DISCUSSION PAPER SERIES

No. 9845

LEARNING FROM EXPERIENCE IN THE STOCK MARKET

Anton Nakov and Galo Nuño

FINANCIAL ECONOMICS and INTERNATIONAL MACROECONOMICS

Centre for Economic Policy Research

www.cepr.org

Available online at: **www.cepr.org/pubs/dps/DP9845.php**

ISSN 0265-8003

LEARNING FROM EXPERIENCE IN THE STOCK MARKET

Anton Nakov, European Central Bank and CEPR
Galo Nuño, European Central Bank

Discussion Paper No. 9845
February 2014

Centre for Economic Policy Research
77 Bastwick Street, London EC1V 3PZ, UK
Tel: (44 20) 7183 8801, Fax: (44 20) 7183 8820
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 9845

February 2014

ABSTRACT

Learning from Experience in the Stock Market*

New evidence suggests that individuals "learn from experience," meaning they learn from events occurring during their own lifetimes as opposed to the entire history of events. Moreover, they weigh more heavily the more recent events compared to events occurring in the more distant past. This paper analyzes the implications of such learning for stock pricing in a model with finitely-lived agents. Individuals learn about the rate of change of the stock price and of dividends using a weighted decreasing-gain algorithm. Information is dispersed across age cohorts with older agents having larger information sets than younger ones. In the model, the stock price exhibits stochastic fluctuations around the rational expectations equilibrium due to successive waves of optimism and pessimism. We demonstrate how this heterogeneous-beliefs model can be approximated by an economy with a representative agent who updates his beliefs following a constant-gain learning scheme. The aggregate gain parameter of the approximation is a nonlinear function of the survival rate and of the individual gain parameters.

JEL Classification: D83, D84 and G12
Keywords: asset pricing, constant-gain learning, dispersed beliefs, heterogeneous beliefs and OLG

Anton Nakov
European Central Bank
DG-Research
Frankfurt
GERMANY

Email: anton.nakov@gmail.com

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=160916

Galo Nuño
European Central Bank
Kaiserstrasse 29
D-60311 Frankfurt
GERMANY

Email: galo.nuno@ecb.int

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=167674

*We are grateful for comments and suggestions to Klaus Adam, George-Marios Angeletos, Gunter Coenen, Alistair Dieppe, Oliver de Groot, Cars Hommes, Miguel Angel Iraola, Christian Matthes, Albert Marcet, Alessio Moro, Giulio Nicoletti, Kristoffer Nimark, John Roberts, and seminar participants at ECB, Bank of Spain, University of Surrey, and University of Cagliari, and conference participants at the Society for Computational Economics 2011 and the Bank of Spain Workshop on Expectations in Macroeconomics 2013. Anton Nakov thanks the Bank of Spain and the Federal Reserve Board for their hospitality during earlier drafts of this paper. The views expressed here are those of the authors only and do not necessarily coincide with the views of the ECB or the Eurosystem.

Submitted 21 January 2014

1 Introduction

The key role of expectations about the future is well understood in economics. The rational expectations hypothesis (REH) has been a major step forward allowing rigorous formalization of the process of expectations formation. Yet it has often been criticized for endowing people with "too much" knowledge about their environment.[1] Empirical studies of individuals forming expectations about aggregate economic variables does not, in general, corroborate the REH. In particular, Malmendier and Nagel (2011, 2013) find evidence that, contrary to the REH, people "learn from experience," meaning that they are more influenced by observations from their own lifetimes than by earlier historical events. More specifically, Malmendier and Nagel (2011) find that individuals who experienced low stock market returns during their lives are more pessimistic about future stock returns and invest a lower fraction of their liquid assets in stocks. In addition, Malmendier and Nagel (2013) find that young individuals place more weight on recently experienced inflation than older individuals do. The by-product of this is that learning may take forever if history "gets lost" as new generations replace older ones.

In this paper, we explore how replacing the REH with "learning from experience" affects the dynamics of a simple general equilibrium model of the stock market. We are interested in the dynamics of heterogeneous beliefs and in the feedback loop that arises when individuals learn about variables which are the result of their collective actions given their beliefs.[2] To this end, we extend the basic Lucas-tree asset pricing model to a stochastic Blanchard-Yaari overlapping generations (OLG) setup in which individuals learn the parameters of the *endogenous* evolution of the stock price as well as the *exogenous* process for dividends. The equilibrium we study is "internally rational" in the sense of Adam and Marcet (2011), that is, agents in it maximize expected utility given their subjective probability distributions.

Specifically, we assume that a small random fraction $1 - \phi$ (with $\phi \lesssim 1$) of individuals exit the stock market every period, and an equal measure of new individuals enter the market. As in Brown and Rogers (2009), each new entrant inherits the assets but does not inherit the accumulated knowledge of his parent about the economy. Instead, children learn from their own experience, updating their beliefs with information about stock prices and dividends which they observe during their own lifetimes.[3] As in Malmendier and Nagel (2013), agents use a decreasing-gain learning scheme with gain parameter θ. A value of $\theta = 1$ implies that individuals assign the same weight to all observations they witness. In contrast, if $\theta > 1$, as the evidence of Malmendier and Nagel suggests, then individuals weigh more heavily recent events compared to older events.

[1]See, for example, Blume et. al. (1982), Arrow (1986), and Adam and Marcet (2011).

[2]See Eusepi and Preston (2011) who emphasize this type of self-referentiality.

[3]Thus, information in the model is dispersed across age cohorts, with older generations observing longer time series than younger ones. See Angeletos and La'O (2009, 2013) for the role of *geographically* dispersed information for macroeconomic dynamics.

To analyze the model quantitatively, we propose a method that allows us to solve for the equilibrium with heterogeneous agents. In the model, the equilibrium stock price equals the reservation price of the marginal stock holder. Individuals who are more optimistic than the marginal agent choose to hold the maximum possible amount of the stock, while those who are more pessimistic choose not to hold the stock. Since the asset holding decision of each individual depends on the current stock price (which is used to forecast future prices), the solution method involves finding a fixed-point for the market-clearing price given the non-linear pricing function.

The introduction of "learning from experience" has several novel implications. First, we find that, even if the retirement rate $1-\phi$ is quite low, so that in any given period only a small fraction of individuals are novice, the asset price fails to converge to the rational expectations equilibrium (REE). Two forces create the oscillating dynamics. On the one hand, there is "momentum" rooted in the information loss due to the retirement of individuals from the market and to the stronger discounting of older information relative to the more recent one. As a result, beliefs about dividends and the stock price are biased towards extrapolation of the more recent past, and trading on these beliefs pushes the asset price further away from the fundamental.

On the other hand, there is a force of reversal toward the REE trend. Namely, when the stock price rises too far above the fundamental value, individual asset exposure constraints begin to bind. Because any given individual (including the optimistic types) can afford to buy less of the stock, the asset price must decline to the valuation of less optimistic individuals for the market to clear. The same reflecting force works also "from below", when the stock price falls far below the fundamental value. The combination of these two factors – momentum and trend reversal – results in boom-and-bust cycles, which are only loosely related to dividends and are mainly due to speculation about the future course of the stock price, in the spirit of Harrison and Kreps (1978).

A second finding is that, although individual expectations in our framework are not model-consistent, the agents' forecasting performance is quite close to that under rational expectations. In particular, the mean forecasting error (averaged across cohorts) is similar to that in the rational expectations model. Likewise, the root mean square errors (again averaged across cohorts) are not too far from the ones obtained under rational expectations.

A third finding of our paper is that the heterogeneous-beliefs economy can be approximated reasonably well by an economy with a representative agent who updates his beliefs with a constant-gain learning (CGL) scheme. The approximation involves two steps. In a first step, we show that the dynamics of average beliefs can be approximated by a CGL scheme in which the social gain parameter is a nonlinear function of the survival rate ϕ and of the individual gain parameter θ. This implies that memories of the distant past are lost with the passage of time as a result of population turnover combined with "learning from experience." In a second step, we show that the evolution of the stock price can be approximated using the evolution of the average (rather than the marginal) beliefs of the population.

Usually CGL is derived from the assumption that a representative agent uses the Kalman filter as in Ljung and Soderstrom (1983), Sargent (1999, ch. 8) or McCulloch (2007). This type of learning has received much attention in the literature due to its ability to produce realistic model features, such as amplification of the persistence of macroeconomic variables in response to aggregate shocks.[4] Yet Malmendier and Nagel (2013) show empirically that even though individual learning follows a decreasing-gain scheme, the average learning-from-experience forecast can be approximated quite closely with a CGL algorithm. Complementary to their work, we propose a theoretical model of stock pricing which reproduces this feature and we provide an expression for the approximation error. We verify numerically that in our model the approximation error is relatively small. We further show that the social gain parameter is increasing in the individual gain parameter θ, and in the rate of generational turnover, $1-\phi$.

Finally, we compare the behavior of the price-dividend ratio in the heterogeneous-agents model with its behavior in a representative-agent CGL approximation. We find that the representative agent model generates more volatility in the price-dividend ratio, but it exhibits a cyclical pattern which is broadly similar to that of the heterogeneous-agents model.

Our paper is related to several strands of research. First, it relates to the emerging literature on learning with heterogeneous agents, such as Cogley, Sargent and Tsyrennikov (2012), Giannitsarou (2003), Branch and McGough (2004), Branch and Evans (2006), Honkapohja and Mitra (2006), and Graham (2011). In contrast to these papers, individuals in our economy use the same decreasing-gain learning scheme, have the same preferences, and observe the same public variables. The only source of heterogeneity in our model is in the individual information sets used to update beliefs, with younger cohorts focusing on a subset of the observations used by older generations.

Second, our work is related to the literature on bounded rationality with heterogeneous beliefs, following Brock and Hommes (1998). In this literature agents switch between heterogeneous expectations based on the short-run profitability of the investment strategies. These models have been estimated on different financial time series as in Boswijk, Hommes and Manzan (2007).[5]

Third, a related line of research analyzes the dynamics of asset prices under learning by a representative agent. Timmermann (1994), Weitzman (2007), and Cogley and Sargent (2008), among others, explain some puzzling asset pricing phenomena based on rational learning by a representative agent. Unlike our setup, individuals in their models use all available past information and know *ex ante* the correct mapping between asset prices and fundamentals. Hence, they only need to learn about the latter in order to achieve convergence to the REE.

[4]The value of the gain parameter typically is estimated or calibrated to yield the smallest mean-squared forecasting error. See, e.g., Milani (2007), Carceles-Poveda and Giannitsarou (2008), Branch and Evans (2011), Adam, Marcet, and Nicolini (2008).

[5]For a survey, see Hommes (2006).

Finally, a recent line of research focuses on the role of higher-order expectations for asset prices. For example, Allen, Morris, and Shin (2006) analyze a linear model with asymmetric information. They find that, in the absence of common knowledge about higher-order beliefs, asset prices generally depart from the market consensus of the expected fundamental value, typically reacting more sluggishly to changes in fundamentals.

The rest of our paper is organized as follows. Section 2 presents the model. In section 3, we calibrate it and analyze the properties of "learning from experience." In section 4, we show to what extent the benchmark model can be approximated by a representative agent with CGL. Conclusions are presented in section 5.

2 The model

The economy is populated by N risk-neutral ex-ante identical dynasties, with N large. Members of each dynasty have stochastic lifetimes with death (or retirement) occurring with a constant exogenous probability, $1-\phi$. Thus, in each period, the number of dynasts of age $s \in \mathbb{N}_0$ is constant and equal to $f_s = N(1-\phi)\phi^s$. Upon retirement, a successor inherits the assets of the former dynast but not his accumulated knowledge about the processes governing the stock price and dividends. Instead, successors embark on their own learning experience "from scratch", starting with the identical initial belief that their predecessors had at birth, namely the belief consistent with REE.

The dynasts trade among themselves a single divisible stock which is in fixed supply, normalized to N. Each individual decides how much to invest in the asset based on inter-temporal arbitrage. Note, however, that the relevant arbitrage is not the one between selling the stock and holding it forever for its dividends. Instead, the condition that governs savings decisions is a one-period-ahead comparison between the value of the stock in the current period and the subjective expected payoff in the following trading period.

The equilibrium stock price in our model equals the *marginal* asset holder i's *subjectively* expected payoff from holding the stock for one period – that is, the present value of *his* expected dividend $E_{it}(D_{t+1})$ plus *his* expected price $E_{it}(P_{t+1})$ in the following period. Because expectations about future prices and dividends differ across individuals, the law of iterated expectations does not apply, and the pricing conditions of individuals do not aggregate to the familiar asset pricing formula with a representative agent.

2.1 Preferences and constraints

The head of dynasty $i \in \{1, ..., N\}$ receives utility from consumption $u(C_{it}) = C_{it}$ per period. He discounts future consumption by factor $\beta\phi$, where $\beta < 1$ is a time preference parameter and $\phi < 1$ is a constant probability of survival. The expected value of lifetime utility for dynast i is thus

$$E_{i0} \sum_{t=0}^{\infty} (\beta\phi)^t \, u(C_{it}), \tag{1}$$

where E_{i0} is individual i's expectation formed at time 0.

Individual i faces the period budget constraint

$$C_{it} + P_t S_{it} \leq (P_t + D_t) S_{it-1} + Y_{it}, \tag{2}$$

where S_{it} denotes his stock holdings, P_t is the asset price, D_t is the dividend, and Y_{it} is a per period income endowment. We assume for simplicity that $Y_{it} = Y$.

In addition, the individual faces constraints on the minimum and the maximum *asset exposure*, defined as the maximum value in terms of consumption that he stands to lose (or gain if short-selling) if the stock price falls to zero.

$$\underline{L}_t \leq P_t S_{it} \leq \bar{L}_t. \tag{3}$$

Constraints (3) imply that an individual investor cannot go arbitrarily short or long in the stock. In a more detailed model, these limitations can be derived from underlying credit constraints that prevent agents from borrowing unlimited amounts of resources. Instead, we will simply assume that $\underline{L}_t = 0$ and $\bar{L}_t = \lambda D_t > 0$, where parameter $\lambda > 0$ (which we loosely refer to as the permissible "leverage") is the maximum multiple of the current dividend that an individual can maintain invested in the risky stock. Our specification of the stock holding constraints puts effective bounds on the price-to-dividend ratio, without the need for a "projection facility" that mechanically constrains beliefs to a pre-specified neighborhood.

Dividends follow the exogenous stochastic process

$$\log\left(D_t/D_{t-1}\right) = \mu + \varepsilon_t, \quad \varepsilon_t \sim N(0, \sigma^2), \tag{4}$$

where $\mu > 0$ and $\sigma^2 > 0$ are, respectively, the mean and the variance of the growth rate of dividends and where D_{-1} is known.

Given the information set available to individual i, his problem is to choose consumption and equity holdings so as to maximize lifetime utility (1), subject to the budget constraint (2). The first-order optimality conditions (FOC) of the individual's problem are:

$$\text{if } P_t < P_{it}, \text{ then } S_{it} = \bar{L}_t / P_t \tag{5a}$$

$$\text{if } P_t = P_{it}, \text{ then } S_{it} \in [\underline{L}_t / P_t, \bar{L}_t / P_t], \tag{5b}$$

$$\text{if } P_t > P_{it}, \text{ then } S_{it} = \underline{L}_t / P_t, \tag{5c}$$

$\forall t$, where

$$P_{it} = \beta\phi E_{it} \left(P_{t+1} + D_{t+1} \right), \tag{6}$$

is individual i's "reservation price". Because the objective function is linear and the feasible set is closed, a maximum exists (and generally is a corner solution). We assume that Y is large enough so that the condition $C_{it} \geq 0$ is never binding.

The FOC can also be written as

$$P_t = \beta\phi E_{it} \left(P_{t+1} + D_{t+1} \right) + \mu_{it}, \tag{7}$$

where $\mu_{it} \in \mathbb{R}$ is the sum of the Lagrange multipliers associated with the exposure constraints (3).

The market clearing condition is

$$\sum_{i=1}^{N} S_{it} = \sum_{s=0}^{\infty} f_s S_{st} = N. \tag{8}$$

If all individuals share the same model consistent expectations, $E_{it}(\cdot) = E_t(\cdot)$ and the transversality condition $\lim_{T\to\infty} (\beta\phi)^T E_t(P_T) = 0$ is satisfied, the REE solution is

$$P_t^{REE} = \frac{\beta\phi e^{\mu+\sigma^2/2}}{1 - \beta\phi e^{\mu+\sigma^2/2}} D_t. \tag{9}$$

Further imposing the parameter restrictions

$$\beta\phi e^{\mu+\sigma^2/2} < 1 \text{ and } \sigma^2/2 \approx 0$$

would imply that the stock price is finite, and that it does not depend on the variance of dividends: $P_t^{REE} = \frac{\beta\phi e^{\mu}}{1-\beta\phi e^{\mu}} D_t$.

2.2 Learning from experience

We depart from model-consistent expectations by assuming that individuals have only limited information about the world they live in. In particular, they do not know anything about other market participants' preferences or constraints. However, they do know their own objectives and constraints and have a prior belief about parameters μ and σ^2 governing the dividend process (4). In the absence of common knowledge, from an individual's perspective, *the price of the asset itself is a stochastic process* affecting optimal savings decisions much like dividends do. Hence individuals try to forecast both the dividend *and* the stock price, conditioning their forecasts on the history of past dividends and stock price realizations.

Individuals update their beliefs about the mean growth rate of the stock price and of dividends, μ. Given P_{t-1} and D_{t-1}, the perceived law of motion (PLM) is

$$x_t = m + \epsilon_t, \tag{10}$$

where

$$x_t = \begin{bmatrix} \log\left(P_t/P_{t-1}\right) \\ \log\left(D_t/D_{t-1}\right) \end{bmatrix}, \quad m = \begin{bmatrix} m^P \\ m^D \end{bmatrix}, \quad \epsilon_t \sim N(0, \Sigma), \quad \Sigma = \begin{bmatrix} \sigma_P^2 & \sigma_{PD}^2 \\ \sigma_{DP}^2 & \sigma_D^2 \end{bmatrix}. \tag{11}$$

This specification allows for beliefs about the growth rates in the share price and dividends to take on different values and their innovations to be imperfectly correlated.

Individuals are assumed to "learn from experience," that is, the information set x_s^t of an agent of age s consists of the realizations of stock prices and dividends observed during his lifetime:

$$x_s^t = \{\log(P_\tau/P_{\tau-1}), \log(D_\tau/D_{\tau-1})\}_{\tau=t-s}^{t}.$$

Individuals are assumed to be born with identical beliefs, centered on the REE outcome in which the asset price grows in lockstep with dividends, $\left(m_0^P, m_0^D\right) = (\mu, \mu)$.

As in Malmenadier and Nagel (2011, 2013), individuals estimate m recursively using past information x_s^t. Let $m_{s,t}$ be the estimator of m using information x_s^t, $m_{s,t} = E[m|x_s^t]$:

$$m_{s,t} = m_{s-1,t-1} + \gamma_{s,t}\left(x_t - m_{s-1,t-1}\right), \quad \text{where } m_{0,t} = [\mu, \mu]'. \tag{12}$$

This is a particular case of the formulation employed in Malmenadier and Nagel (2013). The sequence of gains $\gamma_{s,t}$ determines the rate of updating of cohort s. With $\gamma_{s,t} = 1/t$ it corresponds to the special case of recursive least squares, which uses all available data up to time t with equal weights. Instead, with $\gamma_{s,t}$ set to a constant, it is a constant-gain learning algorithm, which weighs past data with exponentially decaying weights. In the case of "learning from experience" the gain depends on the age s of each cohort. As a result, individuals of distinct ages differ in their forecasts and adjust forecasts differently in response to changes in prices and dividends. The particular decreasing-gain specification which we adopt from Malmenadier and Nagel is

$$\gamma_{s,t} = \gamma_s \equiv \begin{cases} \frac{\theta}{s}, & \text{if } s \geq \theta, \\ 1, & \text{if } s < \theta, \end{cases} , \tag{13}$$

where $\theta \geq 1$ is a constant parameter that determines the shape of the implied function of weights on past prices and dividends.

2.3 Equilibrium

Investor i's subjective expectation $E_{i0}(\cdot)$ is defined on a probability space (Ω, Ψ, Π_i), where Ω is the space of realizations, Ψ the corresponding σ-algebra, and Π_i is a subjective probability measure over (Ω, Ψ). The space of realizations is

$$\Omega \equiv \Omega_P \times \Omega_D, \tag{14}$$

where Ω_P contains all possible sequences of stock prices and Ω_D contains all possible dividend sequences. Individuals can thus condition their investment decision on all possible combinations of dividend and stock price realizations. Denote by Ω^t the set of histories up to period t, and let $x^t \in \Omega^t$. When investor i chooses his stock holding in period t, he takes as given Π_i and his choice is contingent on x^t. Investors have "a consistent set of beliefs", meaning that (Ω, Ψ, Π_i) is a proper probability space and that Π_i satisfies all standard probability axioms and gives proper joint probabilities for all possible dividend and stock price realizations on any set of dates.

The definition of equilibrium in this model is that of an internally rational expectations equilibrium as defined by Adam and Marcet (2011):

Definition 1 (Equilibrium) *An internally rational expectations equilibrium (IREE) consists of a sequence of equilibrium price functions $\{P_t\}_{t=0}^{\infty}$ where $P_t : \Omega_D^t \rightarrow R_+$ for each t, contingent choices $\{C_{it}, S_{it}\}_{t=0}^{\infty}$ where $(C_{it}, S_{it}) : \Omega^t \rightarrow R^2$, and probability beliefs Π_i for each agent i, such that*

1. *All agents $i = 1, ..., N$ choose a function (C_{it}, S_{it}) to maximize their expected utility (1) subject to the budget constraint (2), taking as given the probability measure Π_i.*

2. *Markets clear (8).*

3 Simulation results

In this section, we explore the implications of heterogeneity due to agents being born on different dates and focusing on data realizations from their own lifetimes, rather than on all historical data.

3.1 Calibration

The model's parameters are calibrated to match the U.S. stock market evidence as documented by Shiller (2005). We assume that each period in the model is a quarter.

Dynasts discount future consumption by the factor $\beta\phi$, where β is a time preference parameter and ϕ is the probability of survival. The survival rate is set equal to $\phi = 0.9958$, implying an "average life on the market" of about 60 years. We use Shiller's (2005) stock market dataset covering the S&P index from January 1871 to June 2013 to calibrate our model.[6] In particular, consistent with Shiller's data, we set the mean growth rate of dividends to $\mu = 0.0034$ per quarter, and its standard deviation to $\sigma = 0.0334$. We set the time preference parameter to $\beta = 0.9007$, consistent with a price-to-(quarterly)-dividend ratio in the REE case of 9, as in the data. The leverage ceiling parameter is set to $\lambda = 27$. Note that, by imposing a limit on each individual's investment in the stock, λ affects the measure of households who hold the asset. Setting $\lambda = 27$ is consistent with an average stock market participation rate of around 33 percent, which is the estimate reported by Poterba et. al. (1995) for the average U.S. households in 1992.[7] The value of $\theta = 3.044$, which controls the rate of learning of each agent, comes from the empirical analysis of Malmendier and Nagel (2013). For our numerical simulations, we truncate the maximum number of cohorts to $S = 480$ quarters (120 years), which includes most of the mass of the distribution. The number of Monte Carlo simulations is set to 100.

3.2 Numerical algorithm

The source of heterogeneity in this model is that generations are born one at a time and individuals learn from life experience, which means that different age groups form different beliefs. The key difficulty in solving the model lies in finding the stock price in each period given the current dividend realization and the entire distribution of beliefs about price and dividend growth. Updating dividend beliefs $m^D_{s,t}$ following (12) poses no major problem since dividends D_t are exogenous and known at time t. The problem lies in the computation of P_t, since it determines the beliefs $m^P_{s,t}$, which in turn are used to compute the reservation price of each agent

$$P_{s,t} = \beta\phi\left[e^{m^P_{s,t}(P_t)}P_t + e^{m^D_{s,t}(D_t)}D_t\right], \quad s \in \mathbb{N}^+, \tag{15}$$

and the quantities of stock $S_{s,t}(P_t)$ demanded by agents given FOCs (5a) – (5c). The equilibrium stock price should clear the market as per equation (8).

We find the equilibrium stock price P_t by solving a fixed point problem. Appendix B provides

[6] Shiller's data are monthly. In order to convert to quarterly, we sum up the three dividend payments every quarter and consider the quarterly price as the one in the last month. Results are robust to alternative definitions.

[7] Including mutual fund and 401(k) plan participation.

a sketch of the algorithm. In order to satisfy the FOCs, the price that clears the market must be that of the marginal stock holder. The cohorts in which agents are more optimistic ($P_t < P_{it}$) face condition (5a), $S_{it} = \bar{L}_t/P_t = \lambda D_t/P_t$, whereas the cohorts in which agents are more pessimistic ($P_t > P_{it}$) face condition (5c), $S_{it} = \underline{L}_t/P_t = 0$. Given an initial price guess P_t^{guess}, we have a series of reservation prices $P_{s,t}(P_t^{guess})$ as per (15). We sort these reservation prices in decreasing order and index them by $j \in 0, ..., s-1$, where $j = 0$ corresponds to the highest reservation price, $j = 1$ to the second highest, and so on. Proceeding down the list from the highest reservation price, we find the reservation price P_t^* of the marginal cohort n, $P_t^* = P_{n,t}$. The marginal cohort is the one for which

$$\frac{1}{N}\sum_{j=0}^{n-1} f_j \bar{S}_t < 1, \text{ and } \frac{1}{N}\sum_{j=0}^{n} f_j \bar{S}_t \geq 1, \tag{16}$$

where $\bar{S}_t = \lambda \frac{D_t}{P_t^{guess}}$. This means that cohorts which are more optimistic than the marginal one ($j < n$) choose to hold the stock to the upper limit $\bar{S}_t$, while cohorts which are more pessimistic ($j > n$) choose not to hold the stock at all, satisfying the FOCs. The error in the guess is $P_t^{guess} - P_t^*(P_t^{guess})$. We use a standard numerical algorithm to find the fixed point $P_t = P_t^*(P_t)$ starting from an initial guess $P_t^{guess} = P_{t-1}$.

Figure 1 illustrates the equilibrium computation under the baseline calibration. The dark (blue) line in the upper panel represents the reservation prices of the different cohorts arranged in declining order for some arbitrarily chosen period t, $P_{j,t}$ (the stock prices have been normalized by the price under rational expectations, P_t^{REE}). The most optimistic cohort has a reservation price of 1.013 times P_t^{REE}, whereas the most pessimistic one has $0.965 P_t^{REE}$. The marginal cohort is in position $n = 252$ and its reservation price is $0.970 P_t^{REE}$, equal to the equilibrium stock market price, depicted by the light (blue) line. The lower panel displays the cumulative normalized asset demand $S(n) = \frac{1}{N}\sum_{j=0}^{n} f_j \bar{S}_t$. The equilibrium is achieved for $S(252) = 1$, where demand for the stock equals its supply of unity.

3.3 Heterogeneous beliefs and speculative bubbles

Figure 2 illustrates the behavior of the asset price in our model. The black line in the upper panel shows one particular simulated path of the ratio of the stock price in the baseline economy to the price under rational expectations. Measured on the left scale this ratio oscillates between 0.95 and 1.15, meaning that stock price fluctuations in the heterogeneous agents model are amplified relative to the REE case.

The stochastic oscillations of the stock price around the REE are related to the dynamics of learning. To see this, the black line in the lower panel plots the evolution of the mean expected price across generations relative to the REE price, $\bar{E}_t\left[P_{t+1}\right]/P_t^{REE}$, where $\bar{E}_t(\cdot) \equiv \frac{1}{N}\sum_{s=0}^{\infty} f_s E_{st}(\cdot)$

and $E_{st}(\cdot)$ is the expected value using the information set of cohort s at time t. We also plot the expectations of the youngest ($s = 1$, green line) and the oldest ($s = 480$, thick cyan line) cohorts.[8] Notice that individuals' beliefs regarding the rate of change of the stock price do not converge to the REE value. Instead, they go through successive waves of optimism and pessimism.

The upper panel also shows the price-dividend ratio. In the case of REE this ratio is constant and equal to 9, whereas under learning from experience it displays boom-and-bust cycles around the REE value. It is important to notice how the price-dividend ratio tracks mean price expectations. This is a feature emphasized recently by Adam, Beutel, and Marcet (2013) and that rational expectations models typically fail to reproduce.

Two elements of our model are responsible for the oscillating dynamics. On the one hand, there is a force of momentum, which is rooted in the infrequent resetting of the learning of successive cohorts of individuals as well as in the fact that agents discount older data more heavily than more recent information. Thus, at any given date, a fraction of young individuals enters the market whose learning path initially is more strongly influenced by the more recent stock price and dividend realizations. Their forecasts inform their trading activities, and, through trade, affect the realized stock price, pulling the beliefs of older generations toward the more recent price change realizations. On the other hand, there is a force of trend-reversion, emanating from the constraints on individual risky asset exposure. Namely, as the stock price rises far above the REE, the upper bound in (3) implies that optimistic investors can buy less shares for any given dividend realization. Because, in equilibrium, all shares must be held by someone, the stock price has to fall to the valuation of less optimistic investors. The same reflecting force operates "from below", when the stock price falls too far beneath the REE.[9] The combination of the two factors – momentum and trend reversion – results in boom-and-bust cycles that are only loosely related to dividends.

Indeed, similar to Harrison and Kreps (1978), asset price cycles in our model are partially the result of speculation about the future course of the asset price. Naturally, shocks to dividends do have an influence on the stock price, although the link is not nearly as direct as in the case of REE. Recall that in the REE model, the percentage change in the stock price tracks one-for-one the change in dividends, inheriting the persistence of dividend growth. In contrast, in the OLG model with "learning from experience," a sequence of positive dividend surprises has an escalating effect on asset price changes. This amplification occurs because, through trade, the overreaction to more recent information affects the stock price and, progressively, the beliefs of other individuals, creating a non-linear feedback, which reinforces the effects of dividend shocks on the stock price.

In Table 1 we show the first two moments of the growth rate of stock prices and of dividends, and the price-dividend ratio, in the data and in alternative model simulations. The first column

[8] Recall that the belief of newborns (cohort $s = 0$) is set to the REE solution.

[9] Note that trend reversal kicks in *before* the aggregate leverage constraint $P_t/D_t = \lambda$ becomes binding. Thus, the turning points of the stock price cycles are endogenous in the model.

reports the actual data, the second column shows the REE model, while columns three to six show different parameterizations of our heterogeneous-agents overlapping-generations (HA-OLG) model. One advantage of the HA-OLG model is that it generates dynamics of the price-dividend ratio, unlike the REE model. This is due to the extra volatility of the growth rate of stock prices in the HA-OLG model compared to the REE model.

In the fourth column Table 1 shows the effect of increasing the leverage ceiling parameter λ to 36. The main change is an increase in the volatility of the growth rate of stock prices and of the price-dividend ratio. The reason is that now the marginal agent on average is more optimistic, which explains the increase in the average price-dividend ratio. Next, we look at the effect of reducing the individual gain parameter θ. In the baseline model, it is set to 3.044 as in Malmendier and Nagel (2013). A value of θ above unity implies that agents tend to "forget" the distant past, putting more weight on the more recent observations. Another reason why distant past observations are "lost" is of course the fact that some agents exit the market and are replaced by new generations who start learning anew. To disentangle these two effects, we simulate the model with $\theta = 1$, implying that individuals place the same weight on all observations during their lifetime. The results are displayed in the fifth column of Table 1. In this case the volatility of both the price-dividend ratio and the growth in prices is reduced with respect to the baseline model with $\theta = 3.044$. This is intuitive since in this case the effect of "individual forgetting" is absent and all the dynamics come from the "social forgetting" as young generations replace older ones. Figure 3 is the equivalent of figure 2 for the case of $\theta = 1$. As can be seen in the top panel, the deviation of the stock price from the REE is smaller (between 0.97 and 1.05) which is related to the fact that beliefs are less volatile, as shown in the bottom panel.

Another possibility we look at is that people live longer, reducing the rate of population turnover and allowing individuals to base their decisions on longer time series of data. This experiment is performed in the last column of Table 1 by setting parameter ϕ to 0.9967, consistent with an average life expectancy of 75 years (instead of 60 years as in the baseline). The effect is again a small reduction in the volatility of prices and the price-dividend ratio. The reason is that now the amount of "social forgetting" is reduced.

3.4 Forecasting errors

Learning from experience implies that individuals' forecasts in our model are not fully rational because agents do not take into account all the available information. The important question however is to what extent the suboptimality of agents' forecasts is detectable from the data. Figure 4 plots the mean and the standard deviation of the 1-period ahead forecasting error in the

HA-OLG model as a function of the cohort age,

$$e_{s,t} = \begin{bmatrix} e^P_{s,t} \\ e^D_{s,t} \end{bmatrix} = \begin{bmatrix} \log(P_t/P_{t-1}) - m^P_{s,t-1} \\ \log(D_t/D_{t-1}) - m^D_{s,t-1} \end{bmatrix},$$

alongside the forecasting error of the REE model. The two top panels show that in both models forecasts are unbiased in the sense that mean forecasting errors for both prices and dividends are very close to zero. The two bottom panels show the root mean square error, which, given unbiasedness, equals the standard deviation of the error. The standard deviations of the forecasting errors in the HA-OLG model for both dividends and stock prices are larger than in the REE model. In the REE case the two root mean square errors are the same for prices and dividends, equal to σ. In contrast, in the HA-OLG model, the volatility of prices is considerably higher than that of dividends and hence the volatility of the forecasting errors for prices is higher. This occurs because the stock price depends on market expectations, creating self-referential dynamics as emphasized by Eusepi and Preston (2011). Conversely, in the REE model, uncertainty about prices and dividends is the same because agents coordinate ex-ante onto "the right model" for asset pricing.

Younger cohorts display more significant forecasting error variances in both prices and dividends than older cohorts. This is because older cohorts observe longer histories of data. Of course, the decreasing gain scheme reduces the effect of experience since all agents tend to forget older data. This is confirmed in Figure 5 for the case of $\theta = 1$. In this case the standard deviations of the forecasting errors are significantly smaller than in the baseline case.

4 Approximate aggregate dynamics: constant-gain learning

This section explores the possibility of analyzing the approximate aggregate dynamics of our economy without having to deal with the entire distribution of beliefs across agents. We show how the average belief can be approximated by a representative-agent constant-gain learning (RA-CGL) scheme.

4.1 Theoretical result

Let us define the average belief across age cohorts,

$$\bar{m}_t = \left[\bar{m}^P_t, \bar{m}^D_t\right]' = \frac{1}{N}\left[\sum_{s=0}^{\infty} f_s m^P_{s,t}(P_t), \sum_{s=0}^{\infty} f_s m^D_{s,t}(D_t)\right]'. \tag{17}$$

The evolution of the average belief is given by the following proposition:

Proposition 2 (Average market beliefs) *The average market belief is given by*

$$\bar{m}_t = \bar{m}_{t-1} + \bar{\gamma}\left(x_t - \bar{m}_{t-1}\right) + \xi_t, \tag{18}$$

where

$$\bar{\gamma}(\phi,\theta) \equiv \frac{1}{N}\sum_{s=1}^{\infty} f_s \gamma_s = (1-\phi)\left[\sum_{s=1}^{\lfloor\theta\rfloor} \phi^s \left(1 - \frac{\theta}{s}\right) - \theta \log(1-\phi)\right], \tag{19}$$

is the average gain factor across age cohorts and

$$\xi_t \equiv (1-\phi)\left(m_0 - \bar{m}_{t-1}\right) + (1-\phi)\sum_{s=1}^{\infty} \phi^s \gamma_s (\bar{m}_{t-1} - m_{s-1,t-1}),$$

is a residual term.

Proof. See Appendix A. ■

Therefore, average beliefs about price and dividend growth are updated approximately according to a constant gain learning (CGL) scheme plus a residual term ξ_t. The approximation will be good provided that ξ_t is small compared to the rest of the terms in (18), an issue discussed below. CGL can thus be viewed as an approximate aggregation of the learning of individuals who learn from experience, using data realized in their lifetimes. Notice that the CGL algorithm differs from the actual learning scheme of any of the individual agents because individual learning happens with a decreasing gain, as shown in (13). The population as a whole, however, learns approximately with a constant gain.

The social gain parameter $\bar{\gamma}$ is a non-linear function of the survival probability ϕ and the individual gain θ. Figure 6 plots $\bar{\gamma}(\phi,\theta)$. In the case $\theta = 1$, this is just $\bar{\gamma}(\phi, 1) = -(1-\phi)\log(1-\phi)$. The value of $\bar{\gamma}$ is increasing and concave in θ and $1-\phi$. Under our baseline calibration, the social gain is equal to 0.0588, whereas in the case with $\theta = 1$ it is 0.0228. These numbers are larger, but of the same order of magnitude than existing estimates of the constant-gain parameter both from macro time series data and from surveys. Milani (2007) estimates the gain-parameter in a representative agent model to be 0.0183 in U.S. data, which is very close to the 0.0180 estimated by Malmendier and Nagel (2013).

4.2 Open-loop results

Proposition 2 claims that the average beliefs follow a CGL updating scheme provided that ξ_t is small relative to $\bar{m}_{t-1}$ and $\bar{\gamma}\left(x_t - \bar{m}_{t-1}\right)$. In this section we analyze whether this condition holds for the baseline calibration considered here. To do so, we simulate the baseline model (HA-OLG) and construct the vector x_t using the simulated prices and dividends. Then we run the CGL algorithm (18) and compute $\bar{m}_t$ and ξ_t. We call this 'open-loop' as there is no feedback from the CGL beliefs to the model, that is, these are the results taken the series of prices and dividends as exogenous for the representative-agent CGL (RA-CGL) algorithm.

Figure 7 compares the HA-OLG model with the RA-CGL approximation. Both expectations about price and dividend growth are well approximated by the CGL algorithm. This is confirmed by results in the first column of Table 2 – in particular the correlation between the two series is very high. The mean of ξ_t is close to zero, meaning that the approximation is relatively unbiased.

We explore how sensitive these results are to changes in θ. The second column of Table 2 shows how a reduction in θ, which implies a lower value of $\bar{\gamma}$, slightly improves the approximation by reducing the volatility of ξ_t more than that of $\bar{m}_t$. Given these results, we are confident that in this model the CGL algorithm provides a good approximation to a HA-OLG economy with learning from experience.

4.3 Closed-loop results

We next analyze the 'closed-loop' solution which involves also a feedback *from* beliefs *to* prices. This amounts to analyzing an independent representative-agent economy with CGL. Taking into account the market clearing condition (8), we can compute the asset price as a function of the average belief

$$\begin{aligned} P_t &= \beta\phi \frac{1}{N}\sum_{s=0}^{\infty} f_s E_{st}\left(P_{t+1}+D_{t+1}\right) + \frac{1}{N}\sum_{s=0}^{\infty} f_s \mu_{it} \\ &= \beta\phi \bar{E}_t\left(P_{t+1}+D_{t+1}\right) + \mu_t, \end{aligned} \tag{20}$$

where $\mu_t = \frac{1}{N}\sum_{s=0}^{\infty} f_s \mu_{st}$ is the average Lagrange multiplier. Equation (20) states that the market price is equal to the discounted average belief of next period payoff plus a term reflecting the deviation of the average from the marginal asset holder's belief. Numerical simulations show that in our problem the term μ_t follows a zero-mean iid process and thus we ignore it. This implies that the dynamics of the marginal asset holder's beliefs are similar to those of the average stock holder.

Provided that the growth rate of beliefs is small enough, the asset price can be approximated by the average beliefs across cohorts

$$
\begin{aligned}
P_t &\approx \beta\phi \bar{E}_t\left(P_{t+1}+D_{t+1}\right)=\beta\phi\left\{P_t\sum_{s=0}^{\infty} f_s e^{m^P_{s,t}(P_t)}+D_t\sum_{s=0}^{\infty} f_s e^{m^D_{s,t}(D_t)}\right\} \\
&\approx \beta\phi\left\{P_t \exp\left[\sum_{s=0}^{\infty} f_s m^P_{s,t}(P_t)\right]+D_t \exp\left[\sum_{s=0}^{\infty} f_s m^D_{s,t}(D_t)\right]\right\} \\
&\approx \beta\phi\left[P_t e^{\bar{m}^P_t(P_t)}+D_t e^{\bar{m}^D_t(D_t)}\right],
\end{aligned} \tag{21}
$$

where $\left[\bar{m}^P_t, \bar{m}^D_t\right]$ are a function of the current level of prices and dividends, respectively, and are updated according to the CGL rule (18). The system of equations (4, 18, 21) and the variables $(D_t, P_t, \bar{m}_t)$ define an independent representative-agent economy with CGL beliefs (RA-CGL). In order to solve it, we use numerical methods to find a price vector P_t such that

$$
P_t=\frac{\beta\phi e^{\bar{m}^P_t(D_t)}}{1-\beta\phi e^{\bar{m}^P_t(P_t)}} D_t .
$$

Notice that P_t is an argument of $\bar{m}^P_t(P_t)$.

Figure 7 also displays the difference between the beliefs in the HA-OLG case and the RA-CGL. In the case of dividends the approximation is as good as it was in the open-loop case, dividends being an exogenous variable. In the case of prices, the RA-CGL beliefs match the general pattern of the HA-OLG economy, but they are more volatile. This is confirmed by the numerical results in the third column of Table 2. The volatility of the approximation error ξ^P_t is larger (0.0043) than in the open-loop case (0.0025). The correlation between the RA-CGL and HA-OLG price series is still very high (0.9560).

The upper panel of Figure 8 displays the price-dividend ratio of the RA-CGL (closed-loop) and the HA-OLG models. The feedback from beliefs to prices increases the volatility of the price-dividend ratio in the RA-CGL case compared to the HA-OLG. This is the reason why price beliefs are also more volatile in this case. Notice that in the RA-CGL case there is no leverage or short-selling constraint operating.

Similar to the open-loop case, a reduction in the value of θ improves the fit of the approximation, as shown in the last column of Table 2. It also reduces the volatility of the price-dividend ratio, as shown in the lower panel of Figure 8.

The main conclusion of this exercise is that CGL provides a reasonably accurate approximation, albeit more volatile, to an OLG economy with learning from experience, with a social gain parameter that depends on both the survival probability and the individual gain parameters of the agents.

5 Conclusions

In order to coordinate *a priori* to a REE, individuals must be endowed with incredible amounts of information not only about the structure of the economy and the exogenous shocks but also about the higher-order beliefs of all other market participants. If individuals lack this information, the law of iterated expectations is no longer valid and "beauty contest" dynamics may emerge as individuals embark on speculative trading as in Harrison and Kreps (1978). In particular, empirical research by Malmendier and Nagel (2009, 2011) suggests that expectations are not "externally rational" in the sense of Adam and Marcet (2011); rather, they find evidence that people "learn from experience," giving more weight to recent events realized during their lives than to older ones.

We consider a Lucas-tree stochastic OLG setup and analyze the effects of "learning from experience." The fact that different generations of individuals hold different beliefs leads to boom-and-bust cycles of the stock price around the REE. The aggregate market dynamics can be approximated by a representative-agent model with CGL. Despite the fact that individuals learn with decreasing gain, learning by the population as a whole can be approximated by a constant gain. The social gain parameter is a nonlinear function of the survival rate and the individual gain parameters, reflecting both the fact that historical data is lost when older generations are replaced by young ones and the higher weight that each individual gives to more recent events. This result provides an alternative justification for the use of CGL algorithms in macroeconomic models instead of the more widely used rational expectations. Besides achieving more realism in modeling the expectations formation process, our approach provides discipline by tying the gain parameter to the survival rate and to survey-based information on individual learning gains.

References

[1] Adam, K., Beutel, J., and Marcet, A. 2013. Stock Price Booms and Expected Capital Gains. Unpublished manuscript, University of Mannheim.

[2] Adam, K. and A. Marcet. 2011. "Internal Rationality, Imperfect Market Knowledge and Asset Prices," *Journal of Economic Theory*, In Press.

[3] Adam, K., A. Marcet, and J. P. Nicolini. 2008. "Stock Market Volatility and Learning," European Central Bank Working Paper No. 862.

[4] Allen, F., S. Morris, and H. S. Shin. 2006. "Beauty Contests and Iterated Expectations in Asset Markets," *Review of Financial Studies*, vol. 19 (3), pp. 719–52.

[5] Angeletos, G. M., and J. La'O. 2010. Noisy business cycles. In NBER Macroeconomics Annual 2009, Volume 24 (pp. 319-378). University of Chicago Press.

[6] Angeletos, G. M. and J. La'O. 2011. "Decentralization, Communication, and the Origins of Fluctuations," Working paper, MIT

[7] Arrow, K.J. 1986. "Rationality of Self and Others in an Economic System," *The Journal of Business*, vol. 59 (4), S385–99.

[8] Blume, L. E., M. M. Bray, and D. Easley. 1982. "Introduction to the Stability of Rational Expectations Equilibrium," *Journal of Economic Theory*, vol. 26 (2), pp. 313–17.

[9] Boswijk, H. P., Hommes, C. H., & Manzan, S. 2007. Behavioral heterogeneity in stock prices. Journal of Economic Dynamics and Control, 31(6), 1938-1970.

[10] Branch, W. A., and G. W. Evans. 2006. "Intrinsic Heterogeneity in Expectation Formation," *Journal of Economic Theory*, vol. 127, pp. 264–95.

[11] Branch, W., and G. W. Evans. 2011. "Learning about Risk and Return: A Simple Model of Bubbles and Crashes," *American Economic Journal: Macroeconomics*, vol. 3 (3), pp. 159–91.

[12] Branch, W. and B. McGough. 2004. "Multiple Equilibria in Heterogeneous Expectations Models," *The B. E. Journal of Macroeconomics*, Berkeley Electronic Press, vol. 0 (1), pp. 12.

[13] Brock, William A., and Cars H. Hommes. 1998. "Heterogeneous beliefs and routes to chaos in a simple asset pricing model." Journal of Economic Dynamics and Control 22.8-9: 1235-1274.

[14] Brown, A., and L. Rogers. 2009. "Heterogeneous Beliefs with Finite-Lived Agents," University of Cambridge, Statistical Laboratory, Working Paper.

[15] Carceles-Poveda, E., and C. Giannitsarou. 2007. "Adaptive Learning in Practice," *Journal of Economic Dynamics and Control*, vol. 31 (8), pp. 2659-97.

[16] Cogley, T. and T. J. Sargent. 2008. "The Market Price of Risk and the Equity Premium: A Legacy of the Great Depression?," *Journal of Monetary Economics*, vol. 55 (3), pp. 454–76.

[17] Cogley, T., T. J. Sargent and V. Tsyrennikov. 2012. "Market Prices of Risk with Diverse Beliefs, Learning, and Catastrophes," *American Economic Review*, vol. 102 (3), pp. 141-46.

[18] DeGroot, M. H. 1970. *Optimal Statistical Decisions*. McGraw-Hill Press.

[19] Eusepi, S., and B. Preston. 2011. "Expectations, Learning and Business Cycle Fluctuations," *American Economic Review*, forthcoming.

[20] Evans, G. W. and S. Honkapohja. 2003. *Learning and Expectations in Macroeconomics*, Princeton University Press.

[21] Ferrero, G. 2007. "Monetary Policy, Learning and the Speed of Convergence," *Journal of Economic Dynamics and Control*, vol. 31 (9), pp. 3006–41.

[22] Giannitsarou, C. 2003. "Heterogeneous Learning," *Review of Economic Dynamics*, vol. 6 (4), pp. 885–906.

[23] Graham, L. 2011. "Learning, Information and Heterogeneity," Working paper.

[24] Harrison, J. M. and D. M. Kreps. 1978. "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations," *Quarterly Journal of Economics*, vol. 92 (2), pp. 323–36.

[25] Hommes, C. H. 2006. Heterogeneous agent models in economics and finance. Handbook of computational economics, 2, 1109-1186.

[26] Honkapohja, S., and K. Mitra. 2006. "Learning Stability in Economies with Heterogeneous Agents," *Review of Economic Dynamics*, vol. 9 (2), pp. 284–309.

[27] Malmendier, U. and S. Nagel. 2011. "Depression Babies: Do Macroeconomic Experiences Affect Risk-Taking?," *Quarterly Journal of Economics*, vol. 126 (1), pp. 373–416.

[28] Malmendier, U. and S. Nagel. 2009. "Learning from Inflation Experiences," Unpublished manuscript, U.C. Berkley

[29] Marcet, A. and T. J. Sargent. 1989. "Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models," *Journal of Economic Theory*, vol. 48 (2), pp. 337–68.

[30] Milani, F., 2007. "Expectations, Learning and Macroeconomic Persistence," J*ournal of Monetary Economics*, vol. 54, pp. 2065–82.

[31] Poterba, J. M., A. A. Samwick, A. Shleifer, and R. J. Shiller. 1995. "Stock Ownership Patterns, Stock Market Fluctuations, and Consumption," *Brookings Papers on Economic Activity*, vol. 1995 (2), pp. 295–372.

[32] Shiller, R. 2005. *Irrational Exuberance*, 2nd edition, Princeton University Press.

[33] Timmerman, A. 1994. "Can Agents Learn to Form Rational Expectations? Some Results on Convergence and Stability of Learning in the UK Stock Market," *The Economic Journal*, vol. 104 (425), pp. 777–97.

[34] Weitzman, M. L. 2007. "Subjective Expectations and Asset-Return Puzzles," *American Economic Review*, vol. 97 (4), 1102–30.

Appendix A: Proof of Proposition 2

Proof. Let

$$\bar{m}_t = \frac{1}{N}\sum_{s=0}^{\infty} f_s m_{s,t} = \frac{f_0 m_{0,t}}{N} + \frac{1}{N}\sum_{s=1}^{\infty} f_s \left[m_{s-1,t-1} + \gamma_s \left(x_t - m_{s-1,t-1}\right)\right],$$

given the fact that

$$\frac{1}{N}\sum_{s=1}^{\infty} f_s m_{s-1,t-1} = \phi \bar{m}_{t-1},$$

and that

$$\begin{aligned}
\bar{\gamma} &= \frac{1}{N}\sum_{s=1}^{\infty} f_s\gamma_s = \frac{1}{N}\sum_{s=1}^{\lfloor\theta\rfloor} f_s + \frac{1}{N}\sum_{s=\lfloor\theta\rfloor+1}^{\infty} f_s\frac{\theta}{s} = (1-\phi)\left[\sum_{s=1}^{\lfloor\theta\rfloor}\phi^s + \theta\sum_{s=\lfloor\theta\rfloor+1}^{\infty}\frac{\phi^s}{s}\right] \\
&= (1-\phi)\left[\sum_{s=1}^{\lfloor\theta\rfloor}\phi^s + \theta\sum_{s=\lfloor\theta\rfloor+1}^{\infty}\int\phi^{s-1}d\phi\right] = (1-\phi)\left[\sum_{s=1}^{\lfloor\theta\rfloor}\phi^s + \theta\int\left(\sum_{s=\lfloor\theta\rfloor+1}^{\infty}\phi^{s-1}\right)d\phi\right] \\
&= (1-\phi)\left[\sum_{s=1}^{\lfloor\theta\rfloor}\phi^s + \theta\int\frac{\phi^{\lfloor\theta\rfloor}}{1-\phi}d\phi\right] = (1-\phi)\left[\sum_{s=1}^{\lfloor\theta\rfloor}\phi^s - \theta\sum_{s=1}^{\lfloor\theta\rfloor}\frac{\phi^s}{s} - \theta\log(1-\phi)\right] \\
&= (1-\phi)\left[\sum_{s=1}^{\lfloor\theta\rfloor}\phi^s\left(1-\frac{\theta}{s}\right) - \theta\log(1-\phi)\right],
\end{aligned}$$

we have

$$\bar{m}_t = \frac{1}{N}\sum_{s=0}^{\infty} f_s m_{s,t} = (1-\phi)m_0 + \phi\bar{m}_{t-1} + \bar{\gamma}x_t - \frac{1}{N}\sum_{s=1}^{\infty} f_s\gamma_s m_{s-1,t-1},$$

and defining

$$\xi_t \equiv (1-\phi)\left(m_0 - \bar{m}_{t-1}\right) + (1-\phi)\sum_{s=1}^{\infty}\phi^s\gamma_s(\bar{m}_{t-1} - m_{s-1,t-1}),$$

we obtain (19). ■

Appendix B: Simulation algorithm

We briefly sketch the algorithm used to find the equilibrium price of our heterogeneous-beliefs economy. The idea is to simulate the evolution of dividends while keeping track of each agent's stock holdings and beliefs. To compute the equilibrium price given equations (7) and (8), we employ a numerical routine to find a fixed-point for the price which is consistent with agents' beliefs and constraints and which guarantees that the market clears.

Here we describe a single Monte Carlo simulation of the model:

1. Generate an exogenous series for dividends D_t following (4) and assuming that $D_0 = 1$. Set $P_0 = P_0^{REE}$, where P_t^{REE} is given by (9).

2. Initialize the prior beliefs, $\left(m_{s,0}^P, m_{s,0}^D\right) = (\mu, \mu)$, for all cohorts, $s = 0, ..., S$.

3. **Main loop**. At each point in time $t = 1, ..., T$:

   (a) Compute the dividend beliefs across cohorts

   $$m_{s,t}^D = m_{s-1,t-1}^D + \gamma_s \left[\log(D_t/D_{t-1}) - m_{s-1,t-1}^D\right], s = 1, ..., S.$$

   (b) Compute the price P_t as a *fixed-point* given equations (7) and (8). The initial guess $P_t^{guess} = P_{t-1}$. Employ the following subroutine:

      i. Given the guess, compute the price beliefs across cohorts

      $$m_{s,t}^P = m_{s-1,t-1}^P + \gamma_s \left[\log(P_t^{guess}/P_{t-1}) - m_{s-1,t-1}^P\right], s = 1, ..., S.$$

      ii. Compute the reservation price across cohorts

      $$P_{s,t} = \beta\phi \left[e^{m_{s,t}^P} P_t^{guess} + e^{m_{s,t}^D} D_t\right].$$

      iii. Sort the reservation prices $P_{s,t}$ in decreasing order. Index the prices by j, where $j = 0$ corresponds to the highest $P_{s,t}$.

      iv. Proceeding from the highest reservation price, find the reservation price of the marginal cohort $P_t^* = P_{n,t}$. The marginal cohort n is such that

      $$\sum_{j=0}^{n-1} f_j \bar{S}_t < N, \text{ and } \sum_{j=0}^{n} f_j \bar{S}_t \geq N$$

      where $\bar{S}_t = \lambda \frac{D_t}{P_t^{guess}}$ is the amount allocated to each agent given the leverage constraint (3) and f_j is the size of the cohort corresponding to index j.

v. The error in the guess is $\zeta_t = P_t^{guess} - P_t^*$.

vi. Repeat (i) to (v) until $|\zeta_t|$ is less than the error tolerance level.

(c) Given the price, compute the price beliefs across cohorts

$$m_{s,t}^P = m_{s-1,t-1}^P + \gamma_s \left[\log(P_t/P_{t-1}) - m_{s-1,t-1}^P\right], s = 1, ..., S.$$

4. Repeat the main loop (3) for periods $t = 1, ..., T$.

6 Appendix C: Tables and Figures

Table 1. Moments of prices and dividends

	Data	REE	HA-OLG			
			Baseline	$\lambda = 36$	$\theta = 1$	$\phi = 0.9967$
$\log(P_t/P_{t-1})$						
mean	0.0052	0.0034	0.0034	0.0034	0.0034	0.0034
st. dev.	0.0886	0.0334	0.0463	0.0525	0.0371	0.0447
$\log(D_t/D_{t-1})$						
mean	0.0034	0.0034	0.0034	0.0034	0.0034	0.0034
st. dev.	0.0334	0.0334	0.0334	0.0334	0.0334	0.0334
P_t/D_t						
mean	9.00	9.00	9.14	9.25	9.10	9.16
st. dev.	4.66	-	0.38	0.51	0.20	0.37

Note: REE stands for "rational expectations equilibrium."

HA-OLG stands for "heterogeneous-agents overlapping generations."

Table 2. Evaluation of the CGL approximation

	Open-loop		Closed-loop	
	$\theta = 3.044$	$\theta = 1$	$\theta = 3.044$	$\theta = 1$
Prices				
mean $\bar{m}_t^{P,CGL}$	0.0034	0.0034	0.0034	0.0034
st. dev. $\bar{m}_t^{P,CGL}$	0.0072	0.0037	0.0080	0.0040
mean ξ_t^P	0.0000	0.0000	0.0000	0.0000
st. dev. ξ_t^P	0.0025	0.0013	0.0043	0.0016
$\text{corr}\left(\bar{m}_t^P, \bar{m}_t^{P,CGL}\right)$	0.9628	0.9789	0.9560	0.9737
Dividends				
mean $\bar{m}_t^{D,CGL}$	0.0034	0.0034	0.0034	0.0034
st. dev. $\bar{m}_t^{D,CGL}$	0.0058	0.0034	0.0058	0.0034
mean ξ_t^D	0.0000	0.0000	0.0000	0.0000
st. dev. ξ_t^D	0.0020	0.0012	0.0020	0.0012
$\text{corr}\left(\bar{m}_t^D, \bar{m}_t^{D,CGL}\right)$	0.9648	0.9787	0.9648	0.9787

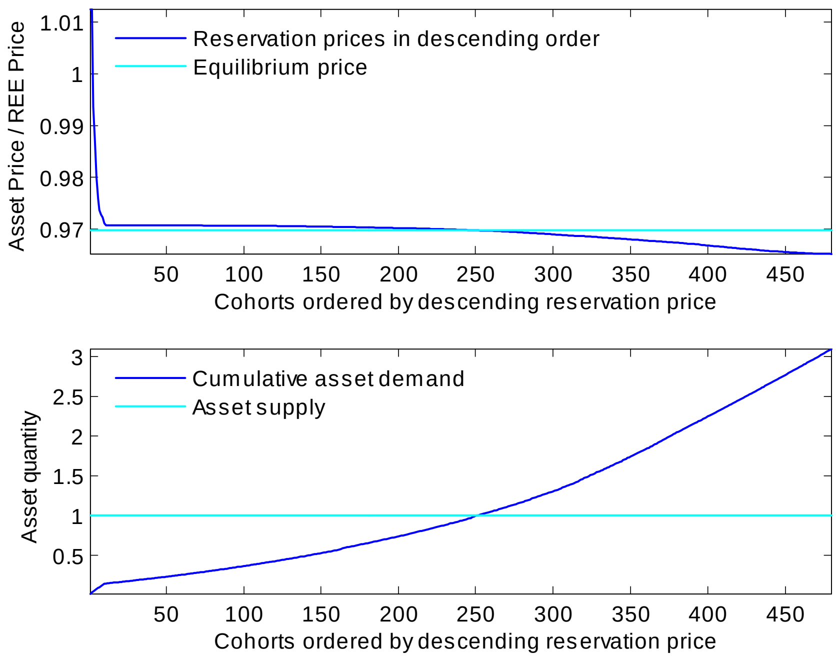

Figure 1: Distribution of reservation prices and asset demand across cohorts

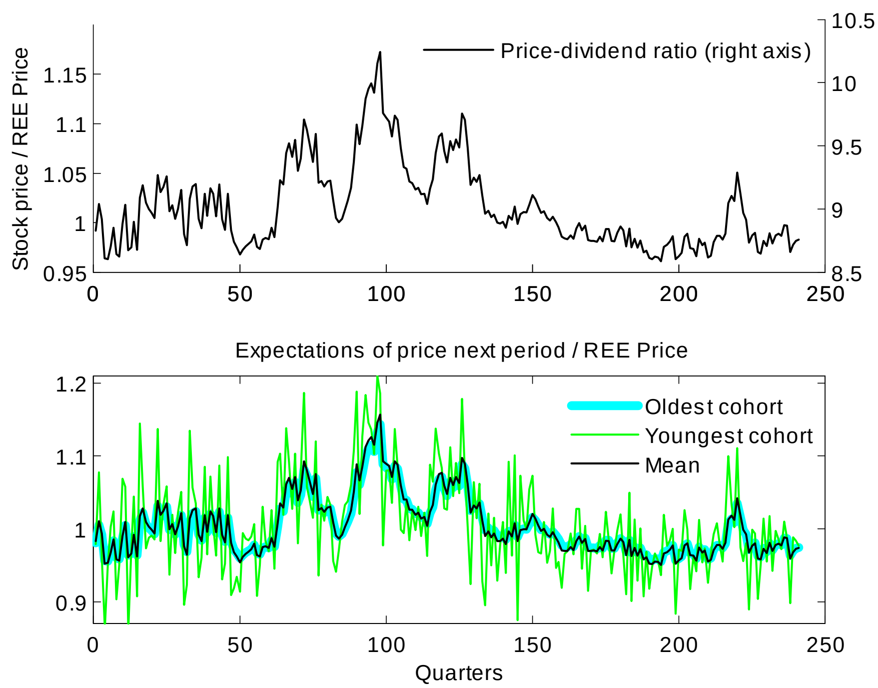

Figure 2: Evolution of prices, the price-dividend ratio, and price expectations ($\theta = 3.044$)

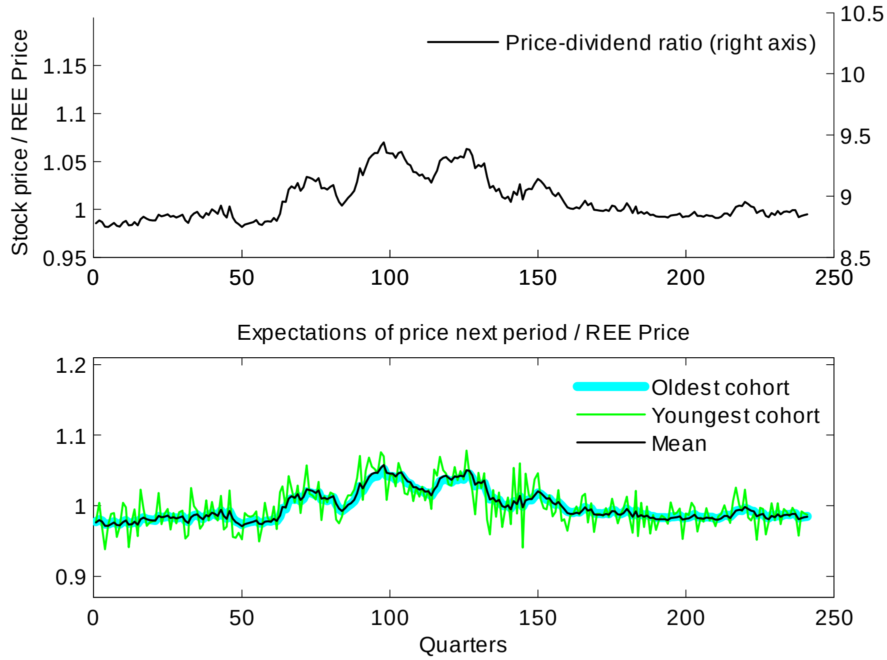

Figure 3: Evolution of prices, the price-dividend ratio, and price expectations ($\theta = 1$)

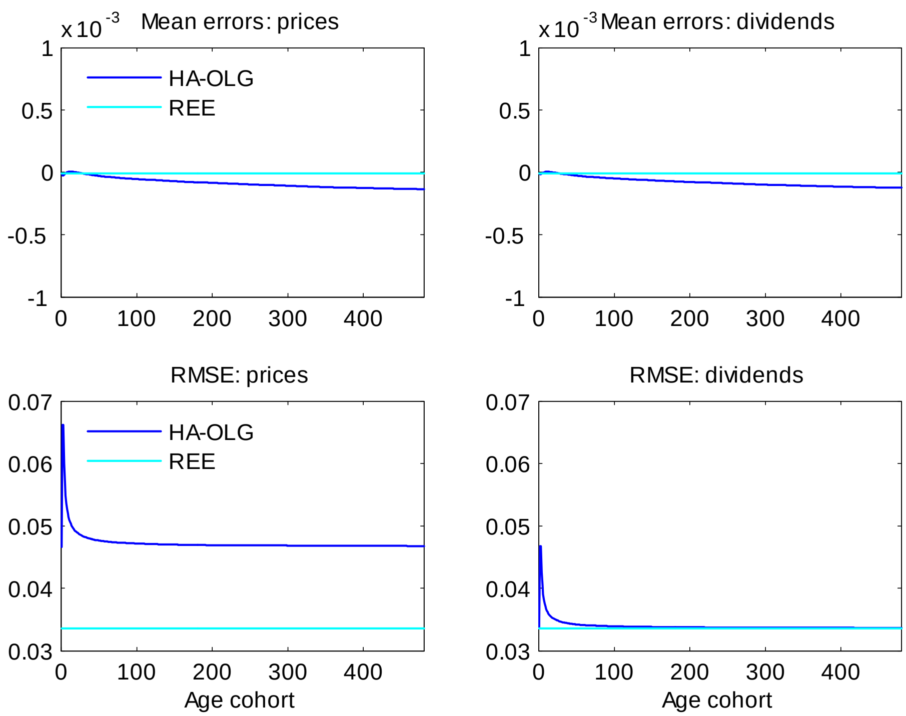

Figure 4: One-period-ahead mean and root mean squared forecasting errors for prices and dividends in the learning from experience model (HA-OLG, $\theta = 3.044$) and the REE

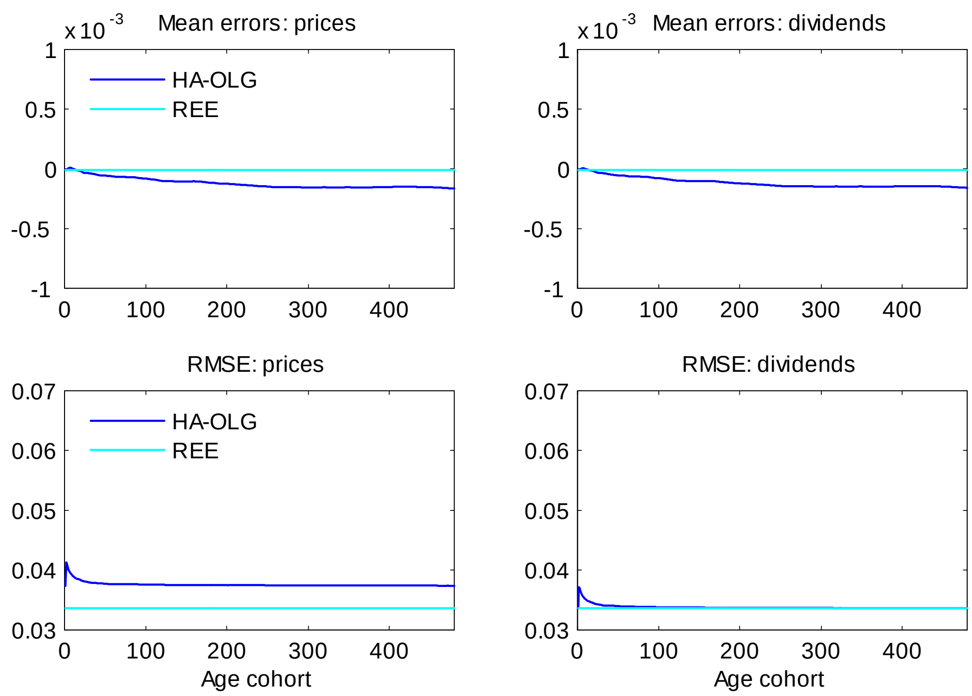

Figure 5: One-period-ahead mean and root mean squared forecasting errors for prices and dividends in the learning from experience model (HA-OLG, $\theta = 1$) and the REE

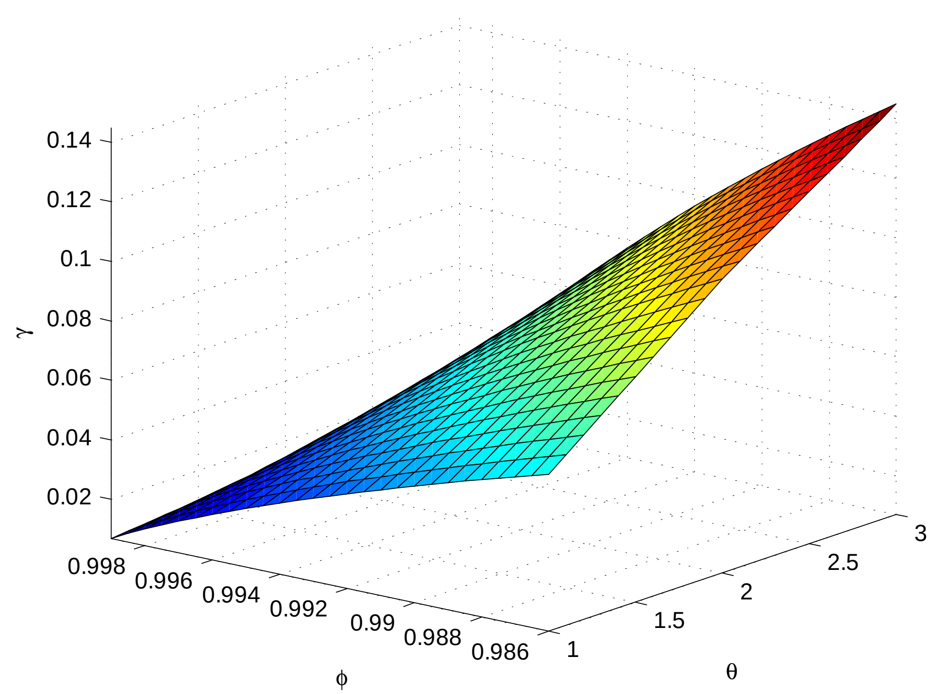

Figure 6: Social gain γ as a function of θ and ϕ

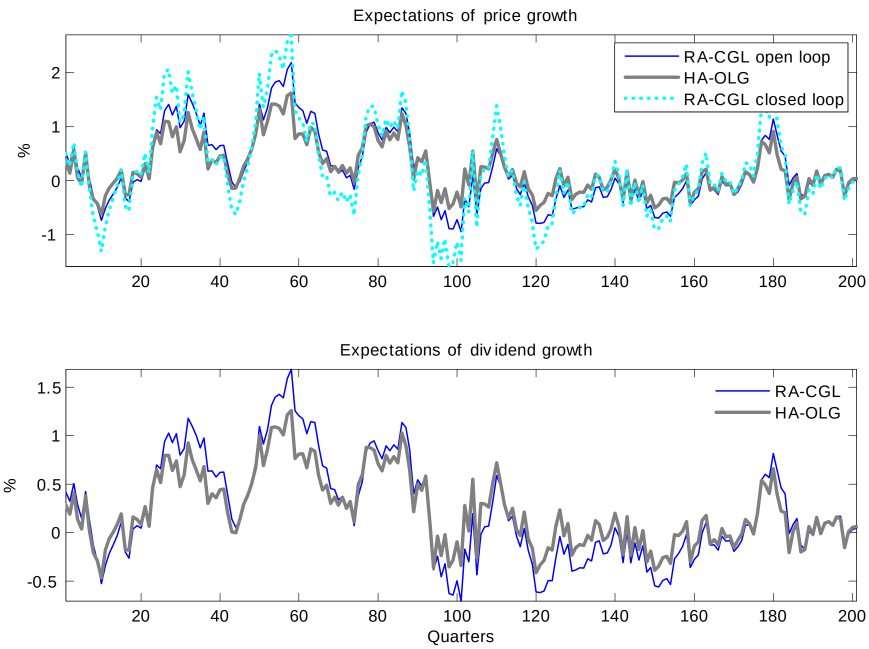

Figure 7: Price and dividend growth beliefs in the heterogeneous agents model (HA-OLG) and in the representative-agent CGL approximation (RA-CGL). Open and closed loop results ($\theta = 3.044$).

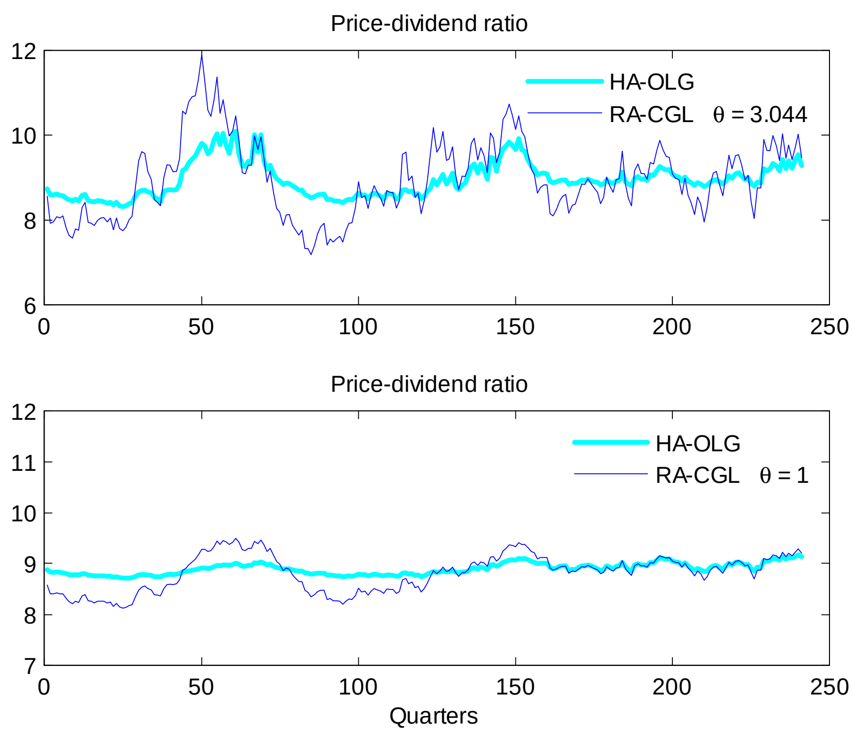

Figure 8: Dynamics of the price-dividend ratio under HA-OLG and RA-CGL (closed loop)