thus, for an arbitrary variable Z_t we have that $z_t = Z_t/C_t$. In this case, the value functions can be written:

$$v_{i,t} = \left[(1 - \beta) c_{i,t}^\rho + \beta \mu_{i,t}^\rho \right]^{1/\rho}, \quad (25)$$

where $\mu_{i,t} \equiv E_t^i \left[v_{i,t+1}^\alpha (C_{t+1}/C_t)^\alpha \right]^{1/\alpha}$ and where the resource constraint is $c_{A,t} + c_{B,t} = 1$. The stochastic discount factor under agent i 's probability measure can then be written (see Epstein and Zin (1989)):

$$\begin{aligned} E_t^i \left[M_{t+1}^i R_{t+1}^j \right] &= 1 \text{ for all } t \text{ and } j \\ M_{t+1}^i &= \beta \left(\frac{c_{i,t+1}}{c_{i,t}} \right)^{\rho-1} \left(\frac{C_{t+1}}{C_t} \right)^{\alpha-1} \left(\frac{v_{i,t+1}}{\mu_{i,t}} \right)^{\alpha-\rho}. \end{aligned} \quad (26)$$

The aggregate state variables in this economy are the ones governing the agents' subjective consumption dynamics ($x_t = [m_{A,t}, m_{B,t}, a_{A,t}, a_{B,t}, T - t]^l$) and the relative wealth of the two agents. Since the relative consumption of agents is monotone in the relative wealth of agents, we use the relative consumption of agent A as the endogenous state variable, $c_{A,t}$. Thus, $v_{i,t} = f_i(x_t, c_{A,t})$.

If the endogenous evolution equation of $c_{A,t}$ is known, we can now easily solve a standard value function iteration problem on a grid for x_t and $c_{A,t} \in (0, 1)$:

$$f_A(x_t, c_{A,t}) = \left[(1 - \beta) c_{A,t}^\rho + \beta E_t^A [f_A(x_{t+1}, c_{A,t+1})^\alpha (C_{t+1}/C_t)^\alpha]^{\rho/\alpha} \right]^{1/\rho}, \quad (27)$$

$$f_B(x_t, c_{A,t}) = \left[(1 - \beta) (1 - c_{A,t})^\rho + \beta E_t^B [f_B(x_{t+1}, c_{A,t+1})^\alpha (C_{t+1}/C_t)^\alpha]^{\rho/\alpha} \right]^{1/\rho}. \quad (28)$$

Thus, the crux of the risk-sharing problem is finding the endogenous evolution equation for c_A for all points in the state-space.

Solving the model at $T - 1$:

At time T , when the economy ends, the value functions reduce to:

$$v_{A,T} = (1 - \beta)^{1/\rho} c_{A,T}, \quad (29)$$

$$v_{B,T} = (1 - \beta)^{1/\rho} (1 - c_{A,T}). \quad (30)$$

Equations (29) and (30) give the boundary conditions for the value functions as a function of

the relative consumption of agent A , $c_{A,t}$. As mentioned earlier, it is convenient to use $c_{A,t}$ as the endogenous state-variable (one could equivalently use relative wealth of, say, agent A), in addition to the exogenous state variables, x_t .

At time $T - 1$, the complete markets requirement that agents IMRS are equalized across states implies that:

$$\begin{aligned} \pi_{T|T-1}^A \beta \left(\frac{c_{A,T}}{c_{A,T-1}} \right)^{\rho-1} \left(\frac{C_T}{C_{T-1}} \right)^{\alpha-1} \left(\frac{v_{A,T}}{\mu_{A,T-1} (v_{A,T} C_T / C_{T-1})} \right)^{\alpha-\rho} &= \dots \\ \pi_{T|T-1}^B \beta \left(\frac{1-c_{A,T}}{1-c_{A,T-1}} \right)^{\rho-1} \left(\frac{C_T}{C_{T-1}} \right)^{\alpha-1} \left(\frac{v_{B,T}}{\mu_{B,T-1} (v_{B,T} C_T / C_{T-1})} \right)^{\alpha-\rho} &. \end{aligned} \quad (31)$$

Here $\pi_{T|T-1}^i$ denotes the probability agent i assigns to a given state at time T given agent i 's beliefs at time $T - 1$. First, define:

$$\begin{aligned} k_{T-1} &= \frac{\mu_{B,T-1} (v_{B,T} C_T / C_{T-1})^{\rho-\alpha}}{\mu_{A,T-1} (v_{A,T} C_T / C_{T-1})^{\rho-\alpha}} \\ &= \frac{E_{T-1}^B ((1-c_{A,T})^\alpha (C_T / C_{T-1})^\alpha)^{\rho/\alpha-1}}{E_{T-1}^A (c_{A,T}^\alpha (C_T / C_{T-1})^\alpha)^{\rho/\alpha-1}}, \end{aligned} \quad (32)$$

where the dependence of k_{T-1} on the current state variables in the economy is implicit. Next, imposing the boundary values as given in Equations (29) and (30), we have that:

$$\begin{aligned} \pi_{T|T-1}^A \beta \left(\frac{c_{A,T}}{c_{A,T-1}} \right)^{\rho-1} \left(\frac{C_T}{C_{T-1}} \right)^{\alpha-1} c_{A,T}^{\alpha-\rho} &= \dots \\ k_{T-1} \pi_{T|T-1}^B \beta \left(\frac{1-c_{A,T}}{1-c_{A,T-1}} \right)^{\rho-1} \left(\frac{C_T}{C_{T-1}} \right)^{\alpha-1} (1-c_{A,T})^{\alpha-\rho} & \\ \Downarrow & \\ \left(\frac{c_{A,T}}{1-c_{A,T}} \right)^{\alpha-1} &= k_{T-1} \frac{\pi_{T|T-1}^B}{\pi_{T|T-1}^A} \left(\frac{c_{A,T-1}}{1-c_{A,T-1}} \right)^{\rho-1}. \end{aligned} \quad (33)$$

Note first that Equation (33) implies that, for a given state of the world at time T and value of state variables at time $T - 1$, $c_{A,T} \in (0, 1)$ is decreasing in k_{T-1} (since $\alpha - 1 < 0$). Thus, for a given k_{T-1} , Equation (33) uniquely determines $c_{A,T}$ for each state of the world at time T and a, say, higher k_{T-1} implies that $c_{A,T}$ is lower in each state of the world. Of course, from Equation (32), we only know k_{T-1} as a function of $c_{A,T}$. However, the right-hand side of Equation (32) is monotone in $c_{A,T}$, which means it is monotone in k_{T-1} . In particular, since a lower k_{T-1} means that $c_{A,T}$ is higher in each state of the world,

$E_{T-1}^A \left(c_{A,T}^\alpha (C_T/C_{T-1})^\alpha \right)$ is decreasing (increasing) in k_{T-1} if $\alpha > 0$ ($\alpha < 0$). The opposite relation holds for $E_{T-1}^B \left((1 - c_{A,T})^\alpha (C_T/C_{T-1})^\alpha \right)$. Since, k_{T-1} is the ratio of these two expectations (taken to the power of $\rho/\alpha - 1$), we have that the right hand side of Equation (32) is indeed monotone in k_{T-1} . In other words, Equations (32) and (33) provide unique solutions for k_{T-1} , and thus for $c_{A,T}$ for each state of the world at time T . It is also immediate from these equations that a solution exists where $c_{A,T-1} \in (0, 1)$ and $k_{T-1} > 0$.

While the fixed point problem for finding k_{T-1} implicit in Equations (32) and (33) must be solved numerically for each point on a grid for the state variables, this is very fast given the monotonicity (e.g., a routine like *zbrent* works very fast). For a particular choice of α , one can solve analytically for $c_{A,T}$ as a function of k_{T-1} and the state variables at $T - 1$. In sum, using $c_{A,T-1}$ as the endogenous state variable, we now have numerically the conditional evolution equation for c_A , from $T - 1$ to time T .

Next, we can now solve numerically for the normalized value functions at time $T - 1$ on a grid for the relevant state variables at time $T - 1$, using:

$$v_{i,t} = \left[(1 - \beta) c_{i,t}^\rho + \beta E_t^i \left[v_{i,t+1}^\alpha (C_{t+1}/C_t)^\alpha \right]^{\rho/\alpha} \right]^{1/\rho}, \quad (34)$$

where the state variables are x_t and $c_{A,t}$. Note that solving numerically for the certainty equivalent of next period's value function requires the use of the evolution equation for c_A .

The second backwards iteration is then at time $t = T - 2$. Again, we start with the requirement that the IMRS is equalized for each state for the two agents:

$$\begin{aligned} \pi_{t+1|t}^A \beta \left(\frac{c_{A,t+1}}{c_{A,t}} \right)^{\rho-1} \left(\frac{C_{t+1}}{C_t} \right)^{\alpha-1} \left(\frac{v_{A,t+1}}{\mu_{A,t} (v_{A,t+1} C_{t+1}/C_t)} \right)^{\alpha-\rho} &= \dots \\ \pi_{t+1|t}^B \beta \left(\frac{1 - c_{A,t+1}}{1 - c_{A,t}} \right)^{\rho-1} \left(\frac{C_{t+1}}{C_t} \right)^{\alpha-1} \left(\frac{v_{B,t+1}}{\mu_{B,t} (v_{B,t+1} C_{t+1}/C_t)} \right)^{\alpha-\rho} &. \end{aligned}$$

Note that $v_{i,t+1} = f_i(x_{t+1}, c_{A,t+1})$ is known from the previous step in the backwards recursion. Now, rewrite the above equation as:

$$\left(\frac{c_{A,t+1}}{1 - c_{A,t+1}} \right)^{\rho-1} \left(\frac{v_{A,t+1}}{v_{B,t+1}} \right)^{\alpha-\rho} = k_t \frac{\pi_{t+1|t}^B}{\pi_{t+1|t}^A} \left(\frac{c_{A,t}}{1 - c_{A,t}} \right)^{\rho-1}, \quad (35)$$

where $k_t = \frac{\mu_{B,t} (v_{B,t+1} C_{t+1}/C_t)^{\rho-\alpha}}{\mu_{A,t} (v_{A,t+1} C_{t+1}/C_t)^{\rho-\alpha}}$. It is clear that the left hand side of Equation (35) is monotone in k_t . Since $\frac{v_{A,t+1}}{v_{B,t+1}}$ is known as a function of x_{t+1} and $c_{A,t+1}$, it is easy to numerically find the value of $c_{A,t+1}$ corresponding to a particular outcome (x_t, ε_{t+1}) , given a value for k_t . It is clear from Equation (35) that the value of $c_{A,t+1}$ given k_t is unique when $\alpha - \rho < 0$, which

is the relevant case in our calibrations. In particular, $v_{A,t+1}/v_{B,t+1}$ is obviously increasing in $c_{A,t+1}$ given the state x_{t+1} , and $c_{A,t+1}/(1 - c_{A,t+1})$ is trivially increasing in $c_{A,t+1}$. Since both $\rho - 1 < 0$ and $\alpha - \rho < 0$, we have that the left hand side of Equation (35) is decreasing in k_t for all states x_{t+1} .

Finally, we need to solve for k_t . Note that the previous equation gives the evolution equation for $c_{A,t+1}$ as a function also of k_t ($c_{A,t+1} = g(X_t, c_{A,t}, k_t, \varepsilon_{t+1})$). Thus, we again find k_t as a fixed point of the equation:

$$k_t = \frac{E_t^B [(v_B(x_{t+1}, 1 - c_{A,t+1}(k_t)))^\alpha (C_{t+1}/C_t)^{\alpha\rho/\alpha-1}]}{E_t^A [(v_A(x_{t+1}, c_{A,t+1}(k_t)))^\alpha (C_{t+1}/C_t)^{\alpha\rho/\alpha-1}]}. \quad (36)$$

Again, Equations (35) and (36) provide a unique solution to the evolution equation for $c_{A,t}$ to $c_{A,t+1}$ as a function of the aggregate exogenous state variables at x_t and x_{t+1} , using the corresponding logic as that for k_{T-1} . The normalized value function at time t can then be found using Equation (34).

Further backwards recursions follow the same algorithm as that given for the case $t = T - 2$. In practice, we find that having $T > 2400$ is sufficient for convergence of typical calibrations (clearly, the time-discount factor β is particularly important in this regard). Note that we do not impose nondegenerate wealth dynamics as the relative wealth of agents implicitly is a state variable (we just chose the relative consumption for convenience).