

Asset Prices and Portfolio Choice with Learning from Experience *

Paul Ehling[†]
BI

Alessandro Graniero[‡]
LBS

Christian Heyerdahl-Larsen[§]
LBS

Current Draft: February 2016

*We would like to thank Suleyman Basak, Peter Bossaerts (FIRN Asset Pricing Group Meeting presenter), Bradyn Breon-Drish (22nd Mitsui Finance Symposium discussant), João Cocco, Bernard Dumas, Engelbert Dockner (WFA discussant), Nicolae Gârleanu (Econometric Society discussant), Francisco Gomes, Ralph Koijen, Anna Pavlova, H el ene Rey, Raman Uppal, Johan Walden, Yajun Xiao (FIRN Asset Pricing Group Meeting discussant), Hongjun Yan (ESSFM discussant), and participants at a LBS Brown Bag, a Stochastic Analysis Seminar at the University of Oslo, a CAPR lunch seminar at BI, Wharton, 22nd Mitsui Finance Symposium, WFA 2014, ESSFM 2014, EEA 2014, FIRN Asset Pricing Group Meeting, Conference on Behavioral Aspects in Macroeconomics and Finance 2014 in Milan, and the 2015 North American Winter Meeting of the Econometric Society for helpful comments and suggestions. Paul Ehling thanks the Centre for Asset Pricing Research (CAPR) at BI for funding support.

[†]Department of Finance, BI Norwegian Business School, Nydalsveien 37, 0484 Oslo, paul.ehling@bi.no

[‡]London Business School, Regent's Park, London, NW1 4SA, agraniero.phd2010@london.edu

[§]London Business School, Regent's Park, London, NW1 4SA, cheyerdahllarsen@london.edu

Abstract

We study asset prices and portfolio choice with overlapping generations where the young disregard history to learn from own experience. Disregarding history implies less precise estimates of consumption growth, which, in equilibrium, leads the young to increase their investment in risky assets after positive returns or act as trend chasers and to lose wealth to the old. Consistent with findings from survey data, the average belief about expected returns in the economy relates negatively to future realized returns and is smoother than objective expected returns. Having especially bad experiences early on in life, cause persistent disagreement and tilt portfolio weights.

Keywords: Learning from Experience, Persistent Disagreement, Survey Based versus Objective Expected Returns, Trend Chasing

JEL Classification: G10, G11, G12, E2

1 Introduction

Recent empirical evidence suggest that experience matters for the formation of beliefs. If this is the case, then personal experiences with stock markets and the macroeconomy in general should manifest through savings and investment decisions and impact asset prices. Therefore, what are, broadly speaking, the implications of experience driven financial decisions? More specifically, how do young agents with short life experience share consumption risk with the older agents with longer life experience? How do differences in experience impact the wealth distribution through trading in financial securities? How do asset prices change when more optimistic young or pessimistic young trade with experienced agents?

In response to these questions, we study an overlapping generations economy with learning in which cohort specific experience drives beliefs about consumption growth and through that affects equilibrium outcomes. The main findings are that old agents have more precise estimates than younger agents and their beliefs react less to news, the young act as trend chasers, wealth shifts from young to old, the market price of risk is countercyclical and the risk-free rate of return is procyclical, the average belief about the stock return is negatively correlated with and less volatile than the actual expected return, and experiencing especially bad outcomes when young leads to significant and persistent disagreement in the economy and to a persistent tilt in portfolio weights.

For experience to matter in the model, we endow young agents with a prior variance about expected consumption growth that is greater than zero. Yet, once born, agents use Bayes rule to update their beliefs about growth. As a consequence, young agents respond more to news than older and more experienced agents whose posterior variance has declined over time. In the model, we see that when consumption booms, new generations become optimists through learning. We also see the opposite, when consumption declines period after period then new generations learn to be pessimists. The experience driven learning has important implications for the portfolio choice of agents. As young agents revise their expectations about consumption growth more in response to news, they become relatively

more optimistic after a positive shock. In equilibrium, this translates into increased demand for the risky asset. Therefore, they act optimally as trend chasers, given their beliefs.

In equilibrium, the market price of risk and the risk-free rate are given by standard formulas from an economy with complete information plus a correction term that captures the difference between the “market view,” as measured by the consumption share weighted average belief, and the true growth rate. The “market view” increases after a positive shock to consumption, as all agents in the economy revise their expectations upwards due to learning. This has implications for the real interest rate and market price of risk. Specifically, in times when the “market view” is higher than the actual growth, we see that the real interest rate is higher and the market price of risk is lower relative to the complete information benchmark. In addition, as the “market view” increases after a positive shock, we see a negative relation between consumption growth and the market price of risk or, put differently, the market price of risk increases in downturns.

As young agents are trend chasers, they increase their exposure to the risky asset after a positive shock. At the same time, the actual market price of risk declines. Hence, we see that young agents tend to buy the risky asset when the expected return is low. This has implications for the relative performance of young versus old in our economy. That is, young and inexperienced investors underperform relative to old or experienced investors because their estimates about growth vary more. We show that the welfare cost of learning from experience is between 72% to 275% of the welfare cost of aggregate fluctuations when varying the time discount factor between 0.001 and 0.02. Moreover, the fact that young agents underperform is broadly consistent with the evidence in Agarwal, Driscoll, Gabaix, and Laibson (2009). First, Agarwal, Driscoll, Gabaix, and Laibson (2009) use a proprietary database to provide evidence for the hypothesis that older adults make fewer financial mistakes than younger adults.¹ Second, Agarwal, Driscoll, Gabaix, and Laibson (2009) show

¹More specifically, Agarwal, Driscoll, Gabaix, and Laibson (2009) conclude that financial mistakes follow a U-shaped pattern. Our analysis, however, abstracts from cognitive impairment that plagues about half of the population between ages 80 and 89.

that older and experienced investors have greater investment knowledge. Third, the survey based analysis in Arrondel, Calvo-Pardo, and Tas (2014) suggests that investors' measure of information "increases with past experience."

According to Malmendier and Nagel (2011), individuals who have experienced high stock market or bond market returns are more likely to take on further financial risks, i.e., are more likely to participate in the stock market or bond market, and allocate a higher proportion of their liquid assets to stocks or bonds. Malmendier and Nagel (2011) also provide empirical evidence that point to the importance of experience for beliefs. According to their view, experience effects could be the result of attempts to learn from experiences where all available historical data is used by not entirely trusted.² Further, in a follow up paper, Malmendier and Nagel (2014), show that individuals adapt their inflation forecasts to new data but overweight inflation realized during their life-times. They also show that young individuals update expectations more strongly in the direction of recent surprises than the old. Moreover, learning from experience explains substantial disagreement between young and old individuals in periods of high surprise inflation. Such empirical motivated differences in expectations between young and old is the departure point for our theoretical work where we study the role of experience on asset prices, consumption, and portfolio choice in a consumption based overlapping generations economy. Specifically, we show that experiencing a series of negative or positive shocks to consumption early on in life impacts the beliefs of the young strongly. What is surprising in our artificial booms or depressions is that after 30 years only half of the effect on beliefs fades away. Therefore, the model showcases how especially good or bad experiences early on in life can be a source of persistent disagreement. It also shows that such persistent disagreement can have a long lasting impact on the optimal portfolio choice of agents, consistently with, for example, the evidence in Knüpfer, Rantapuska, and Sarvimäki (2014), who show that labor market shocks can have a long-lasting impact on portfolio choice. In addition, we show that particularly good or bad experiences in one mar-

²Using age as measure of managers' investment experience, Greenwood and Nagel (2009) show that young managers trend-chase in their technology stock investments, while old managers do not.

ket significantly influences the portfolio policy in that market with little spillover effects to the allocations to other risky assets. This is consistent with Malmendier and Nagel (2011), who show that investors who experience bad stock market returns invest less in stocks, but do not significantly alter their investment in bonds.

Using a host of data sources to improve on previous work, Greenwood and Shleifer (2014) show that survey based measures of expectations are positively correlated with past stock returns, but negatively correlated with future returns. The negative correlation suggests that when survey respondents expect high returns, then the realized future returns will be low. Kojen, Schmeling, and Vrugt (2015) present empirical analysis that highlights how pervasive the evidence of a negative relation between survey based expectations and future returns is. Specifically, they show that survey expectations of returns negatively predict future returns in the cross-section of countries and in the time-series of global equities, currencies, and global fixed income. Motivated by Greenwood and Shleifer (2014), Barberis, Greenwood, Jin, and Shleifer (2015) propose a model with extrapolative expectations to reconcile the evidence on expectations with the evidence on volatility and predictability. Our parsimonious model complements Barberis, Greenwood, Jin, and Shleifer (2015) as it ties together extrapolative expectations with inexperience in a way that is consistent with both the evidence in Greenwood and Shleifer (2014) and in Malmendier and Nagel (2011). In our model, all agents learn optimally in a Bayesian sense, given their experience based beliefs, and equilibrium expected returns decline after a series of positive returns. Yet, surveying agents in our model produces a mean forecast that shows a positive relation with past realized returns and a negative relation with future realized returns. In addition, the model survey produces a smoother average belief than the true expected return, which is consistent with Piazzesi, Salomao, and Schneider (2015), who find that in the bond market statistical measures of the risk premium are more volatile than survey based measures of the risk premium.

The results in the paper are robust to the extension of within-cohort heterogeneity in

priors about expected endowment growth at birth. Specifically, we show how a Gaussian distribution for the priors reproduces our equilibrium results with homogeneous priors and that the presence of a fraction of agents who use all past historical information right after birth to immediately learn the true expected growth rate of consumption still reproduces our qualitative equilibrium results. The last extension we consider is the case with more than one risky security such as a nominal long-term bond exposed to the risk of inflation.

Our paper extends the literature on disagreement and asset prices by studying differences in beliefs in an overlapping generations environment. This literature, initiated by Harrison and Kreps (1978) and Detemple and Murthy (1994), among others, directly or indirectly assumes that agents learn from all available data. More recent examples of models with disagreement that employ belief structures similar to ours include Zapatero (1998), Basak (2000), Basak (2005), Dumas, Kurshev, and Uppal (2009), Xiong and Yan (2010), Cvitanic and Malamud (2011), Cvitanic, Jouini, Malamud, and Napp (2012), and Bhamra and Uppal (2014). The majority of the papers in this strand of the literature considers infinitely lived agents. This has important consequences for the stationarity of the models, since agents with more accurate beliefs accumulate wealth and eventually dominate the economy as shown in the market selection literature.³ Although in many cases the market selection process is slow as illustrated by Yan (2008) and Dumas, Kurshev, and Uppal (2009), it can be quite powerful if the state space is large and agents have access to many securities as in Fedyk, Heyerdahl-Larsen, and Walden (2013). In our model with overlapping generations, one cohort will never dominate the economy completely because new agents are continuously born and old agents die. Hence, in contrast to most of the previous literature, we have a stationary model with heterogeneous beliefs.

Our paper relates to Gârleanu and Panageas (2015), who also study a continuous-time overlapping generations economy with demographics as in Blanchard (1985). Their focus is on the implications of preference heterogeneity, instead of beliefs heterogeneity, for asset

³See, for instance, Blume and Easley (1992), Sandroni (2000), and Kogan, Ross, and Westerfield (2006).

pricing. Using recursive preferences, they show that without heterogeneity in risk aversion asset pricing moments are constant instead of countercyclical as in the data. Further, for a given amount of heterogeneity in risk aversion, heterogeneity in the elasticity of intertemporal substitution can still impact asset price moments significantly. Their model matches standard asset pricing moments.

We discuss four related papers that study the impact of an experience bias on asset prices. Perhaps the first paper in this new strand of literature is Schraeder (2015). Her discrete-time and finite-horizon economy with up to eight overlapping generations produces heterogeneity among agents and through the heterogeneity connects trading volume to volatility, excess volatility, overreaction, and price reversals, among other results.

In contemporaneous work Collin-Dufresne, Johannes, and Lochstoer (2015) consider an experience driven learning bias with two dynasties and recursive preferences. While their focus is on matching asset pricing moments, we derive implications for individual behavior that is consistent with Malmendier and Nagel (2011) and replicate the relation between survey data and future returns as in Greenwood and Shleifer (2014). In addition, in their model, due to preference for early resolution of uncertainty, young agents behave as more risk averse than the old. Hence, experience has a dual role as inexperienced agents behave as more risk averse and have less accurate beliefs. In our model, experience driven beliefs do not interplay with the effective risk aversion. Further, our model has a more general cohort and demographic structure, more general structure for priors, and multiple risky securities but still allows for closed form solutions.

Recently, Buss, Uppal, and Vilkov (2015) study the role of experience in a general equilibrium model with two risky assets where one of the risky assets is an alternative asset that is opaque and illiquid. They show that inexperienced agents initially tilt their portfolio away from the alternative asset, but eventually increase their position as they accumulate experience. Somewhat surprisingly, lower transaction costs for the alternative asset might amplify the initial portfolio tilt, as it is less costly to rebalance towards the alternative asset

when the investor becomes more experienced.

In more recent work, Malmendier, Pouzo, and Vanasco (2015) consider a discrete-time overlapping generations model with CARA utility, consumption from terminal wealth, and learning from experience. They show that stock price volatility and autocorrelations are higher when more agents rely on recent observations. Moreover, when the disagreement across generations is high, then there is higher trading volume in the stock market.

2 The Model

2.1 Demographics

We consider a continuous-time overlapping generations economy in the tradition of Blanchard (1985). Every period a fraction ν of the population dies to be replaced by newborn agents of the same mass, $\int_{-\infty}^t \nu e^{-\nu(t-s)} ds = 1$ for all t , and the time t size of the cohort born at time $s < t$ is given by $\nu e^{-\nu(t-s)} ds$.

2.2 Endowments

Agents receive endowment, $y_{s,t}$, continuously from birth at time s until death. Endowments do not depend on the time of birth, i.e., $y_{s,t} = Y_t$ for all $s \leq t$. Endowments evolve as follows:

$$dY_t = Y_t (\mu_Y dt + \sigma_Y dz_t), \tag{1}$$

where z_t is a shock modeled as a standard Brownian motion. Aggregate endowment at time t is

$$\int_{-\infty}^t \nu e^{-\nu(t-s)} y_{s,t} ds = Y_t \int_{-\infty}^t \nu e^{-\nu(t-s)} ds = Y_t, \tag{2}$$

thus, the dynamics of aggregate endowment coincide with the dynamics of individual endowments.

2.3 Information, Learning, and Disagreement

For experience to matter, we make the following assumptions regarding information, learning, and disagreement. Agents observe individual endowments, and hence aggregate endowment, but do not know the value of μ_Y .⁴ To simplify, we assume that all agents start with the correct prior for μ_Y .⁵ An agent born at time s believes the expected endowment growth is normally distributed with mean $\hat{\mu}_{s,s} = \mu_Y$ and variance $\bar{V} > 0$. Starting from birth, agents use Bayes' rule to update beliefs about expected aggregate endowment growth. By standard filtering theory, the dynamics of the expected endowment growth, $\hat{\mu}_{s,t}$, as perceived by an agent born at time s , and its posterior variance are

$$d\hat{\mu}_{s,t} = \frac{V_{s,t}}{\sigma_Y} dz_{s,t}, \quad V_{s,t} = \frac{\sigma_Y^2 \bar{V}}{\sigma_Y^2 + \bar{V}(t-s)}, \quad (3)$$

respectively, and where $z_{s,t}$ denotes a Brownian motion under the belief of an agent born at time s with associated probability P^s and information set (or sigma algebra) $\mathcal{F}_{s,t}^Y = \sigma(Y(t), s \leq t)$. The volatility coefficient, $\frac{V_{s,t}}{\sigma_Y}$, is decreasing over time because agents learn about the true parameter value, implying that young agents update more than old agents in response to a shock.

Perceived shocks relate to z_t through

$$dz_{s,t} = dz_t - \Delta_{s,t} dt, \quad (4)$$

where $\Delta_{s,t} = \frac{\hat{\mu}_{s,t} - \mu_Y}{\sigma_Y}$. The process $\Delta_{s,t}$ summarizes the standardized estimation error of agents born at time s relative to the objective probability measure.

⁴The specification of priors and optimal learning is similar to the models in Detemple and Murthy (1994) and Basak (2000).

⁵In Subsection 4.1.1 we endogenize this assumption by allowing for within cohort heterogeneity that is normally distributed and unbiased on average.

Proposition 1. *The estimation error at time t of the cohort born at time s is*

$$\Delta_{s,t} = \frac{\bar{V}(z_t - z_s)}{\sigma_Y^2 + \bar{V}(t - s)}. \quad (5)$$

Moreover, we have that $\Delta_{s,s} = 0$ and $\lim_{t \rightarrow \infty} \Delta_{s,t} = 0$.

2.4 Security Markets and Prices

Agents trade in an instantaneously risk-free asset, which is in zero net supply. Its dynamics are given by

$$dB_t = r_t B_t dt, \quad (6)$$

where r_t denotes the equilibrium real short rate.

For tractability, we employ an infinitely lived risky asset in zero net supply,⁶ which evolves according to

$$\frac{dS_t}{S_t} = (\mu_t^S dt + \sigma^S dz_t) = (\mu_{s,t}^S dt + \sigma^S dz_{s,t}), \quad (7)$$

where the last part of the equation represents the dynamics perceived by agents born at time s . In Equation (7), we have $\mu_{s,t}^S = \mu_t^S + \sigma^S \Delta_{s,t}$, in which μ_t^S is determined in equilibrium, while the volatility coefficient, σ^S , is taken as exogenous since it defines the risky security.

Annuity contracts complete the set of available securities as in Yaari (1965). They entitle to an income stream of $\nu W_{s,t}$ per unit of time. In return, the competitive insurance industry receives all financial wealth when the agent dies. Entering such a financial contract is optimal for all agents.

It is convenient to summarize the price system in terms of the stochastic discount factor. Since agents have different beliefs, they have individual stochastic discount factors. The

⁶The risky asset can be interpreted as continuously resettled contracts (e.g., futures contracts). The same asset structure is used, for example, in Basak (2000) and Karatzas, Lehoczky, and Shreve (1994). We discuss this assumption in Section 5.

dynamics of the stochastic discount factor as perceived by an agent born at time s follows

$$d\xi_{s,t} = -\xi_{s,t}(r_t dt + \theta_{s,t} dz_{s,t}), \quad (8)$$

while the dynamics of the stochastic discount factor under the actual probability measure is

$$d\xi_t = -\xi_t(r_t dt + \theta_t dz_t). \quad (9)$$

Thus, we have that the relation between the market price of risk under the objective probability measure, θ_t , and the market price of risk as perceived by the cohort born at time s , $\theta_{s,t}$, is

$$\theta_{s,t} = \theta_t + \Delta_{s,t}. \quad (10)$$

The relation between the stochastic discount factor under the objective measure and the belief of an agent born at time s is captured by the disagreement process, $\eta_{s,t}$, through the relation $\xi_t = \eta_{s,t}\xi_{s,t}$. Formally, $\eta_{s,t}$ is the Radon Nikodym derivative that allows to move from the probability measure of an agent born at time s to the actual probability measure and vice versa. The dynamics of the disagreement process, $\eta_{s,t}$, is

$$d\eta_{s,t} = \Delta_{s,t}\eta_{s,t}dz_t. \quad (11)$$

2.5 Preferences and Individual Optimization

Agents maximize lifetime utility given by

$$E_{s,s} \left[\int_s^\tau e^{-\rho(t-s)} \log(c_{s,t}) dt \right], \quad (12)$$

where τ is the stochastic time of death. In the above, the first time subscript in the expectation operator denotes under which probability measure the expectation is taken. We use the

convention that expectation operators with one time subscript are taken under the objective probability measure. By integrating out the stochastic time of death, the expected life-time utility can be written as

$$E_{s,s} \left[\int_s^\infty e^{-(\rho+\nu)(t-s)} \log(c_{s,t}) dt \right]. \quad (13)$$

The dynamics of financial wealth of an agent born at time s follows

$$dW_{s,t} = (r_t W_{s,t} + \pi_{s,t} (\mu_{s,t}^S - r_t) + \nu W_{s,t} + y_{s,t} - c_{s,t}) dt + \pi_{s,t} \sigma^S dz_{s,t}, \quad W_{s,s} = 0, \quad (14)$$

where $\pi_{s,t}$ denotes the dollar amount held in the risky asset. Since $W_{s,s} = 0$, agents are born without any financial wealth.

All agents maximize expected utility from life-time consumption, Equation (12), subject to the wealth dynamics in Equation (14).

2.6 Equilibrium

We start with defining equilibrium.

Definition 1. *Given preferences, endowments, and beliefs, equilibrium is a collection of allocations $(c_{s,t}, \pi_{s,t})$ and prices (r_t, μ_t^S) such that the processes $(c_{s,t}, \pi_{s,t})$ are optimal when agents maximize Equation (12) subject to the dynamic budget constraint in Equation (14) and markets clear:*

$$\int_{-\infty}^t \nu e^{-\nu(t-s)} c_{s,t} ds = Y_t, \quad (15)$$

$$\int_{-\infty}^t \nu e^{-\nu(t-s)} \pi_{s,t} ds = 0, \quad (16)$$

$$\int_{-\infty}^t \nu e^{-\nu(t-s)} (W_{s,t} - \pi_{s,t}) ds = 0. \quad (17)$$

We solve the individual optimization problems by martingale methods as in Cox and Huang (1989). Consider an agent born at time s . The static optimization problem for this

agent can be written as

$$\max_{c_s} E_{s,s} \left[\int_s^\infty e^{-(\rho+\nu)(t-s)} \log(c_{s,t}) dt \right]$$

s.t.

$$E_{s,s} \left[\int_s^\infty e^{-\nu(t-s)} \xi_{s,t} c_{s,t} dt \right] = E_{s,s} \left[\int_s^\infty e^{-\nu(t-s)} \xi_{s,t} y_{s,t} dt \right].$$

From the first order conditions (FOCs), we have

$$\frac{e^{-(\rho+\nu)(t-s)}}{c_{s,t}} = \kappa_s e^{-\nu(t-s)} \xi_{s,t}, \quad (18)$$

where κ_s denotes the Lagrange multiplier. For $s \leq u \leq t$ the FOCs imply

$$e^{-(\rho+\nu)(t-u)} \left(\frac{c_{s,u}}{c_{s,t}} \right) = e^{-\nu(t-u)} \frac{\xi_{s,t}}{\xi_{s,u}}. \quad (19)$$

The total wealth at time $u \geq s$ of an agent born at time s is the sum of the value of endowment, $H_{s,u} = \frac{1}{\xi_{s,u}} E_{s,u} \left[\int_u^\infty e^{-\nu(t-s)} \xi_{s,t} y_{s,t} dt \right]$, and financial wealth, $W_{s,t}$. Let the total wealth at time u be $\hat{W}_{s,u} = H_{s,u} + W_{s,u}$. Using the static budget constraint, we obtain

$$c_{s,u} = (\rho + \nu) \hat{W}_{s,u}. \quad (20)$$

Equation (20) shows that the standard constant wealth to consumption ratio with log utility holds in our overlapping generations setup with incomplete information. Using the market clearing conditions and the individual aggregate endowments, we have

$$Y_t = \int_{-\infty}^t \nu e^{-\nu(t-s)} c_{s,t} ds = \int_{-\infty}^t \nu e^{-\nu(t-s)} (\rho + \nu) \hat{W}_{s,t} ds = (\rho + \nu) \hat{W}_t, \quad (21)$$

and, consequently, aggregate wealth, \hat{W}_t , is given by

$$\hat{W}_t = \frac{Y_t}{\rho + \nu}. \quad (22)$$

Total aggregate wealth, \hat{W}_t , equals the value of aggregate endowments, H_t . The wealth of an agent born at time s corresponds exactly to the aggregate wealth, i.e., $H_{s,s} = \frac{Y_s}{\rho + \nu}$. Combining this with Equation (20), we have that the consumption of an agent born at time s equates with endowment

$$c_{s,s} = Y_s. \quad (23)$$

The initial consumption is given by the output at time of birth. This can also be interpreted as a state dependent Pareto weight. The reason why the Pareto weight is stochastic, is that agents cannot hedge against output fluctuations prior to their birth. However, once born, they face complete markets.

Using the relation in Equation (23), we obtain the following proposition:

Proposition 2. *Optimal consumption at time t of agents born at time $s \leq t \leq \tau$, where τ denotes the stochastic time of death, is*

$$c_{s,t} = Y_s e^{-\rho(t-s)} \begin{pmatrix} \eta_{s,t} \\ \eta_{s,s} \end{pmatrix} \begin{pmatrix} \xi_s \\ \xi_t \end{pmatrix}. \quad (24)$$

The next proposition characterizes the stochastic discount factor.

Proposition 3. *In equilibrium, the stochastic discount factor is*

$$\xi_t = \frac{X_t}{Y_t}, \quad (25)$$

where X_t solves the integral equation

$$X_t = \int_{-\infty}^t \nu e^{-(\rho+\nu)(t-s)} X_s \frac{\eta_{s,t}}{\eta_{s,s}} ds. \quad (26)$$

The fraction of aggregate output at time t consumed by an agent born at time s is $\frac{c_{s,t}}{Y_t}$, and since the measure of agents born at time s equals $\nu e^{-\nu(t-s)}$, we have that the fraction of aggregate output at time t consumed by agents born at time s is

$$f_{s,t} = \nu e^{-\nu(t-s)} \frac{c_{s,t}}{Y_t} = \nu e^{-\nu(t-s)} \frac{Y_s e^{-\rho(t-s)} \left(\frac{\eta_{s,t}}{\eta_{s,s}} \right) \left(\frac{\xi_s}{\xi_t} \right)}{Y_t} = \nu e^{-(\rho+\nu)(t-s)} \left(\frac{X_s}{X_t} \right) \left(\frac{\eta_{s,t}}{\eta_{s,s}} \right). \quad (27)$$

Next, we introduce a decomposition of X_t :

Proposition 4. *Let*

$$d\bar{\eta}_t = \bar{\Delta}_t \bar{\eta}_t dz_t, \quad (28)$$

where $\bar{\Delta}_t$ is the consumption share weighted average disagreement in the economy given by

$$\bar{\Delta}_t = \int_{-\infty}^t f_{s,t} \Delta_{s,t} ds, \quad (29)$$

then

$$X_t = e^{-\rho t} \bar{\eta}_t. \quad (30)$$

The process $\bar{\eta}_t$ captures the change of measure from the “market view” as measured by the consumption share weighted average belief, $\bar{\mu}_t = \int_{-\infty}^t f_{s,t} \hat{\mu}_{s,t} ds$, to the actual probability measure. The stochastic discount factor can then be decomposed in the following way:

$$\xi_t = \underbrace{\bar{\eta}_t}_{\text{Effect from disagreement}} \times \underbrace{\frac{e^{-\rho t}}{Y_t}}_{\text{Log utility discount factor}}. \quad (31)$$

The stochastic discount factor in Equation (31) can be interpreted as a stochastic discount factor with a representative agent with beliefs given by $\bar{\mu}_t$. Although, every agent in the economy has a Kalman gain that is decreasing over time, the overlapping generation structure leads to a stationary and stochastic gain for the representative agent.

The next proposition characterizes the real short rate and the market price of risk.

Proposition 5. *In equilibrium, the real short rate is*

$$r_t = \rho + \mu_Y - \sigma_Y^2 + \sigma_Y \bar{\Delta}_t, \quad (32)$$

and the market price of risk is

$$\theta_t = \sigma_Y - \bar{\Delta}_t. \quad (33)$$

The expression for the risk-free rate is intuitive: The real rate is the standard real rate with log utility, $\rho + \mu_Y - \sigma_Y^2$, plus a correction for the disagreement in the economy. In times when the “market view” of the expected growth in the economy is higher than the actual growth, this translates into demand for borrowing to smooth consumption across time. As aggregate consumption is fixed, the real interest rate increases to clear markets. Using the definition for $\bar{\Delta}_t$, we express the real rate as $r_t = \rho + \bar{\mu}_t - \sigma_Y^2$. We see that the expression for the real rate equals to the one in the standard log utility case, except that the expected aggregate endowment growth, μ_Y , is replaced by the market view, $\bar{\mu}_t$. The expression for the market price of risk is also intuitive: It is given by the standard log-utility market price of risk adjusted for disagreement. Again it is useful to rewrite the expression for the market price of risk as $\theta_t = \sigma_Y - \frac{1}{\sigma_Y} (\bar{\mu}_t - \mu_Y)$. We see that when the market is relatively optimistic about the expected growth, i.e., $\bar{\mu}_t > \mu_Y$, then the market price of risk is low. Indeed, under the objective probability measure, the risky asset is expensive and, thus, the market price of risk must be low.

Proposition 6. *We have that after a positive shock the risk-free rate increases, i.e., $\frac{\partial r_t}{\partial z} > 0$, and the market price of risk decreases, i.e., $\frac{\partial \theta_t}{\partial z} < 0$.*

The intuitions for Proposition 6 follow from the dynamics of $\bar{\Delta}_t = \frac{\bar{\mu}_t - \mu_Y}{\sigma_Y}$. When there is a positive shock to aggregate output, i.e., $dz_t > 0$, all agents in the economy revise their expectations upwards and the disagreement process $\bar{\Delta}_t$ increases. Since the real short rate increases and the market price of risk decreases we refer to the risk-free rate as procyclical and the market price of risk as countercyclical.

Before proceeding to the optimal portfolio policies, it is convenient to derive the dynamics for the optimal individual consumption.

Proposition 7. *The dynamics of individual consumption is*

$$dc_{s,t} = c_{s,t} (\mu_{c_{s,t}} dt + \sigma_{c_{s,t}} dz_t), \quad (34)$$

where the drift and the diffusion are given by

$$\mu_{c_{s,t}} = \mu_Y + (\sigma_Y - \bar{\Delta}_t) (\Delta_{s,t} - \bar{\Delta}_t), \quad \sigma_{c_{s,t}} = \sigma_Y + \Delta_{s,t} - \bar{\Delta}_t. \quad (35)$$

Proposition 7 shows that the difference between the individual estimation error, $\Delta_{s,t}$, and the consumption share weighted average disagreement in the economy, $\bar{\Delta}_t$, drive the diffusion of the individual consumption.

The next proposition characterizes the optimal portfolio policy of an agent born at time s .

Proposition 8. *The optimal dollar amount invested in the risky asset for an agent born at time s is*

$$\pi_{s,t} = \frac{\Delta_{s,t} - \bar{\Delta}_t}{\sigma^S} \hat{W}_{s,t} + \frac{\sigma_Y}{\sigma^S} W_{s,t}. \quad (36)$$

The optimal portfolio has two components. The sign of the first one is determined by the relative disagreement between the agent and the “market view” scaled by the variance of the risky asset. To understand the second component, it is useful to look at the optimal exposure of total wealth to the Brownian motion, which is $\hat{W}_{s,t} \theta_{s,t} = \hat{W}_{s,t} (\sigma_Y + \Delta_{s,t} - \bar{\Delta}_t)$. All agents have intrinsic exposure to the shock through their endowment given by $\sigma_Y H_{s,t}$. The optimal portfolio allocation to the risky asset adjusts for the difference between the optimal exposure, $\hat{W}_{s,t} (\sigma_Y + \Delta_{s,t} - \bar{\Delta}_t)$, and the intrinsic exposure, $\sigma_Y H_{s,t}$, taking into account that the local volatility of the risky asset in general differs from the volatility of endowments.

It is of interest to inspect the optimal portfolio allocation of a new born agent

$$\pi_{s,s} = -\frac{\bar{\Delta}_s}{\sigma^S} H_{s,s}. \quad (37)$$

From Equation (37), we see that the optimal portfolio for a new born depends exclusively on the relative disagreement in the market since the agent is born with the correct prior about the expected aggregate endowment growth. Further, the dollar amount that he invests in the stock market depends exclusively on his endowment since financial wealth at birth is zero. In the case where disagreement between the market view coincides with the actual growth rate, the portfolio allocation is trivially zero. If the market is more optimistic (pessimistic) than the actual growth rate $\bar{\mu}_t > \mu_Y$ ($\bar{\mu}_t < \mu_Y$), then the newborn will short (go long in) the risky asset.

3 Numerical Illustrations

To strengthen the intuition for our results, we simulate an economy populated by a large number of cohorts where one cohort is born every period. One period in the simulation represents one month. After 6000 burn-in periods, we obtain an economy with 6000 cohorts. We use the final values from the burn-in simulation as starting values for simulating this economy forward for another 1200 periods. We generate data from 10,000 simulations, each with 1200 periods or 100 years.

In the numerical illustrations, we use the following parameter values: We follow Gârleanu and Panageas (2015) and set the time discount factor, ρ , at 0.1% and the birth and death intensity, ν , at 2%. The drift, μ_Y , and volatility, σ_Y , of aggregate endowment are set to 2% and 3.3%, respectively, which is consistent with the long sample in Campbell and Cochrane (1999). Prior variance of the expected aggregate endowment growth, \bar{V} , is 0.01^2 . The diffusion term for the risky asset, σ^S , equals 15%. With these values, the average risk-free rate and expected return on the risky asset are 2.00% and 2.50%, respectively.

Proposition 6 shows that under the true probability measure the risk-free rate increases after a positive shock while the market price of risk decreases. As $\sigma^S > \sigma_Y$, the expected return on the risky asset is decreasing after a positive shock.⁷ However, the objective probability measure will, in general, be different from the probability measure of individual agents. Specifically, young agents revise their expectations about the expected growth of endowment much more aggressively than older agents. Under the objective measure, the expected endowment growth is constant. This has implications for how agents with different experience perceive shocks to their endowment, and, hence shocks to the risky asset. Specifically, the expected return under the true probability measure is $\mu_t^S = r_t + \sigma^S \sigma_Y - \sigma^S \bar{\Delta}_t$ while for an agent born at time s it is


$$\mu_{s,t}^S = \underbrace{r_t + \sigma^S \sigma_Y - \sigma^S \bar{\Delta}_t}_{\text{Expected return under the true measure}} + \sigma^S \Delta_{s,t}. \quad (38)$$

We have that $\Delta_{s,t}$ and $\bar{\Delta}_t$ increase after positive shocks to endowment. As $\Delta_{s,t}$ has a positive sign and $\bar{\Delta}_t$ has a negative sign in the perceived expected return, the true and perceived expected return can be negatively correlated. Especially for young agents, for which $\Delta_{s,t}$ reacts strongly to shocks, the true and perceived expected return on the risky asset move in opposite directions. When agents accumulate experience, the variance of the estimation error, $\Delta_{s,t}$, declines and eventually the true and the perceived expected return correlate positively. To illustrate this, we plot the correlation between the objective expected return, μ_t^S , and the perceived expected return, $\mu_{s,t}^S$, by cohort age in Figure 1. The figure shows that the correlation is strongly negative with a value of -0.8 for young agents, from where it increases monotonically in age. As the expected return under the true measure decreases after positive shocks to endowment, Figure 1 shows that for young and inexperienced agents the perceived expected return increases after positive shocks. This implies that young agents form expectations that mimics return extrapolation.⁸

⁷The expected excess return is always decreasing after positive shocks.

⁸We discuss return extrapolation in more detail in Section 5.


Figure 1: *Objective and Perceived Returns*. The figure plots the correlation between the expected return under the objective measure and the perceived expected return by cohort lifespan. Each observation is calculated using a window of 60 non-overlapping observations (5 years). The figure is averaged from 10,000 simulations with 1200 periods or 100 years per simulation.



Our results regarding perceived expected returns provide a direct view at optimal portfolio allocations. A positive shock to the stock market increases young agents expectations about the future stock market return. In turn, their demand for the risky asset increases. Old agents reduce their expectations about expected stock market relative to the young and, coherently, they, therefore, reduce their portfolio holdings in the risky asset. Figure 2, which shows the correlation between portfolio allocations and stock market shocks by cohort age, confirms this intuition. The correlation increases during the first years and then declines monotonically over time reaching -1 in ripe old age. The reason for the correlation to approach -1 in the long run can be understood from considering the general equilibrium properties of the model. For the market to clear, old agents counter-balance the portfolio allocations of the young. Specifically, consider the optimal portfolio allocation for an agent as shown in Equation (8). Using the definitions of $\bar{\Delta}_t$ and $\Delta_{s,t}$, we express the optimal portfolio as $\pi_{s,t} = \frac{\hat{\mu}_{s,t} - \bar{\mu}_t}{\sigma_Y \sigma^S} \hat{W}_{s,t} + \frac{\sigma_Y}{\sigma^S} W_{s,t}$, from which we see that the portfolio allocation is driven

by the difference between the agents' and the market view, $\hat{\mu}_{s,t} - \bar{\mu}_t$. Following a positive shock to the stock market, $dz_t > 0$, the young revise their expectations about aggregate consumption growth more than the revision in the market view, since $Var(\hat{\mu}_{s,t}) \geq Var(\bar{\mu}_t)$. Hence, the young increase their allocation in the risky asset. Old agents revise their expectations less than the market view, since $Var(\hat{\mu}_{s,t}) \leq Var(\bar{\mu}_t)$. Therefore, they counter-balance the behavior of the young by reducing their demand for the risky asset, thereby the market clears.

Figure 2: *Portfolios and Shocks*. The figure plots the correlation between portfolio allocations and stock market shocks by cohort lifespan. Each observation is calculated using a window of 60 non-overlapping observations (5 years). The figure is averaged from 10,000 simulations with 1200 periods or 100 years per simulation.



The model with learning from experience can shed light on several stylized features of survey data. Specifically, we can conduct a survey in the model and calculate the mean forecast. We define the average belief or survey forecast in the economy as

$$\hat{\mu}_t^S = \int_{-\infty}^t \nu e^{-\nu(t-s)} \mu_{s,t}^S ds. \quad (39)$$

Figure 3 shows a typical simulated path of the true expected return, μ_t^S , and the survey forecast, $\hat{\mu}_t^S$. First, the true expected return is more volatile than the average belief. The unconditional standard deviation of the true expected return is 1.2% while it is only 0.3% for the survey forecast.⁹ This finding is consistent with Piazzesi, Salomao, and Schneider (2015), who find that statistical measures of the bond risk premium are more volatile than the bond risk premium measured from survey data. Second, Piazzesi, Salomao, and Schneider (2015) also find that while statistical measures of the bond risk premium are countercyclical, survey measures of the bond risk premium are not very cyclical. This is also consistent with Figure 3, where the true expected return is negatively correlated with the survey measure of the expected return.¹⁰ In addition, Greenwood and Shleifer (2014) show that average beliefs from survey measures are negatively correlated with future realized returns. Hence, when the average belief about future returns are high, on average realized returns tend to be low. The model with learning from experience is also consistent with this negative correlation. To formally test the relation in the model, we follow Greenwood and Shleifer (2014) and use the survey forecast to predict future returns. Indeed, Table 1 shows that the survey forecast predicts returns with a negative sign.

Table 1: *Average Belief and Future Realized Returns.* The table shows coefficient estimate from a regression of future realized returns on the average expected returns: $R_{t+1} = a + b\hat{\mu}_t^S + e_{t+1}$. The regression uses data from 10,000 simulations with 1200 periods (monthly observations) or 100 years per simulation. The table reports the average coefficients and t-statistics across the 10,000 sample paths.


$\hat{\mu}_t^S$	-1.50
	[-5.68]
constant	0.06
	[10.07]

The intuition for this result is that the actual expected return is driven by the consumption share weighted average belief and not the average belief, $\hat{\mu}_t^S$. Thus, the average belief puts too much weight on the young and inexperienced agents with low wealth. Be-

⁹Unconditional moments are calculated as the average values across the 10,000 paths of 100 years of monthly observations.

¹⁰The unconditional correlation between the true expected return and the average belief about the expected return is -0.32 .

Figure 3: *True Expected Return versus Mean Forecast from Survey*. The figure shows a representative path of the true expected return, μ_t^S , and the survey measure, $\hat{\mu}_t^S$, with 1200 periods (monthly observations) or 100 years.




cause young agents perceive high future returns after a series of positive shocks to the stock market, the survey forecast reflects their view and, therefore, predicts future returns with a negative sign.

Above we investigate the way young individuals make mistakes in estimating expected returns by comparing the perceived expected returns with the objective expected returns. Figure 4 shows how such estimation mistakes affect expected log consumption growth dynamics. In the top plot, we see that cohort specific expected log consumption growth falls for young agents but steadily increases as agents learn and age. When young, any shock induces a large revision in estimates for growth as well as large revisions to portfolios. Because the young trade with experienced agents, their consumptions growth is low. Figure 4 shows that the economic cost of learning from experience are potentially large as losses accumulate during the early phases of life and that only after an extended time of learning will an agent's consumption growth catch up. One measure of the welfare cost is to calculate the expected lifetime utility of an agent that is learning from experience using the true probability measure. For the baseline parameters, the welfare cost of learning from experience amounts to 72% of the welfare cost of aggregate fluctuations in endowments. While the

welfare costs of fluctuations in aggregate endowments are independent of the time discount factor, the welfare cost of learning from experience is not. If the time discount factor is increased from $\rho = 0.001$ to $\rho = 0.02$, then the welfare cost of learning from experience is 2.75 times as high as the welfare cost of aggregate fluctuations.¹¹ The bottom plot of Figure 4 shows that cohort specific consumption volatility also falls for young agents and steadily increases as agents learn and age. However, cohort consumption volatility at all ages is always significantly above the aggregate consumption volatility of 3.3%.

Figure 4: *Cohort Specific Consumption Growth and Volatility.* The figure plots the drift term of log consumption under the objective measure and the volatility of consumption by cohort lifespan. The figure is averaged from 10,000 simulations with 1200 periods or 100 years per simulation.




When agents learn from experience, large positive or negative realizations of aggregate endowment growth impact the belief of the young strongly. A natural question is whether a large negative or positive shocks to endowment growth has a lasting effect on the beliefs of the young. Malmendier and Nagel (2011) show that empirically this is indeed the case as the effect will fade away only after 30 years. Similarly, Knüpfer, Rantapuska, and Sarvimäki (2014) show that labor market shocks can permanently reduce the willingness to invest in

¹¹We assume that young agents start with the correct prior about μ_Y . This assumption reduces the welfare cost of learning from experience as young agents are initially correct.


risky assets.¹² To understand how particularly good or bad consumption trajectories in the model affect perceived expected returns, we sort on average consumption growth during the first 5 years after the 6000 burn-in periods. We define the top 5% and bottom 5% average consumption growth sample paths as good and bad times, respectively. Figure 5 shows the objective and the perceived expected returns in good and bad times. We see that in good (bad) times the objective expected return declines (increases) over the consumption boom (depressed) period but gradually increases (declines) back to the unconditional expected return of roughly 2.5%. In contrast, in good (bad) times the perceived expected return remains significantly above (below) the unconditional expected return even after 100 years and after 30 years only half of the effect fades away. Therefore, the model is consistent with the empirical evidence on how significant experiences can be a source of persistent disagreement. Figure 6 shows the portfolio policies in good and bad times. One can see that large negative or positive realizations early on in life have a long lasting impact on the optimal portfolio choice of agents.

Figure 5: *Objective and Perceived Returns in Good and Bad Times.* The figure plots the objective expected return, μ_t^S , and the perceived expected return, $\mu_{s,t}^S$, by cohort age in good and bad times. Good (bad) times are defined as the top (bottom) 5% average consumption growth during the first 5 years after the 6000 burn-in periods. Each observation is calculated using a window of 60 non-overlapping observations (5 years). The figure is averaged from 10,000 simulations with 1200 periods or 100 years per simulation.



¹²Perhaps the first paper that theoretically considers the idea that a persistent shock such as the Great Depression might lead to significant and persistent, but not constant, pessimism and through that drive up the risk premium is Cogley and Sargent (2008). They, however, do not consider an overlapping generations model setting hence expectations converge to rational expectations.

Figure 6: *Portfolio Policies in Good and Bad Times*. The figure plots the portfolio policy, $\pi_{s,t}$, by cohort age in good and bad times. Good (bad) times are defined as the top (bottom) 5% average consumption growth during the first 5 years after the 6000 burn-in periods. Each observation is calculated using a window of 60 non-overlapping observations (5 years). The figure is averaged from 10,000 simulations with 1200 periods or 100 years per simulation.



4 Extensions

4.1 Within-Cohort Heterogeneity

In the model in Section 2, we assume that all agents within a cohort are identical. One question that naturally arises is how robust the results are to within cohort heterogeneity. Therefore, in this subsection, we extend the baseline model to allow for heterogeneity in beliefs within cohorts. We assume that in every period there is a distribution of initial beliefs. We index agent types by a and assume that agents within a cohort are distributed following a generic distribution $g(a)$ defined over the domain $[a, \bar{a}]$, such that $\int_a^{\bar{a}} g(a) da = 1$. An agent of type a has a prior mean and variance given by $\hat{\mu}_{a,s,s}$ and $V_{a,s,s}$, respectively. Here, we do allow the prior mean and variance to potentially depend on the time when the agent is born. Given the prior mean and variance, we calculate the estimation error, $\Delta_{a,s,t}$,

using standard filtering. The next proposition states the estimation error.

Proposition 9. *The estimation error at time t of an agent of type a born at time s is*

$$\Delta_{a,s,t} = \frac{\sigma_Y}{\sigma_Y^2 + V_{a,s,s}(t-s)} (\hat{\mu}_{a,s,s} - \mu_Y) + \frac{V_{a,s,s}(z_t - z_s)}{\sigma_Y^2 + V_{a,s,s}(t-s)}. \quad (40)$$

Moreover, we have that $\Delta_{a,s,s} = \frac{\hat{\mu}_{a,s,s} - \mu_Y}{\sigma_Y}$ and $\lim_{t \rightarrow \infty} \Delta_{s,t} = 0$.

The dynamics of the disagreement process is

$$d\eta_{a,s,t} = \Delta_{a,s,t} \eta_{a,s,t} dz_t. \quad (41)$$

An agent of type a born at time s maximizes

$$E_{a,s,s} \left[\int_s^\infty e^{-(\rho+\nu)(t-s)} \log(c_{a,s,t}) dt \right], \quad (42)$$

where the expectation is taken with respect to the belief of an agent of type a born at time s . The equilibrium can be solved using the same approach as in Section 2. The next proposition characterizes the equilibrium real short rate and market price of risk.

Proposition 10. *In equilibrium, the real short rate is*

$$r_t = \rho + \mu_Y - \sigma_Y^2 + \sigma_Y \bar{\Delta}'_t, \quad (43)$$

and the market price of risk is

$$\theta_t = \sigma_Y - \bar{\Delta}'_t, \quad (44)$$

where

$$\bar{\Delta}'_t = \int_{-\infty}^t \int_{\underline{a}}^{\bar{a}} f_{a,s,t} \Delta_{a,s,t} da ds, \quad (45)$$

and

$$f_{a,s,t} = \nu e^{-\nu(t-s)} g(a) \left(\frac{X_s}{X_t} \right) \left(\frac{\eta_{a,s,t}}{\eta_{a,s,s}} \right), \quad (46)$$

with X_t solving the integral equation

$$X_t = \int_{-\infty}^t \nu e^{-(\rho+\nu)(t-s)} X_s \left(\int_{\underline{a}}^{\bar{a}} g(a) \frac{\eta_{a,s,t}}{\eta_{a,s,s}} da \right) ds. \quad (47)$$

From Proposition 10, we see that the real short rate and market price of risk take similar forms as in the baseline model in Section 2. In the next two subsections, we discuss two specific examples of within-cohort heterogeneity. In the first example, we assume that the distribution of initial prior beliefs about the mean is Gaussian. In the second example, we let a fraction of the population learn from all historical data and, hence, these agents know the true mean instantaneously.

4.1.1 Gaussian Type Distribution

In this subsection, we follow Atmaz (2015) and specify the distribution function, $g(a)$, to be Gaussian, that is

$$g(a) = \frac{1}{\sqrt{2\pi\nu_0^2}} e^{-\frac{1}{2}\frac{a^2}{\Sigma}}, \quad (48)$$

with mean zero and variance Σ .¹³ The belief of an agent of type a when born is $\hat{\mu}_{a,s,s} = \mu_Y + a$. Agents are homogeneous with respect to the prior variance, which is given by \bar{v} . As the bias parameter a has mean zero, the average agent is born with the correct prior for the mean. By standard filtering, the error process at time t of an agent born at time s with initial belief a is

$$\Delta_{a,s,t} = \frac{\sigma_Y}{\sigma_Y^2 + \bar{v}(t-s)} a + \frac{\bar{v}(z_t - z_s)}{\sigma_Y^2 + \bar{v}(t-s)}. \quad (49)$$

The dynamics of the disagreement process is

$$d\eta_{a,s,t} = \eta_{s,a,t} \Delta_{s,a,t} dz_t. \quad (50)$$

¹³Atmaz (2015) derives equilibrium with a continuum of agents differing in their beliefs. His agents are infinitely lived however, so there are no generation specific beliefs.

By aggregating total consumption within a cohort born at time s , we have

$$\begin{aligned} c_{s,t} &= \int_{-\infty}^{\infty} c_{a,s,t} da = \int_{-\infty}^{\infty} g(a) Y_s e^{-\rho(t-s)} \frac{\eta_{a,s,t} \xi_s}{\eta_{a,s,s} \xi_t} da \\ &= Y_s e^{-\rho(t-s)} \frac{\xi_s}{\xi_t} \int_{-\infty}^{\infty} g(a) \frac{\eta_{a,s,t}}{\eta_{a,s,s}} da. \end{aligned} \quad (51)$$

From Equation (51), we see that the only term that differs between the agent types born at time s is the disagreement process. Define the aggregate disagreement process for the cohort born at time s as

$$\eta_{s,t} = \int_{-\infty}^{\infty} g(a) \frac{\eta_{a,s,t}}{\eta_{a,s,s}} da. \quad (52)$$

By following the same approach as in Atmaz (2015), one can show that

$$d\eta_{s,t} = \Delta_{s,t} \eta_{s,t} dz_t, \quad (53)$$

where

$$\Delta_{s,t} = \frac{(\bar{v} + \Sigma)(z_t - z_s)}{\sigma_Y^2 + (\bar{v} + \Sigma)(t - s)}. \quad (54)$$

Equation (54) has the same form as in the base case in Proposition 1 with $\bar{V} = \bar{v} + \Sigma$. Hence, we can interpret the base case model as a model with heterogeneous initial beliefs that are normally distributed with zero bias on average. Here, the dynamics of the cohort specific belief behave similar to the base case even without learning, i.e., when $\bar{v} = 0$. In this case, $\bar{V} = \Sigma$; thus, it only depends on the within cohort cross-sectional heterogeneity. Consequently, the convergence of the cohort specific belief does not occur because of learning, but due to market selection. Agents that start with a relatively more correct initial belief have a higher consumption growth and eventually dominate the cohort. However, when $\Sigma > 0$ and $\bar{v} = 0$, the individual agents do not change their beliefs and, therefore, a simple average of the beliefs in the economy is constant. This contrasts with the case when $\bar{v} > 0$ since the beliefs of individual agents do change over time in the baseline model in Section 2.

4.1.2 Introducing Rational Learners


Assuming that only a fraction $a \in [0, 1]$ of a cohort exclusively learn from realizations of the endowment during their own life yields another form of within-cohort heterogeneity. The remaining agents in the cohort use all past historical information and, therefore, know the true expected growth rate of the endowment. For the agents who learn from experience, the dynamics of the disagreement process, $\eta_{a,s,t}$, is same as in the baseline case in Section 2. For the agents who know the true expected growth rate the disagreement process, $\eta_{1-a,s,t}$, is equal to one for all states and times. Hence, we have $X_t = \int_{-\infty}^t \nu e^{-(\rho+\nu)(t-s)} X_s \left(a \frac{\eta_{a,s,t}}{\eta_{a,s,s}} + (1-a) \right) ds$ with $\frac{d\eta_{a,s,t}}{\eta_{a,s,t}} = \Delta_{s,t} dz_t$ and $\Delta_{s,t}$ as in Proposition 1.

Figure 7 shows the correlation between the expected return as perceived by an agent of a particular age that is learning from experience and the true expected return for four economies with different fractions of agents who learn from experience. We see that as the fraction of the population that learns from experience is reduced, the older the agents who learn from experience have to be before the correlation between their perceived expected return and the true expected return is positive. This can be understood by comparing the market price of risk of an agent that is learning from experience with that of the true market price of risk

$$\begin{aligned} \theta_t &= \sigma_Y - \bar{\Delta}'_t = \sigma_Y - a \int_{-\infty}^t \nu e^{-(\rho+\nu)(t-s)} \frac{X_s}{X_t} \frac{\eta_{a,s,t}}{\eta_{a,s,s}} ds, \\ \theta_{a,s,t} &= \sigma_Y - \bar{\Delta}'_t + \Delta_{s,t}. \end{aligned} \tag{55}$$

Here, $\bar{\Delta}'_t$ and $\Delta_{s,t}$ are positively correlated and, therefore, the perceived expected return is only positively correlated with the true expected return if $\Delta_{s,t}$ is sufficiently small relative to $\bar{\Delta}'_t$. However, $\bar{\Delta}'_t$ is increasing in the fraction of agents that are learning from experience. Hence, if the fraction of the population who learn from experience is small, then $\Delta_{s,t}$ has to be smaller for the correlation to be positive or the agents have to learn for a longer period of time.

Figure 7: *Objective and Perceived Returns*. The figure plots the correlation between the expected return under the objective measure and the perceived expected return of the agents that are learning from experience as a function of lifespan. Each observation is calculated using a window of 60 non-overlapping observations (5 years). The fraction of each cohort that is learning from experience is $a = (1, 0.75, 0.5, 0.25)$. The figure is averaged from 10,000 simulations with 6000 periods or 500 years per simulation.



4.2 Multiple Risky Assets

In the model in Section 2, there is only one source of risk. Agents observe the endowment process and from that learn about the expected growth rate. Malmendier and Nagel (2011) show that individuals who experience high stock market returns invest more in the stock market, but do not alter their bond portfolios. Conversely, investors who experience high bond market returns invest more in bonds, but do not change their investment in stocks. Hence, the experienced return is asset specific. In this subsection, we extend the baseline model to two risky assets to show that a model with learning from experience is consistent with an asset specific learning from experience. Rather than setting up the multi-asset model in full generality, we focus on two sources of risk: endowment risk and inflation risk. We do this for two reasons. First, Malmendier and Nagel (2014) show using survey data that consumers exhibit learning from experience behavior when forecasting inflation. Second, focusing on learning from inflation allows for the interpretation of the second risky asset as

a nominal bond or some other asset for which inflation is an important driver of returns.

We follow Xiong and Yan (2010) and Ehling, Gallmeyer, Heyerdahl-Larsen, and Illeditsch (2015) and assume that inflation is exogenous and independent of aggregate endowment. Specifically, we keep the same structure as in Section 2, but in addition to learning about the expected growth rate of endowment agents also learn about the expected inflation. We model the price level as

$$d\Pi_t = \Pi_t (\mu_\Pi dt + \sigma_\Pi dz_t^\Pi). \quad (56)$$

Since agents do not know the expected inflation, μ_Π , they learn about it from observing the price level, Π_t , during their own life. As endowments and the price level are independent, the learning about the endowment growth can be separated from the learning about the expected inflation. To simplify, agents start with the correct value for the prior mean and the prior variance about inflation is \bar{V}_Π . The error process for the expected endowment growth, $\Delta_{s,t}$ is given in Proposition 1. The error process for expected inflation at time t for a cohort born at time s is

$$\Delta_{s,t}^\Pi = \frac{\bar{V}_\Pi (z_t - z_s)}{\sigma_\Pi^2 + \bar{V}_\Pi (t - s)}. \quad (57)$$

Similar to Proposition 1, we have that $\Delta_{s,s}^\Pi = 0$ and $\lim_{t \rightarrow \infty} \Delta_{s,t}^\Pi = 0$. The relation between the true Brownian motion and the perceived Brownian motion is $dz_{s,t}^\Pi = dz_t^\Pi - \Delta_{s,t}^\Pi dt$. The disagreement process, $\eta_{s,t}$, decomposes into two parts

$$\eta_{s,t} = \eta_{s,t}^Y \eta_{s,t}^\Pi, \quad (58)$$

where $d\eta_{s,t}^Y = \Delta_{s,t} \eta_{s,t}^Y dz_t$ and $d\eta_{s,t}^\Pi = \Delta_{s,t}^\Pi \eta_{s,t}^\Pi dz_t^\Pi$.

In addition to the asset specified in Section 2, the agents trade a risky asset in zero net supply with price P_t that is locally perfectly correlated with the shock to inflation. We refer to this asset as the long-term bond. The dynamics of it is

$$\frac{dP_t}{P_t} = (\mu_t^P dt + \sigma^P dz_t^\Pi) = (\mu_{s,t}^P dt + \sigma^P dz_{s,t}^\Pi). \quad (59)$$

The equilibrium is similarly to the equilibrium in Section 2 and that the real short rate has the exact same form as in the case with only endowment risk. However, there are now two priced sources of risk in the economy. The market prices of endowment shocks, θ_t , and inflation shocks, θ_t^Π , are

$$\begin{aligned}\theta_t &= \sigma_Y - \bar{\Delta}_t, \\ \theta_t^\Pi &= -\bar{\Delta}_t^\Pi,\end{aligned}\tag{60}$$

where $\bar{\Delta}_t$ has the same form as in Equation (29) and $\bar{\Delta}_t^\Pi = \int_{-\infty}^t f_{s,t} \Delta_{s,t}^\Pi ds$. Let $\pi_{s,t}^S$ and $\pi_{s,t}^P$ be the dollar amount invested in the risky assets correlated with endowment shocks and inflation shocks, respectively. The next proposition characterizes the optimal portfolio.

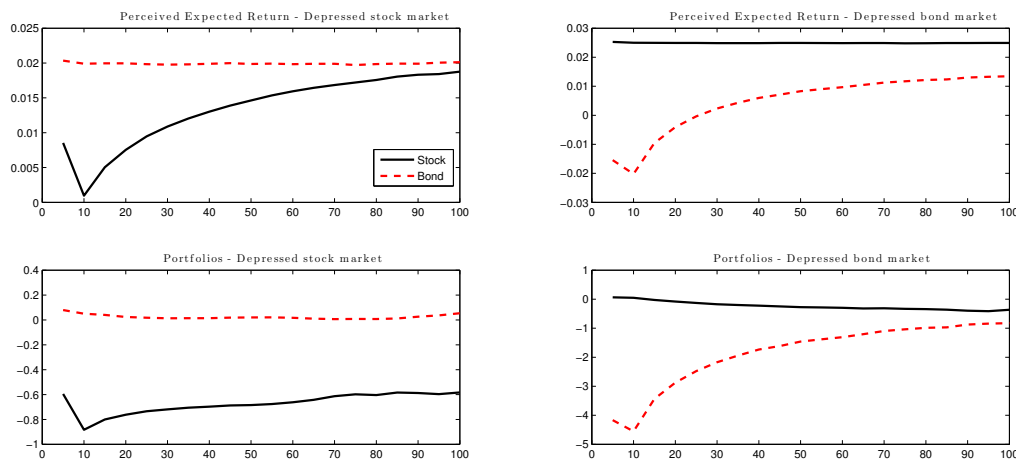
Proposition 11. *The optimal dollar amount invested in the risky assets for an agent born at time s are*

$$\begin{aligned}\pi_{s,t}^S &= \frac{\Delta_{s,t} - \bar{\Delta}_t}{\sigma^S} \hat{W}_{s,t} + \frac{\sigma_Y}{\sigma^S} W_{s,t}, \\ \pi_{s,t}^P &= \frac{\Delta_{s,t}^\Pi - \bar{\Delta}_t^\Pi}{\sigma^P} \hat{W}_{s,t}.\end{aligned}$$

To examine the case with two risky assets, we reuse the parameters from the baseline model. For the risky bond, we assume that the volatility, σ_P , is 5%. We simulate two sets of paths. In the first set, we draw 200,000 five year paths for the endowment shocks, dz . Next, we pick the 10,000 worst realizations for the Brownian motion after five years. We call this set of paths depressed stock market. For the second set of paths, we draw the shocks to inflation rather than endowment. We refer to these paths as high inflation period or depressed bond market. Figure 8 shows the expectation of the stock market and bond market (top plots) and portfolio positions in the stock and the bond (bottom plots) for an agent born just before the bad realizations. The left plots show the depressed stock market and the right plots show the depressed bond market. We see from the plot that portfolio

positions and expectations are asset specific, i.e., an agent who experiences a depressed stock market shorts the stock, but does not change the position in the bond relative to an agent who does not experience a depressed stock market.

Figure 8: *Expected Returns and Portfolio Policies in Depressed Markets*. The figure plots the perceived expected stock market return, $\mu_{s,t}^S$, and expected bond market return, $\mu_{s,t}^P$, (top plots) and the portfolio position in the stock, $\pi_{s,t}^S$ and the bond, $\pi_{s,t}^P$, (bottom plots) by cohort age in a depressed stock market (left-hand plots) and a depressed bond market (right-hand side). A depressed (stock and bond) market is defined as the bottom 5% realizations of the shocks during the first 5 years after the 6000 burn-in periods. Each observation is calculated using a window of 60 non-overlapping observations (5 years). The figure is averaged from 10,000 simulations with 1200 periods or 100 years per simulation.



5 Discussion

In this section, we discuss some of the key features of the model and compare it to related models.

5.1 Model Assumptions

Since we assume a continuum of agents with different beliefs, one has to condition on the whole trajectory of past shocks (and beliefs) to solve for the equilibrium. The key assumptions ensuring the tractability of the model are: 1) homogeneous endowments across cohorts, 2) log utility, and 3) that risky assets are in zero net supply. We now discuss the role for each

of these assumptions. First, since endowment is homogeneous across all agents, regardless of when they were born, the model rules out life-cycle profiles for the endowment process. Because markets are complete, homogeneous endowments imply that agents smooth their lifetime earnings across time and states. Importantly, homogeneous endowment has, when combined with log utility and a risky asset in zero net supply, the property that an agent is born with the same wealth as the total wealth in the economy. In fact, the present value of the endowment stream is the same for every agent. The only source of heterogeneity comes from the financial wealth, which by assumption is zero in the aggregate. Second, we assume that agents have log utility, which has the desired property that the propensity to consume out of wealth is constant. By imposing market clearing, this implies that the total wealth is proportional to the total endowment in the economy. This is also where the third assumption plays a crucial role. Because the risky asset is in zero net supply, the total wealth is simply the value of total endowment that agents are born with. Therefore, all the three assumptions are required for a tractable solution of the model economy.

5.2 Asset Structure

The assumption that the risky asset is in zero net supply is crucial for the solution of the model. While this is a restrictive assumption, its equilibrium implications are limited. Importantly, the discount factor does not depend on the particular form of the risky asset, as long as it completes the market. For the risky asset, we fix the diffusion coefficient exogenously and solve for the expected return. This is the same approach as in Basak (2000) in a heterogeneous belief setting and Karatzas, Lehoczky, and Shreve (1994) in a general multi-agent setting with homogeneous beliefs. Since the volatility is constant, the risk premium varies only due to variations in the market price of risk. An alternative approach, used in Duffie and Huang (1985), is to price an exogenously defined dividend stream, while still maintaining the assumption that the asset is in zero net supply. In this case, the stock price diffusion coefficient is determined endogenously. Hence, one must verify

that the resulting stock price dynamics completes the market, which, in general, is very difficult to perform without closed form solutions.¹⁴ One particular simple case is to assume that the dividend stream is proportional to the total endowment, Y_t . The price of this asset is $\alpha \frac{Y_t}{\rho}$, where α is the proportionality factor; its dynamics are

$$dG_t = dS_t + \alpha Y_t dt = S_t (\mu_t^S dt + \sigma^S dz_t), \quad (61)$$

where $\sigma^S = \sigma_Y$. For this case, we see that the asset price has the same dynamics as in our model with $\sigma^S = \sigma_Y$. Yet, in our setting we allow the risky asset to be more volatile than the endowment, which is consistent with the data. Another alternative, in the spirit of Abel (1999), is to assume that the risky asset is a claim to Y_t^α , where α is the leverage factor. The drawback of using the leverage factor is that there does not exist any closed form solution for the value of the claim in our setting. To calculate the expected return, one has to solve for the diffusion coefficient, σ_t^S . Applying Malliavin calculus, we see that the expression for σ_t^S is

$$\sigma_t^S = \alpha \sigma_Y + \frac{E_t \left[\int_t^\infty \xi_u Y_u^\alpha \left(\int_t^u \mathcal{D}_t \bar{\Delta}_s dz_s - \int_t^u \bar{\Delta}_s \mathcal{D}_t \bar{\Delta}_s ds \right) du \right]}{E_t \left[\int_t^\infty \xi_u Y_u^\alpha du \right]}, \quad (62)$$

where in the above $\mathcal{D}_t \bar{\Delta}_s$ denotes the Malliavin derivative of $\bar{\Delta}_s$ with $s > t$. Here, unlike in our case, the volatility of the risky asset is endogenously stochastic.

5.3 Return Extrapolation

In our model, agents learn about μ_Y by observing the endowment, Y_t . When an agent observes a shock to endowment that exceeds his expectation, the agent revise the expectation upwards. The behavior of agents and the resulting asset prices are such that the equilibrium appears as if agents are return extrapolating.¹⁵ To see this, define the instantaneous return

¹⁴For conditions for the financial market to be complete see Anderson and Raimondo (2008) and Hugonnier, Malamud, and Trubowitz (2010) in a single good setting and Ehling and Heyerdahl-Larsen (2015) in a multiple good setting.

¹⁵For an equilibrium model where return extrapolators coexist with rational agents see Barberis, Greenwood, Jin, and Shleifer (2015).

under the true probability measure as

$$dR_t = \frac{dS_t}{S_t} = \mu_t^S dt + \sigma^S dz_t. \quad (63)$$

Agents agree on the realized return, but disagree about its decomposition. For instance, the perceived return of a cohort born at time s is

$$dR_t = \mu_{s,t}^S dt + \sigma^S dz_{s,t}. \quad (64)$$

The expected return as perceived by cohort s is

$$\frac{1}{dt} E_{s,t}(dR_t) = \mu_{s,t}^S = \mu_t^S + \beta(t-s) \int_s^t (dR_u - E_u(dR_u)), \quad (65)$$

where $\beta(t-s) = \frac{\bar{V}}{\sigma_Y^2 + \bar{V}(t-s)}$. From Equation (65), we see that the return perceived by the cohort born at time s equates with the true expected return plus a term that depends on past returns relative to the true expected returns. Hence, whenever the actual return on the risky asset exceeds the true expectation, agents revise their expectations upwards. Effectively, they extrapolate returns. The strength of the return extrapolation is determined by $\beta(t-s)$. This term decreases when an agent ages; in the limit, it declines to zero. Consequently, young agents exhibit stronger return extrapolation than older, more experienced, agents. It is useful to express the true expected return in terms of past return surprises through

$$\mu_t^S = A - (\sigma^S - \sigma_Y) \int_{-\infty}^t f_{s,t} \left(\beta(t-s) \int_s^t (dR_u - E_u(dR_u)) \right) ds, \quad (66)$$

where $A = \rho + \mu_Y - \sigma_Y^2 + \sigma_Y \sigma^S$. For $\sigma^S > \sigma_Y$, we see that the true expected return depends negatively on the consumption share weighted average return extrapolation of the agents in the economy.¹⁶ The negative dependance of the expected return on past return surprises

¹⁶In Section 3, we set $\sigma^S = 15\%$ and $\sigma_Y = 2\%$. This is to reflect that stock returns are more volatile than fundamentals. If instead we focus on excess returns, the expected excess return always depends negatively on the consumption share weighted average return extrapolation.

captures the fact that the Sharpe ratio is high in bad times. In our model, all agents extrapolate returns relative to the true expected return. However, the total adjustment after a positive return surprise is negative for old agents. This follows from the fact that old agents have a lower sensitivity to past return surprises than the consumption share weighted average. Hence, after a negative return surprise the old agents revise their expectations upwards. In contrast, younger agents revise their expectations downwards after negative return surprises. Therefore, their expectation is negatively correlated with the true expected return.

5.4 Stationarity

In models where infinitely lived agents agree to disagree, except for knife edged cases, one class of agents takes over in the long run.¹⁷ In our model, since agents have finite life, consumption shares are not degenerate in the long-run. Still, market selection is at work because young agents with less precise beliefs lose out to older more experienced agents on average. Nevertheless, no cohort ever completely dominates the economy because new agents are constantly born and old agents die.

5.5 Learning versus Learning from Experience

Finally, we contrast our learning from experience channel to that of an infinitely lived representative agent who learns about the growth rate of aggregate consumption. The expressions for the risk free rate and the market price of risk in Proposition 5 are quite general. For instance, in a model with CRRA preferences where there is a difference between the belief of the representative agent (captured by the “market view” in our economy) and the econometrician, there is a wedge as captured by $\bar{\Delta}$. However, the $\bar{\Delta}$ dynamics differ. In a representative agent setting, the wedge converges to zero. Such a convergence does not occur

¹⁷See, for instance, Blume and Easley (1992), Sandroni (2000), Kogan, Ross, and Westerfield (2006), and Yan (2008).

in our setting because old agents die and new agents are born. Hence, the dynamics of the stochastic discount factor substantially differ from that of a representative agent who learns about the expected growth rate. This difference drives the predictability of excess returns. Our model replicates the negative correlation between the average belief and actual realizations of excess returns in the future. In an economy with an infinitely lived representative agent, there is no predictability. To demonstrate this, we inspect the market price of risk as perceived by the representative agent

$$\theta_t = \gamma\sigma_Y. \tag{67}$$

Indeed, Equation (67) shows that under the belief of the representative agent there is no predictability of excess returns. Therefore, in contrast to our economy with learning from experience, a representative agent economy cannot rationalize the negative relation between expected excess returns and realized excess returns, as documented empirically. Further, there is no correlation between past excess returns and the beliefs of the representative agent. Although the representative agent is extrapolating from fundamentals in the standard learning model, this does not translate into extrapolation from excess returns, as documented in the empirical literature. In our OLG economy with learning from experience, young agents extrapolate from returns. However, this is an equilibrium outcome stemming entirely from the interaction of young agents with old agents, due to heterogeneous beliefs, and does not arise through the extrapolation from fundamentals.

6 Conclusions

We make a modest departure from rational expectations concerning how young agents learn from data: First, to isolate the influence of experience, we assume that all agents share the true prior about consumption growth. Second, agents are unsure about their estimate for growth and not all data is used for learning about consumption growth. Therefore, early

on in life new data can have a large impact on expectations. Hence, the young are more likely to be overly optimistic or pessimistic than the old and experienced. Specifically, these two assumptions imply in an overlapping generations economy with optimal learning, given beliefs, that old agents have more precise estimates than younger agents and their beliefs react less to news, the young act as trend chasers, wealth shifts from young to old, the market price of risk is countercyclical and the risk-free rate of return is procyclical, the average belief about the stock return is negatively correlated with and less volatile than the actual expected return, and experiencing especially bad outcomes when young leads to significant and persistent disagreement in the economy and to a persistent tilt in portfolio weights.

A Proofs of Propositions

A.1 Proof of Proposition 1

Following standard filtering theory, Liptser and Shiryaev (1974a,b), the dynamics of the expected consumption growth as perceived by an agent born at time s are given by Equation (3). Defining the disagreement process as

$$\Delta_{s,t} = \frac{\hat{\mu}_{s,t} - \mu_Y}{\sigma_Y}, \quad (68)$$

and applying Ito's lemma to it, we have

$$d\Delta_{s,t} = -\frac{\bar{V}}{\sigma_Y^2 + \bar{V}(t-s)}\Delta_{s,t}dt + \frac{\bar{V}}{\sigma_Y^2 + \bar{V}(t-s)}dz_t. \quad (69)$$

The solution to this stochastic differential equation is found by applying Ito's lemma to

$$\Delta_{s,t} = \frac{\bar{V}(z_t - z_s)}{\sigma_Y^2 + \bar{V}(t-s)}, \quad (70)$$

which, then, yields the desired result. By the strong law of large numbers, we have that

$$\lim_{t \rightarrow \infty} \frac{z_t}{t} = 0, \quad (71)$$

and hence $\lim_{t \rightarrow \infty} \Delta_{s,t} = 0$.

A.2 Proof of Proposition 2

We solve for equilibrium using the martingale method, Cox and Huang (1989). An agent born at time s solves the following static optimization problem

$$\begin{aligned} \max_{c_s} E_{s,s} \left[\int_s^\infty e^{-(\rho+\nu)(t-s)} \log(c_{s,t}) dt \right] \\ \text{s.t.} \\ E_{s,s} \left[\int_s^\infty e^{-\nu(t-s)} \xi_{s,t} c_{s,t} dt \right] = E_{s,s} \left[\int_s^\infty e^{-\nu(t-s)} \xi_{s,t} y_{s,t} dt \right]. \end{aligned}$$

From the first order conditions (FOCs), we have

$$\frac{e^{-(\rho+\nu)(t-s)}}{c_{s,t}} = \kappa_s e^{-\nu(t-s)} \xi_{s,t}, \quad (72)$$

where κ_s denotes the Lagrange multiplier. Note that for $s \leq t$ the FOCs imply

$$e^{-(\rho+\nu)(t-s)} \left(\frac{c_{s,s}}{c_{s,t}} \right) = e^{-\nu(t-s)} \frac{\xi_{s,s}}{\xi_{s,t}}. \quad (73)$$

Rearranging leads to

$$c_{s,t} = c_{s,s} e^{-\rho(t-s)} \frac{\xi_{s,s}}{\xi_{s,t}}. \quad (74)$$

Using Equation (23) and the Radon-Nikodym derivative to move from the probability measure of an agent born at time s to the actual probability measure, we get the optimal consumption

at time t of an agent born at time $s \leq t$

$$c_{s,t} = Y_s e^{-(\rho+\nu)(t-s)} \frac{\eta_{s,t} \xi_s}{\eta_{s,s} \xi_t}, \quad (75)$$

where

$$\eta_{s,t} = \frac{\xi_t}{\xi_{s,t}}. \quad (76)$$

A.3 Proof of Proposition 3

The expression for the equilibrium stochastic discount factor is obtained taking the market clearing condition for the goods market, Equation (15), and plugging the optimal consumption at time t of an agent born at time s , Equation (24), into it. The resulting expression is

$$Y_t = \int_{-\infty}^t \nu e^{-(\rho+\nu)(t-s)} Y_s \frac{\xi_s}{\xi_t} \frac{\eta_{s,t}}{\eta_{s,s}} ds. \quad (77)$$

Then, defining

$$X_t = \int_{-\infty}^t \nu e^{-(\rho+\nu)(t-s)} Y_s \xi_s \frac{\eta_{s,t}}{\eta_{s,s}} ds, \quad (78)$$

and rearranging Equation (77) leads to

$$\xi_t = \frac{1}{Y_t} \int_{-\infty}^t \nu e^{-(\rho+\nu)(t-s)} Y_s \xi_s \frac{\eta_{s,t}}{\eta_{s,s}} ds. \quad (79)$$

Substituting into the integrand above the expression $X_s = Y_s \xi_s$ yields the result.

A.4 Proof of Proposition 4

To prove Proposition 4, we start by obtaining the dynamics of X_t . Applying Ito's lemma to Equation (26), we have

$$\frac{dX_t}{X_t} = -\rho dt + \int_{-\infty}^t \nu e^{-(\rho+\nu)(t-s)} \frac{X_s}{X_t} \frac{d\eta_{s,t}}{\eta_{s,s}} ds = -\rho dt + \bar{\Delta}_t dz_t \quad (80)$$

where

$$\bar{\Delta}_t = \int_{-\infty}^t f_{s,t} \Delta_{s,t} ds, \quad (81)$$

and where

$$f_{s,t} = \nu e^{-\nu(t-s)} \frac{Y_s e^{-\rho(t-s)} \left(\frac{\eta_{s,t}}{\eta_{s,s}} \right) \left(\frac{\xi_s}{\xi_t} \right)}{Y_t} = \nu e^{-(\rho+\nu)(t-s)} \left(\frac{X_s}{X_t} \right) \left(\frac{\eta_{s,t}}{\eta_{s,s}} \right) = \nu e^{-(\rho+\nu)(t-s)} \frac{C_{s,t}}{Y_t}, \quad (82)$$

which represents the share of aggregate output at time t that accrues to agents born at time s .

X_t has dynamics satisfying the above stochastic differential equation with well-known analytic solution

$$X_t = X_0 e^{-\int_0^t \left(\rho + \frac{\bar{\Delta}_s^2}{2} \right) ds + \int_0^t \bar{\Delta}_s dz_s}. \quad (83)$$

Assuming that $X_0 = 1$, the expression for X_t can be written as

$$X_t = e^{-\rho t} \bar{\eta}_t, \quad (84)$$

where

$$\bar{\eta}_t = e^{-\frac{1}{2} \int_0^t \bar{\Delta}_s^2 ds + \int_0^t \bar{\Delta}_s dz_s}. \quad (85)$$

Applying Ito's lemma to (85) gives the dynamics of $\bar{\eta}_t$,

$$d\bar{\eta}_t = \bar{\Delta}_t \bar{\eta}_t dz_t. \quad (86)$$

Equations (84) and (86) yield the required results.

A.5 Proof of Proposition 5

From the expression of the equilibrium stochastic discount factor in Equation (25), we have

$$\xi_t = \frac{X_t}{Y_t}. \quad (87)$$

Using Equation (84), the stochastic discount factor can be decomposed in the following way

$$\xi_t = \bar{\eta}_t \frac{e^{-\rho t}}{Y_t}. \quad (88)$$

Then, by applying Ito's lemma to Equation (88), we have

$$d\xi_t = d\left(\bar{\eta}_t \frac{e^{-\rho t}}{Y_t}\right). \quad (89)$$

Using Equations (2) and (86), we get

$$d\left(\bar{\eta}_t \frac{e^{-\rho t}}{Y_t}\right) = \left(\bar{\eta}_t \frac{e^{-\rho t}}{Y_t}\right) [(-\rho - \mu_Y + \sigma_Y^2 - \sigma_Y \bar{\Delta}_t) dt - (\sigma_Y - \bar{\Delta}_t) dz_t]. \quad (90)$$

Matching the drift and diffusion terms of the state price density, Equation (9), with Equation (90), we obtain

$$r_t = \rho + \mu_Y - \sigma_Y^2 + \sigma_Y \bar{\Delta}_t \quad (91)$$

and

$$\theta_t = \sigma_Y - \bar{\Delta}_t. \quad (92)$$

These are the equilibrium real short rate and market price of risk, respectively.

A.6 Proof of Proposition 6

To prove Proposition 6, we need to compute the dynamics of $\bar{\Delta}_t$. Applying Ito's lemma to Equation (29) gives

$$d\bar{\Delta}_t = \frac{\partial \bar{\Delta}_t}{\partial f_{s,t}} df_{s,t} + \frac{\partial \bar{\Delta}_t}{\partial \Delta_{s,t}} d\Delta_{s,t} + \frac{\partial \bar{\Delta}_t}{\partial f_{s,t} \partial \Delta_{s,t}} (df_{s,t} d\Delta_{s,t}) + \frac{1}{2} \frac{\partial \bar{\Delta}_t}{\partial f_{s,t}^2} (df_{s,t}^2) + \frac{1}{2} \frac{\partial \bar{\Delta}_t}{\partial \Delta_{s,t}^2} (d\Delta_{s,t}^2). \quad (93)$$

Thus, to evaluate Equation (93), we need the dynamics of $f_{s,t}$ and $\Delta_{s,t}$. From Equation (29), using Ito's lemma, we obtain

$$df_{s,t} = f_{s,t}[(-\nu + \bar{\Delta}_t^2 - \bar{\Delta}_t \Delta_{s,t}) dt + (\Delta_{s,t} - \bar{\Delta}_t) dz_t]. \quad (94)$$

Next, we compute from Equation (3), by making use of the fact that $\Delta_{s,t} = \frac{\hat{\mu}_{s,t} - \mu_Y}{\sigma_Y}$, the following

$$d\Delta_{s,t} = -\Delta_{s,t} \frac{V_{s,t}}{\sigma_Y^2} dt + \frac{V_{s,t}}{\sigma_Y^2} dz_t. \quad (95)$$

Inserting the dynamics of $f_{s,t}$ and $\Delta_{s,t}$ into Equation (93) and focusing on the diffusion terms, we get

$$d\bar{\Delta}_t = \dots dt + \int_{-\infty}^t \Delta_{s,t} f_{s,t} (\Delta_{s,t} - \bar{\Delta}_t) ds dz_t + \int_{-\infty}^t f_{s,t} \left(\frac{V_{s,t}}{\sigma_Y^2} \right) ds dz_t. \quad (96)$$

The last term of this equation is intuitive. In case of a positive (negative) shock to aggregate output, i.e., $dz_t > 0 (< 0)$, it is always positive (negative). Notice that the first term can be written as

$$\left(\int_{-\infty}^t \Delta_{s,t}^2 f_{s,t} ds - \bar{\Delta}_t^2 \right) dz_t. \quad (97)$$

It follows from the Jensen's inequality that

$$\left(\int_{-\infty}^t \Delta_{s,t}^2 f_{s,t} ds - \bar{\Delta}_t^2 \right) \geq 0. \quad (98)$$

Thus, the integrand will always be greater or equal to zero, hence, the entire term will be greater or equal to zero. Consequently, a positive (negative) shock to aggregate output, i.e., $dz_t > 0 (< 0)$, causes the consumption share weighted average disagreement in the economy to increase (decrease), $d\bar{\Delta}_t > 0 (< 0)$. This drives up (down) the interest rate and down (up) the market price of risk. This proves the proposition.

A.7 Proof of Proposition 7

Applying Ito's lemma to Equation (24), yields the dynamics for optimal consumption,

$$dc_{s,t} = c_{s,t} [(-\rho + r_t + \theta_t^2 + \Delta_{s,t}\theta_t) dt + (\theta_t + \Delta_{s,t}) dz_t]. \quad (99)$$

Substituting the optimal market price of risk, Equation (33), and real short rate, Equation (32), into the PDE above, we get

$$dc_{s,t} = c_{s,t} (\mu_{c_{s,t}} dt + \sigma_{c_{s,t}} dz_t), \quad (100)$$

where

$$\mu_{c_{s,t}} = \mu_Y + (\sigma_Y - \bar{\Delta}_t) (\Delta_{s,t} - \bar{\Delta}_t), \quad \sigma_{c_{s,t}} = \sigma_Y + \Delta_{s,t} - \bar{\Delta}_t, \quad (101)$$

which yields the result.

A.8 Proof of Proposition 8

To derive the optimal portfolio allocation at time t for an agent born at time s , we start with

$$\hat{W}_{s,t} = \frac{1}{\xi_t} E_t \left[\int_t^\infty e^{-\nu(u-t)} \xi_u c_{s,u} du \right]. \quad (102)$$

Using

$$\hat{W}_{s,t} = \frac{c_{s,t}}{\rho + \nu}, \quad (103)$$

substituting in optimal consumption, Equation (24), and after rearranging terms, we have

$$\xi_t \hat{W}_{s,t} = Y_s e^{-(\rho+\nu)(t-s)} \frac{\eta_{s,t}}{\eta_{s,s}} \frac{\xi_s}{\rho + \nu}. \quad (104)$$

Applying Ito's lemma to both sides of the above equation and matching the diffusion terms, leads to

$$\xi_t \left(\pi_{s,t} \sigma^S + H_{s,t} \sigma_Y - \hat{W}_{s,t} \theta_t \right) = \Delta_{s,t} Y_s e^{-(\rho+\nu)(t-s)} \frac{\eta_{s,t}}{\eta_{s,s}} \frac{\xi_s}{\rho + \nu}. \quad (105)$$

Rearranging and substituting the terms yields

$$\pi_{s,t} \sigma^S = \hat{W}_{s,t} (\sigma_Y - \bar{\Delta}_t + \Delta_{s,t}) - H_{s,t} \sigma_Y. \quad (106)$$

Finally, solving for the optimal portfolio, $\pi_{s,t}$,

$$\pi_{s,t} = \frac{\Delta_{s,t} - \bar{\Delta}_t}{\sigma^S} \hat{W}_{s,t} + \frac{\sigma_Y}{\sigma^S} W_{s,t}, \quad (107)$$

yields the result.

References

- Abel, A. B., 1999, "Risk Premia and Term Premia in General Equilibrium," *Journal of Mathematical Economics*, 43, 3–33.
- Agarwal, S., J. C. Driscoll, X. Gabaix, and D. Laibson, 2009, "The Age of Reason: Financial Decisions over the Life Cycle and Implications for Regulation," *Brookings Papers on Economic Activity*, (Fall), 51–117.
- Anderson, R. M., and R. C. Raimondo, 2008, "Equilibrium in Continuous-Time Financial Markets: Endogenously Dynamically Complete Markets," *Econometrica*, 76, 841–907.
- Arrondel, L., H. Calvo-Pardo, and D. Tas, 2014, "Subjective Return Expectations, Information and Stock Market Participation: Evidence from France," *Paris School of Economics-Banque de France and University of Southampton*, 1–59.
- Atmaz, A., 2015, "A Dynamic Model of Short Interest," *London Business School*, 1–46.

- Barberis, N., R. Greenwood, L. Jin, and A. Shleifer, 2015, "X-CAPM: An Extrapolative Capital Asset Pricing Model," *Journal of Financial Economics*, 115, 1–24.
- Basak, S., 2000, "A Model of Dynamic Equilibrium Asset Pricing with Heterogeneous Beliefs and Extraneous Risk," *Journal of Economic Dynamics and Control*, 24, 63–95.
- Basak, S., 2005, "Asset Pricing with Heterogeneous Beliefs," *Journal of Banking and Finance*, 29, 2849–2881.
- Bhamra, H. S., and R. Uppal, 2014, "Asset Prices with Heterogeneity in Preferences and Beliefs," *Review of Financial Studies*, 27, 519–580.
- Blanchard, O., 1985, "Debt, Deficits, and Finite Horizons," *Journal of Political Economy*, 93, 223–247.
- Blume, L., and D. Easley, 1992, "Evolution and Market Behavior," *Journal of Economic Theory*, 58, 9–40.
- Buss, A., R. Uppal, and G. Vilkov, 2015, "Where Experience Matters: Asset Allocation and Asset Pricing with Opaque and Illiquid Assets," *INSEAD, Edhec Business School, and Frankfurt School of Finance & Management*, 1–59.
- Campbell, J. Y., and J. H. Cochrane, 1999, "By Force of Habit: A Consumption Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107, 205–251.
- Cogley, T., and T. Sargent, 2008, "The Market Price of Risk and the Equity Premium: A Legacy of the Great Depression," *Journal of Monetary Economics*, 55, 454–476.
- Collin-Dufresne, P., M. Johannes, and L. A. Lochstoer, 2015, "Asset Pricing when 'This Time is Different'," *University of Lausanne and Columbia Business School*, 1–46.
- Cox, J., and C. Huang, 1989, "Optimal Consumption and Portfolio Policies when Asset Prices Follow a Diffusion Process," *Journal of Economic Theory*, 83, 33–83.

- Cvitanic, J., E. Jouini, S. Malamud, and C. Napp, 2012, “Financial Markets Equilibrium with Heterogenous Agents,” *Review of Finance*, 16, 285–321.
- Cvitanic, J., and S. Malamud, 2011, “Price Impact and Portfolio Impact,” *Journal of Financial Economics*, 100, 201–225.
- Detemple, J. B., and S. Murthy, 1994, “Intertemporal Asset Pricing with Heterogeneous Beliefs,” *Journal of Economic Theory*, 62, 294–320.
- Duffie, D., and C. Huang, 1985, “Implementing Arrow-Debreu Equilibria by Continuous Trading of Few Long-lived Securities,” *Econometrica*, 53, 1337–1356.
- 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., M. Gallmeyer, C. Heyerdahl-Larsen, and P. Illeditsch, 2015, “Disagreement about Inflation and the Yield Curve,” *BI Norwegian Business School, University of Virginia (UVA) - McIntire School of Commerce, London Business School, and Wharton School*, 1–59.
- Ehling, P., and C. Heyerdahl-Larsen, 2015, “Complete and Incomplete Financial Markets in Multi-Good Economies,” *Journal of Economic Theory*, 160, 438–462.
- Fedyk, Y., C. Heyerdahl-Larsen, and J. Walden, 2013, “Market Selection in a Multi-Asset Economy,” *Review of Finance*, 17, 1179–1237.
- Gârleanu, N., and S. Panageas, 2015, “Young, Old, Conservative and Bold: The Implications of Heterogeneity and Finite Lives for Asset Pricing,” *Journal of Political Economy*, 123, 670–685.
- Greenwood, R., and S. Nagel, 2009, “Inexperienced Investors and Bubbles,” *Journal of Financial Economics*, 93, 239–258.

- Greenwood, R., and A. Shleifer, 2014, “Expectations of Returns and Expected Returns,” *Review of Financial Studies*, 27, 714–746.
- Harrison, J., and D. Kreps, 1978, “Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations,” *Quarterly Journal of Economics*, 92, 323–336.
- Hugonnier, J., S. Malamud, and E. Trubowitz, 2010, “Endogenous Completeness of Diffusion Driven Equilibrium Markets,” *Econometrica*, 80, 1249–1270.
- Karatzas, I., J. P. Lehoczky, and S. E. Shreve, 1994, “Existence and uniqueness of multi-agent equilibrium in a stochastic, dynamic consumption/investment model,” *Mathematics of Operations Research*, 15, 80–128.
- Knüpfer, S., E. Rantapuska, and M. Sarvimäki, 2014, “Labor Market Experiences and Portfolio Choice: Evidence from the Finnish Great Depression,” *London Business School and Aalto University*, 1–44.
- Kogan, L., W. J. Ross, S. A., and M. M. Westerfield, 2006, “The Price Impact and Survival of Irrational Traders,” *The Journal of Finance*, 61, 195–229.
- Koijen, R. S. J., M. Schmeling, and E. B. Vrugt, 2015, “Survey Expectations of Returns and Asset Pricing Puzzles,” *London Business School, City University London, and VU University Amsterdam*, 1–66.
- Liptser, S., and N. Shiryaev, 1974a, *Statistics of Random Processes I General Theory* (2d ed.), Springer.
- Liptser, S., and N. Shiryaev, 1974b, *Statistics of Random Processes II Applications* (2d ed.), Springer.
- Malmendier, U., and S. Nagel, 2011, “Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?” *Quarterly Journal of Economics*, 126, 373–416.

- Malmendier, U., and S. Nagel, 2014, “Learnign from Inflation Experiences,” *UC Berkeley and Stanford University*, 1–37.
- Malmendier, U., D. Pouzo, and V. Vanasco, 2015, “A Theory of Experience Effects,” *UC Berkeley and Stanford*, 1–57.
- Piazzesi, M., J. Salomao, and M. Schneider, 2015, “Trend and Cycle in Bond Premia,” *Stanford and University of Minnesota*, 1–27.
- Sandroni, A., 2000, “Do Markets Favor Agents able to Make Accurate Predictions?” *Econometrica*, 68(6), 1303–1341.
- Schraeder, S., 2015, “Information Processing and Non-Bayesian Learning in Financial Markets,” *HEC Lausanne, forthcoming Review of Finance*, 1–31.
- Xiong, W., and H. Yan, 2010, “Heterogeneous Expectations and Bond Markets,” *Review of Financial Studies*, 23, 1433–1466.
- Yaari, M. E., 1965, “Uncertain Lifetime, Life Insurance, and the Theory of the Consumer,” *Review of Economic Studies*, 32, 137–150.
- Yan, H., 2008, “Natural Selection in Financial Markets: Does It Work?” *Management Science*, 54(11), October, 1935–1950.
- Zapatero, F., 1998, “Effects of Financial Innovations on Market Volatility when Beliefs are Heterogeneous,” *Journal of Economic Dynamics and Control*, 22, 597–626.