

**Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.**

Learning from Experience in the Stock Market

Anton Nakov

2012-41

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.

Learning from Experience in the Stock Market*

Anton Nakov

Galo Nuño

Federal Reserve Board

Banco de España

May 22, 2012

Abstract

We study the dynamics of a Lucas-tree model with finitely lived individuals who “learn from experience.” Individuals update expectations by Bayesian learning based on observations from their own lifetimes. In this model, the stock price exhibits stochastic fluctuations around the rational expectations equilibrium. This heterogeneous-agents economy can be approximated by a representative-agent model with constant-gain learning, where the gain parameter is related to the survival rate.

Keywords: learning from experience, OLG, asset pricing, bubbles, heterogeneous agents

JEL codes: G12, D83, D84

*We are grateful for comments and suggestions to Gunter Coenen, John Roberts, Alistair Dieppe, Oliver de Groot, Alessio Moro, Giulio Nicoletti, Miguel Angel Iraola, and seminar participants at ECB, University of Surrey, and University of Cagliari, and conference participants at SCE 2011. Anton Nakov thanks the European Central Bank for its hospitality during the first drafts of this paper. The views expressed here are those of the authors and do not necessarily coincide with the views of the Eurosystem or the Federal Reserve System. Corresponding author: Galo Nuño, International Affairs, Bank of Spain, Alcalá 48, 28014 Madrid, Spain. E-mail: galo.nuno@bde.es

1 Introduction

The crucial role of expectations about the future is well understood in economics. The rational expectations hypothesis (REH) has been an important step forward allowing rigorous formalization of the process of expectations formation. Yet it has been often criticized for endowing people with “too much” knowledge about their environment.¹ Empirical research studying how individuals form expectations about aggregate economic variables does not, in general, corroborate the REH. In particular, Malmendier and Nagel (2009, 2011) find evidence that, contrary to the REH, people “learn from experience,” meaning that individuals are more strongly influenced by data realized during their own lifetimes than by earlier historical data. More specifically, Malmendier and Nagel (2011) find that individuals who experienced low stock market returns during their lives are less likely to participate in the stock market, invest a lower fraction of their liquid assets in stocks, and are more pessimistic about future stock returns. In addition, Malmendier and Nagel (2009) find that young individuals place more weight on recently experienced inflation than older individuals do. The upshot is that learning dynamics may be perpetual if history “gets lost” as new generations replace older ones.

In this paper, we explore how replacing the REH with “learning from experience” modifies the results 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 that are the result of their collective decisions given their beliefs, a type of self-referentiality emphasized by Eusepi and Preston (2011). To this end, we extend the asset pricing model of Adam and Marcet (2011) to a stochastic 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.

Specifically, we assume that a small random fraction of individuals exit the stock market every period with a given probability $1 - \theta$, 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 in a Bayesian way with information about stock prices and dividends which they observe during their own lifetimes. This is our main difference with Adam and Marcet’s (2001) model which has infinitely-lived individuals. Our main difference with Brown and Rogers (2009) is that the mapping between dividend growth and the stock price is not known to the individuals in our model.

We find that, even if the retirement rate $1 - \theta$ 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

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

expectations equilibrium (REE). Instead, a chaotic equilibrium emerges in which the stock price exhibits stochastic cycles (around the REE price), the frequency of which is positively related to the rate of retirement. Two forces create the oscillating dynamics. On the one hand, there is “momentum” rooted in the continuous entry of new individuals. At any given date, a fraction of young individuals discount the experience of their parents and pay more attention to the most recent stock price developments. The latter biases the young’s beliefs about the future course of dividends and stock prices toward simple extrapolation of the recent past, and their trading activities push the asset price away from the fundamental. On the other hand, there is a force of reversal toward the REE trend. When the stock price rises too far above the fundamental value, individual leverage 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 reversion – 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 key finding 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 takes two steps. In a first 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. In a second step, we show that the dynamics of average beliefs can be approximated by a CGL scheme in which the gain parameter is a function of the survival rate θ . This approximation 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.”

CGL is usually motivated based on its ability to produce realistic model features, such as amplification of the persistence of macro variables in response to aggregate shocks.² Rarely is there a discussion of the reasons why all agents should learn in the same suboptimal way. The value of the gain parameter typically is estimated or calibrated to yield the smallest possible mean-squared forecasting error. Our contribution is to provide an alternative justification for using CGL in a representative-agent context. Namely, we see it as a useful shortcut to approximating the aggregate dynamics of an economy populated by many Bayesian learners, each of them using a decreasing gain sequence, under the assumption that they “learn from experience.”

Finally, we analyze the asymptotic behavior of our economy in the limit with infinitely lived agents (taking the limit as $\theta \rightarrow 1$). We show analytically that, in this case, even if traders do

²For example, see Milani (2007), Carceles-Poveda and Giannitsarou (2008), Branch and Evans (2011), Adam, Marcet, and Nicolini (2008).

not know anything about each other, endowing them with long histories of dividend and stock price realizations is sufficient for their beliefs to eventually converge to the REE. We study the properties of the convergence, such as the speed and the shape of the transition path. We find that, if new dividend information arrives monthly, it can take several centuries before the asset price comes close to the REE. In the baseline calibration, after a full one century of trading and learning, the median simulated stock price is still 20 percent higher than its REE counterpart.

Our setup rules out the possibility of a rational bubble, defined as a gap between the stock price and the REE price that grows unboundedly in expectations. We preclude bubbles by assuming that individuals face constraints on their maximum exposure to the stock. Specifically, we cap individual leverage, defined as the multiple of the current dividend that an individual is allowed to maintain invested in the form of stock holdings. In our environment leverage is an important factor affecting the properties of convergence to the REE. In particular, the higher the degree of permissible leverage, the slower is the rate of convergence.

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

Second, a related body of literature analyzes the dynamics of asset prices under Bayesian learning by a representative agent. Timmermann (1994), Weitzman (2007), and Cogley and Sargent (2008), among others, offer an explanation for some interesting 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.

Third, following Radner (1979) and Lucas (1972), a large body of literature studies rational expectations equilibria in economies with asymmetric information. Vives (1993), in particular, analyzes the speed of convergence to REE in a model of rational learning in which the market price is informative about an unknown parameter only through the actions of agents. Vives finds that whenever the average precision of private information is finite, convergence to the REE is slow, at the rate $1/\sqrt{n^{1/3}}$, where n is the number of trading periods.

Fourth, recent literature 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

will 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. In section 2, we recast the model of Adam and Marcet (2011) in an OLG setting. In section 3, we calibrate the model and analyze the properties of “learning from experience.” In section 4, we show how the model can be approximated by a representative agent with CGL. Section 5 explores the case in which the survival probability approaches one, and section 6 concludes.

2 The model

In this section, we recast the model of Adam and Marcet (2011) in an OLG setup. We make some additional changes to their model as follows. First, we assume Bayesian learning of the means *and* the variances of the stock price and dividends.³ Second, we specify a particular market arrangement (a centralized auction), which ensures that information about the current dividend is incorporated into the contemporaneous stock price. We are also explicit about the way the market arrives at the equilibrium asset price. Third, we replace Adam and Marcet’s investment constraints on the *number* of shares an agent can hold with constraints on individual exposure in the stock. More precisely, we assume that there is a ceiling for the maximum *value* an individual can invest in the stock, preventing him from going arbitrarily long in the asset. Likewise, we assume that there is a floor for an individual’s position in the stock, preventing him from engaging in unlimited shorting. These value limits, which can be rationalized by underlying credit constraints, are sufficient to rule out rational bubbles without reliance on a “projection facility.”⁴

Adam and Marcet’s model is interesting to us for three reasons. First, it introduces a meaningful distinction between “internal rationality” and “external rationality.” Internally rational individuals maximize expected utility given consistent beliefs about the future. Externally rational individuals are endowed, in addition, with common knowledge of each other’s preferences and beliefs, for any possible path of dividends. We assume that our economy is populated by individuals who are internally rational but are not externally rational. Second, an appealing feature of the model is its simplicity, allowing us to obtain closed-form analytical expressions for the asset price dynamics. Third, despite its simplicity, the model is rich enough to be contrasted with actual data on stock prices and dividends.

The economy is populated by N risk-neutral ex ante identical dynasties. Members of each

³Adam and Marcet (2011) show that, up to a first-order approximation, Bayesian learning of the means, or decreasing-gain recursive least squares learning, are equivalent to full Bayesian learning in a model with an infinitely lived representative agent. Instead, we simply work with Bayesian learning as in De Groot (1970).

⁴A “projection facility” is a technical assumption that mechanically constrains beliefs to a pre-specified neighborhood.

dynasty have stochastic lifetimes with death (or retirement) occurring with a constant exogenous probability, $(1 - \theta)$. Thus, in each period, the measure of dynasts of age $j \in \mathbb{N}_0$ is constant and equal to $f_j = N(1 - \theta)\theta^j$. 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. However, as emphasized by Adam and Marcet, 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 stock price in our model thus equals the *marginal* asset holder i 's *subjective* expected present value of holding the stock for one period, collecting the dividend D_{t+1} , and selling it in the following period at *his* expected price $E_{it}(P_{t+1})$. Because expectations about future prices generally would 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.

In the following subsections we provide a sketch of the model. We provide more details in Appendix A.

2.1 Preferences and constraints

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

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

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

Individual i faces the period budget constraint

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

where S_{it} denotes his stock holdings, P_t is the asset price, D_t is the dividend, and Y_{it} is a 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.

$$\underline{E}_t \leq P_t S_{it} \leq \bar{E}_t. \quad (3)$$

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 $\underline{E}_t = 0$ and $\bar{E}_t = \lambda D_t > 0$, 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 exposure constraints (3) differ from the stock holding constraints used by Adam and Marcet (2011), namely $0 \leq S_{it} \leq \bar{S}$, which limit the minimum and maximum *number of shares* held by an individual. Their constraints suffice for the maximization problem to be well-defined at the individual level. However, they are not sufficient to prevent agents from collectively holding the entire stock at ever-rising prices.⁵ In contrast, 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

$$\log(D_t/D_{t-1}) = \varepsilon_t \sim N(\gamma, \sigma^2), \quad (4)$$

where $\gamma > 0$ 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), and the exposure constraints (3).

The first-order conditions for an individual are

$$P_t = \beta \theta E_{it} (P_{t+1} + D_{t+1}) + \mu_{it}, \quad (5)$$

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

⁵To see this fact, note that the budget constraint (2) alone does not preclude a rational bubble, because with $C_{it} = 0$ we have that $S_{it} = (1 + D_t/P_t)S_{it-1} + Y/P_t \geq S_{it-1}$. That is, agents are not sufficiently discouraged from holding the stock as the stock price rises.

2.2 Learning from experience

Individuals are assumed to “learn from experience,” that is, the information set $\omega_{i,n}^t$ of agent i of age n consists of the realizations of stock prices and dividends observed during his lifetime,

$$\omega_{i,n}^t = \{P_\tau, D_\tau\}_{\tau=t-n}^t.$$

Dynasts update their beliefs about the mean growth rate of the stock price and dividends, γ_i , as well as the covariance matrix of their innovations, Σ_i . Given P_{t-1} and D_{t-1} , individual i 's perceived law of motion is

$$\begin{bmatrix} \log(P_t/P_{t-1}) \\ \log(D_t/D_{t-1}) \end{bmatrix} = \begin{bmatrix} \varepsilon_{it}^P \\ \varepsilon_{it}^D \end{bmatrix} \sim N(\gamma_i, \Sigma_i), \quad \gamma_i = \begin{bmatrix} \gamma_i^P \\ \gamma_i^D \end{bmatrix}, \quad \Sigma_i = \begin{bmatrix} \sigma_{iP}^2 & \sigma_{iPD}^2 \\ \sigma_{iDP}^2 & \sigma_{iD}^2 \end{bmatrix}. \quad (6)$$

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' prior beliefs about these parameters are of the Normal-Wishart conjugate form,

$$\Sigma_i^{-1} \sim W(\mathbf{\Sigma}_0, n_{i0}) \quad \text{and} \quad \gamma_i | \Sigma_i^{-1} \sim N\left((\gamma_{i0}^P, \gamma_{i0}^D)', \Sigma_i/n_{i0}\right), \quad (7)$$

where the Wishart distribution W with precision matrix $\mathbf{\Sigma}_0$ and $n_{i0} > 3$ degrees of freedom specifies individuals' prior marginal distribution of the inverse of the covariance matrix of innovations. In turn, the normal distribution N specifies individuals' prior belief about the mean growth rates of the stock price and of dividends, conditional on the precision matrix Σ_i^{-1} . The vector $(\gamma_{i0}^P, \gamma_{i0}^D)$ denotes the conditional prior mean, while n_{i0} is the precision of prior beliefs.

Individuals are assumed to be born with identical prior beliefs, centered on the REE outcome in which the asset price grows in lockstep with dividends,

$$(\gamma_{i0}^P, \gamma_{i0}^D) = (\gamma, \gamma), \quad \mathbf{\Sigma}_0 = \sigma^2 \begin{bmatrix} 1 & \delta \\ \delta & 1 \end{bmatrix} (n_0 - 3), \quad \text{where } \delta \rightarrow 1, \delta < 1. \quad (8)$$

The joint distribution of the stock price and dividends is computed as the posterior of (γ_i, Σ_i) conditional on information $\omega_{i,n}^t$ available up to period t . The posterior distribution is also a Normal-Wishart with location parameters $(\gamma_{it}^P, \gamma_{it}^D, \mathbf{\Sigma}_{it}, n_{it})$. Defining the one-step-ahead forecast error as

$$\mathbf{e}_{it} = \begin{bmatrix} \log(P_t/P_{t-1}) - \gamma_{it}^P \\ \log(D_t/D_{t-1}) - \gamma_{it}^D \end{bmatrix}, \quad (9)$$

it follows from DeGroot (1970, ch. 9) that the recursive Bayesian updating scheme is given by

$$\gamma_{it+1} = \gamma_{it} + \frac{\mathbf{e}_{it}}{n_{it} + 1}, \quad \Sigma_{it+1} = \Sigma_{it} + \frac{n_{it}}{n_{it} + 1} \mathbf{e}_{it} \mathbf{e}'_{it}, \quad n_{it+1} = n_{it} + 1. \quad (10)$$

2.3 Timeline of events and market arrangement

Events unfold as follows. At the beginning of period t individuals update their beliefs about (γ_i, Σ_i) based on the stock price and dividends observed in period $t - 1$ using the recursive Bayesian updating scheme (9)-(10). Each individual's expectations about the future stock price and dividends are obtained by projecting his latest estimate of the growth rates of the stock price and dividends into period $t + 1$. Given these expectations, individual i computes his reservation price as⁶

$$\begin{aligned} P_{it} &= \beta\theta E_{it}(P_{t+1} + D_{t+1}) = \beta\theta E_{it} [\exp(\varepsilon_{it}^P + \varepsilon_{it+1}^P) P_{t-1} + \exp(\varepsilon_{it}^D + \varepsilon_{it+1}^D) D_{t-1}] \\ &= \beta\theta \left\{ \exp[2\gamma_{it}^P + \Sigma_{it}(1, 1)/(n_{it} - 3)] P_{t-1} + \exp[2\gamma_{it}^D + \Sigma_{it}(2, 2)/(n_{it} - 3)] D_{t-1} \right\}. \end{aligned} \quad (11)$$

The stock is traded on a multiple-round, sealed-bid, centralized auction where actual exchange occurs only in the very last round. The market-clearing price is established as follows. In the first round, each individual sends his initial sealed bid P_{it} given by (11). An auctioneer sorts all the bids from highest to lowest in an order book, and *notionally* allocates the asset, starting from the top bidder and moving down the order book until the entire stock is allocated.⁷ The auctioneer then announces publicly the time t , round 1, *preliminary* asset price, P_{t1} , as the bid of the marginal investor who would just be willing to hold the asset if that were the final price. Thus, price P_{t1} *would* clear the market if trade were allowed at that point and no new information had become available.

We assume, however, that at the end of the first round the actual dividend for time t becomes publicly known. Hence, in subsequent rounds of the auction, investors revise their bids based on the preliminary price announced by the auctioneer in the preceding round P_{tk-1} , and on the time t dividend,

$$P_{itk} = \beta\theta \left\{ \exp[\gamma_{it}^P + \Sigma_{it}(1, 1)/(2n_{it} - 6)] P_{tk-1} + \exp[\gamma_{it}^D + \Sigma_{it}(2, 2)/(2n_{it} - 6)] D_t \right\}. \quad (12)$$

Bids are collected again, and the asset is notionally allocated to the highest bidders, determining the new *preliminary* price of round k as the price offered by the marginal potential buyer. This process is repeated for a large number of rounds until convergence of the price P_{tk} .⁸ This limiting-

⁶To obtain the last equality, we use that $E \exp(\varepsilon) = \exp[E(\varepsilon) + \text{Var}(\varepsilon)/2]$ when ε is normally distributed.

⁷Recall from (3) that no individual can go infinitely long in the asset.

⁸While the convergence is asymptotic, in practice we will cut off the number of rounds to $K = 2000$.

round price is the actual clearing price in period t at which trade occurs.⁹ At the end of period t , owners of the asset receive the dividend and the clearing price, and the successful bidders receive the stock. In equilibrium, individuals collectively hold the entire stock of the asset, so

$$N = \sum_{j=0}^{\infty} f_j S_{jt}, \quad (13)$$

where recall that f_j is the measure of individuals of age j . The model is completely characterized by the first-order conditions for individual investors (5), the recursive Bayesian learning scheme (10), the market-clearing condition (13), and the exogenous process for dividends (4). The solution algorithm is described in Appendix B.

3 Heterogeneous beliefs and speculative bubbles

In this section, we explore the implications of heterogeneity due to agents being born at 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 month, which represents a compromise: dividends typically are announced quarterly, whereas stock prices are available at a much higher frequency.

Dynasts discount future consumption by the factor $\beta\theta$, where β is a time preference parameter and where θ is the probability of survival. The survival rate is set equal to $\theta = 0.996$, implying an “average life on the market” of about 20 years. We use Shiller’s (2005) stock market dataset covering the S&P index from January 1871 to June 2011 to calibrate our model. In particular, consistent with Shiller’s data, we set the mean growth rate of dividends to $\gamma = 0.0027$ per month, and its standard deviation to $\sigma = 0.0114$. We set the time preference parameter to $\beta = 0.998$, consistent with an average price-to-(monthly)-dividend ratio of around 300, close to Shiller’s number of 320. The leverage ceiling parameter is set to $\lambda = 500$. 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 $\lambda = 500$ is consistent with an average stock market participation rate of around 60 percent, which is the estimate reported by Poterba et. al. (1995) for U.S. households with income over \$250,000. Prior uncertainty (or “confidence”) is parameterized by setting $n_0 = 48$, equivalent

⁹If several investors place the same bid, they receive an equal share of the stock. If there is insufficient demand to clear the market because leverage constraints are binding for everyone, then the market closes recording $P_t = P_{t-1}$. Such forced closing of the market only happens initially in the representative-agent version of the model.

to four years (the duration of an undergraduate economics degree) of stock price and dividend observations. For our numerical simulations, we set the number of agents to $N = 100$ and the number of auction rounds per period to $K = 2000$.¹⁰ We perform 1000 Monte Carlo simulations of 5000 months each, equivalent to more than four centuries of trading.

3.2 Simulation results

Figure 1 illustrates the behavior of the asset price according to the model. The thin solid line plots one particular simulated path of the ratio of the stock price in the OLG economy to the REE price. Notice that the ratio oscillates within a 95 percent confidence interval between 0.5 and 1.5, that is, stock price fluctuations are strongly amplified in the OLG model. Second, the *median* stock price in the OLG model does fairly quickly converge to the REE. In that sense, the REE asset price is a relevant statistic for the OLG model. Third, the 95 percent confidence band does not shrink over time, indicating the lack of asymptotic convergence of individual price histories.

The stochastic oscillations of the stock price around the REE are related to the dynamics of learning. To see this, Figure 2 plots the evolution of price growth beliefs held by the cross-section of households relative to the REE belief γ . We plot the median belief, and a 95% confidence interval at each point in time. Notice that individuals' beliefs regarding the growth rate of the stock price do not converge to γ ; instead, they go through successive waves of optimism and pessimism vis-a-vis γ .

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 schemes of successive cohorts of individuals. Namely, at any given date, a fraction of young individuals enters the market whose learning path initially is strongly influenced by the most recent stock price and dividend realizations. The young's 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.¹¹ The combination of the two factors – momentum and trend reversion – results in boom-and-bust cycles that are only loosely related to dividends.

¹⁰We also report results with $N = 1000$ agents, which are very similar.

¹¹Note that trend reversion 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.

Indeed, similar to Harrison and Kreps (1978), asset price cycles in our model are primarily the result of speculation about the future course of the asset price. To see that, we simulated again our economy under a *constant* realization of the dividend growth process, setting all dividend growth innovations to zero. We found that, even in this case, disagreement necessarily arises in investors' beliefs. In particular, over time, investors' assessments of the variance of asset price growth begin to differ because the prior confidence n_{it} of a random fraction $1 - \theta$ of investors is reset from $n_{it} > n_{i0}$ down to $n_{i,t+1} = n_{i0}$ in the updating scheme(10). This, together with the direct effect of n_{it} on expected future prices in (12), necessarily creates dispersion of beliefs and bid prices, which translates into cycles even in the absence of dividend shocks.

Figure 3 shows the sample periodogram and the time series of the price-dividend ratio in the absence of dividend shocks for two economies: one with $N=100$ traders and another with $N=1000$ traders. In the upper panel, the sample periodogram suggests that the economy with more agents displays more frequent price-dividend fluctuations. In the lower panel, the economy with more agents exhibits a smaller amplitude of stock price fluctuations.

In turn, Figure 4 shows the sample periodogram of the price-dividend ratio as a function of the survival rate θ . In the benchmark case with $\theta = 0.996$, the periodogram indicates a series of local maxima at periods of 8 to 16 years. With a higher survival rate of $\theta = 0.998$, the peak periodicity corresponds to around 33 years. While not shown here, the amplitude of the cycles also depends on the survival rate. Namely, stock price cycles have a smaller amplitude with a higher survival rate.

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, stock price changes track one-to-one changes in dividends, inheriting the persistence of dividend growth (zero by assumption). 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 young's overreaction to current information affects the stock price and, progressively, the beliefs of older generations, creating a non-linear feedback, which reinforces the effects of dividend shocks on the stock price.

4 Approximate aggregate dynamics

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. The approximation involves two parts. One part is to approximate the stock price dynamics for a given evolution of average beliefs; the second part is to approximate the evolution of average beliefs. We discuss each of these parts in turn and then combine them to arrive at a stand-alone representative-agent

model that approximates the behavior of the heterogeneous-agents economy.

4.1 Price dynamics

The equilibrium price is obtained by iterating on (12) as $k \rightarrow \infty$ and can be written as

$$P_{t\infty} = \beta\theta \left\{ \exp \left[\gamma_{jt}^P + \frac{\Sigma_{jt}(1,1)}{(2j-6)} \right] P_{t\infty} + \exp \left[\gamma_{jt}^D + \frac{\Sigma_{jt}(2,2)}{(2j-6)} \right] D_t \right\} + \mu_{jt}. \quad (14)$$

A first-order approximation to the above expression is

$$P_t \approx \beta\theta \left[(1 + \gamma_{jt}^P) P_t + (1 + \gamma_{jt}^D) D_t \right] + \mu_{jt}.$$

Taking the average across all age groups yields

$$\begin{aligned} P_t &= \frac{1}{N} \sum_{j=0}^{\infty} f_j P_t = \frac{1}{N} \sum_{i=0}^{\infty} f_j \left[\beta\theta \left((1 + \gamma_{jt}^P) P_t + (1 + \gamma_{jt}^D) D_t \right) + \mu_{jt} \right] \\ &= \beta\theta \left[(1 + \gamma_t^P) P_t + (1 + \gamma_t^D) D_t \right] + \mu_t D_t, \end{aligned} \quad (15)$$

with

$$\gamma_t^P \equiv \frac{1}{N} \sum_{j=0}^{\infty} f_j \gamma_{jt}^P, \quad \gamma_t^D \equiv \frac{1}{N} \sum_{j=0}^{\infty} f_j \gamma_{jt}^D, \quad \mu_t \equiv \frac{1}{N} \sum_{j=0}^{\infty} f_j \frac{\mu_{jt}}{D_t},$$

where γ_t^P and γ_t^D are the average expectations across individuals and where μ_t is the average Lagrange multiplier normalized by the contemporaneous dividend.

In principle, the average Lagrange multiplier should be a function of the price-dividend ratio. When the price-dividend ratio is close to the maximum leverage λ , the multiplier μ_t should turn negative. This sign of the multiplier reflects the fact that most individuals are constrained, and hence the marginal trader is more pessimistic about the future price than the average one. And when the price-dividend ratio is sufficiently low, μ_t should turn positive as the marginal trader is more optimistic than average. To verify this relationship between the price-dividend ratio and the average Lagrange multiplier, Figure 5 shows a cross-plot of the two variables from data generated by our benchmark model. The negative, quasi-linear, relationship can be approximated well by the linear function

$$\mu_t = a - b \frac{P_t}{D_t} \quad (16)$$

describing the behavior of μ_t as a function of the price-dividend ratio. We estimate the parameters a and b by least squares regression from our simulations and report them in Table 1. Both coefficients are significant at the 1 percent level and the regression's R^2 is 0.83 for the case of 1,000 agents.¹²

¹²A more accurate approximation to the dynamics of the average Lagrange multiplier can be obtained by increas-

Using (16), the dynamics of the price-dividend ratio can be approximated as

$$\frac{P_t}{D_t} = \frac{\beta\theta(1 + \gamma_t^D) + a}{1 + b - \beta\theta(1 + \gamma_t^P)}, \quad (17)$$

which depends only on the average expectations about the growth rates of prices and dividends, and on the parameters a and b governing the average Lagrange multiplier.

In the upper half of Table 2, we evaluate the quality of the approximation implied by equations (15) and (17). In the first two lines under the line “Price approximation,” we take the actual average beliefs γ_t^P and γ_t^D given by the benchmark heterogeneous-agents model. In line “actual μ_t ,” we also take as given the actual value of the average multiplier μ_t ; whereas, in line “approx. μ_t ,” we use the approximate μ_t given by the law of motion (16). We consider two metrics of similarity: the correlation between the price-dividend ratio in the heterogeneous-agents model and the approximate model; and the R^2 , defined as one minus the ratio of the variance of the approximation error to the variance of the P_t/D_t ratio in the benchmark model.¹³ We find that in both cases the approximation is reasonable. For example, when using the approximate multiplier and 1,000 agents, the correlation between approximate and actual P_t/D_t ratio is 0.96. The adequacy of the approximation can also be verified visually in the upper panel of Figure 6, which plots the price-dividend ratio from the benchmark model along with the approximation using (16) and (17).

4.2 Average learning dynamics

Equation (17) links the evolution of the stock price to the *average* market beliefs about the growth rate of the stock price and dividends. For a complete stand-alone approximation, we need to approximate the evolution of average beliefs. We begin with the evolution of the average price growth expectation, which is given by

$$\gamma_t^P = \frac{1}{N} \sum_{j=0}^{\infty} f_j \gamma_{jt}^P = \frac{f_0 \gamma}{N} + \frac{1}{N} \sum_{j=1}^{\infty} f_j \left[\gamma_{j-1t-1}^P + \frac{1}{n_{jt}} (\log(P_{t-1}/P_{t-2}) - \gamma_{j-1t-1}^P) \right],$$

where the prior confidence, n_{jt} , is given by

$$n_{jt} = n_0 + j, \quad j \in \mathbb{N}. \quad (18)$$

ing the number of agents in the economy.

¹³We discard the first 2,000 periods of the simulation to avoid the effect of initial conditions.

Assuming that individual expectations are uncorrelated with age, we obtain¹⁴

$$\frac{1}{N} \sum_{j=1}^{\infty} \frac{f_j}{n_0 + j} \gamma_{j-1t-1}^P \approx \left(\frac{1}{N} \sum_{j=1}^{\infty} \frac{f_j}{n_0 + j} \right) \left(\frac{1}{N} \sum_{j=1}^{\infty} f_j \gamma_{j-1t-1}^P \right),$$

where the second product on the right-hand side equals $\theta \gamma_{t-1}^P$. The first product, in turn, can be approximated using the following

Proposition 1 *In the limit as $\theta \rightarrow 1$,*

$$\frac{1}{N} \sum_{j=1}^{\infty} \frac{f_j}{n_0 + j} \approx (1 - \theta).$$

Proof. See the Appendix. ■

Therefore, for θ close to 1, we have

$$\begin{aligned} \gamma_t^P &\approx (1 - \theta)\gamma + \theta [\gamma_{t-1}^P + (1 - \theta) (\log (P_{t-1}/P_{t-2}) / \theta - \gamma_{t-1}^P)] \\ &\approx \gamma_{t-1}^P + (1 - \theta) (\log (P_{t-1}/P_{t-2}) - \gamma_{t-1}^P), \end{aligned} \quad (19)$$

that is, average beliefs about price growth are updated approximately according to a CGL scheme. CGL can thus be viewed as an approximate aggregation of the learning of individuals who update their beliefs by Bayes' rule, 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 (18). The population as a whole, however, learns approximately with a constant gain.

The evolution of average dividend expectations as $\theta \rightarrow 1$ can be derived symmetrically as

$$\gamma_t^D = \frac{1}{N} \sum_{i=0}^{\infty} f_i \gamma_{it}^D \approx \gamma_{t-1}^D + (1 - \theta) [\log (D_{t-1}/D_{t-2}) - \gamma_{t-1}^D]. \quad (20)$$

Note that the value of the gain parameter, which appears in the approximation, equals the retirement probability (1 minus the survival rate). In our baseline calibration, this is equal to 0.004, corresponding to an expected life on the market of 20 years. In quarterly terms, the retirement probability is 0.012, which is quite close to existing estimates of the constant-gain parameter from macro time series data; for example, Milani (2007) estimates the constant-gain parameter to be 0.018 in U.S. data.

In the upper half of Table 2, the line “Price learning” and the line “Dividend learning” evaluate the adequacy of the assumption that individual expectations are uncorrelated with age. In this

¹⁴We evaluate the adequacy of this assumption at the end of the sub-section.

exercise, we take the actual price and dividend sequences from the benchmark heterogeneous-agents model and construct series for stock price and dividend growth expectations using the approximations in (19) and (20). The table shows the two metrics of similarity: the correlation of the approximate with the true average growth expectations, as well as the R^2 defined in sec. 4.1. By these measures, the approximation of both stock price and dividend learning dynamics are reasonably accurate. The middle and the lower panels of Figure 6 confirm this result visually.

4.3 Representative-agent approximation

We now analyze the quality of the two approximations – of the stock price and of the average learning dynamics – as a unit. Namely, we consider the stand-alone representative-agent model in which the stock price is given by equation (17), with parameters a and b from Table 1 (first column), and in which average beliefs follow (19) and (20). This model can be simulated independently for any given evolution of dividends. The line “Price approximation” in the lower half of Table 2 reports the R^2 and correlation with the evolution of the stock price in the benchmark heterogeneous-agents model. As can be expected, the overall approximation deteriorates because approximation errors in the stock price are compounded with errors in the average expectation dynamics. Nevertheless, Figure 7 shows that the overall approximation is still decent; it definitely approximates the benchmark model much better than the rational expectations model (REM) does.

Table 3 evaluates how well the simulated price-dividend ratio matches with the evidence documented by Shiller (2005). The model fits quite well with the observed autocorrelation of the price-dividend ratio by explaining it as a consequence of the dynamic coordination of heterogeneous beliefs. The representative agent constant-gain learning (RA-CGL) approximation produces a smoother price-dividend ratio than the benchmark heterogeneous-agents model.

Finally, Table 4 compares the one-step-ahead forecast errors (9) generated by the heterogeneous agents overlapping-generations (HA-OLG) model, the RA-CGL approximation, and the REM. The distribution of forecast errors is quite similar between the HA-OLG and the RA-CGL models.¹⁵ In particular, the forecast errors for the stock price are unbiased in the HA-OLG and the RA-CGL models but, in both cases, are more dispersed than in the REM. In addition, in the case of the HA-OLG and RA-CGL models, the distribution of price forecast errors displays more leptokurtosis than the REM.

Thus, the HA-OLG and RA-CGL models provide an unbiased average forecast of the evolution of stock prices and dividends, but the uncertainty about the future evolution of prices is larger than that of dividends. This outcome occurs because the stock price depends on market expectations,

¹⁵Because dividends are exogenous, the distribution of the forecast errors for dividends is essentially identical across the three models.

creating self-referential dynamics as emphasized in Eusepi and Preston (2011). In contrast, in the REM, the uncertainty about prices and dividends is essentially the same because agents coordinate ex ante onto the right model for asset pricing.

5 The case with infinitely lived individuals

In this section, we analyze the limiting case in which the probability of survival is $\theta = 1$. We demonstrate the asymptotic convergence of the model to rational expectations despite the fact that individuals do not know anything about each other. We then analyze two properties of the convergence process: its speed and the shape of the convergence path.

5.1 Convergence to rational expectations

The proof of convergence consists of two steps.¹⁶ In the first step, we establish a contemporaneous relationship between the stock price and the dividend, which depends on the current state of beliefs. In the second step, we take the limit as $t \rightarrow \infty$ to establish the asymptotic convergence. The two steps are summarized by the following two propositions, the proofs of which are in Appendix C.

Proposition 2 *The market-clearing stock price P_t is given by*

$$P_t = \begin{cases} \frac{\beta\pi_t^P}{1-\beta\pi_t^P} D_t, & \text{if } \pi_t^P < 1/\beta \\ \lambda D_t, & \text{if } \pi_t^P \geq 1/\beta, \end{cases} \quad (21)$$

where

$$\pi_t^P = \exp\left(\gamma_t^P + \frac{\Sigma_t(1,1)}{2(n_t-3)}\right) \quad \text{and} \quad \pi_t^D = \exp\left(\gamma_t^D + \frac{\Sigma_t(2,2)}{2(n_t-3)}\right). \quad (22)$$

Proof. See the Appendix. ■

Proposition 3 *The stock price $P_t \rightarrow P_t^{REE} \equiv \frac{\beta\theta \exp(\gamma+\sigma^2/2)}{1-\beta\theta \exp(\gamma+\sigma^2/2)} D_t$.*

Proof. See the Appendix. ■

5.2 Speed and shape of the convergence path

Having established asymptotic convergence, it is useful to know how long it takes for the stock price to converge to the REE.¹⁷ Figure 7 plots one randomly drawn path of the ratio of the stock

¹⁶For a related proof for the case of least squares learning using a projection facility, see Adam, Marcet, and Nicolini (2008).

¹⁷In different contexts, this question has been studied, for example, by Vives (1993), Marcet and Sargent (1992), and Ferrero (2006). Evans and Honkapohja (2003, ch. 15) establish that in recursive least squares learning for gain sequences of the form $t^{-\chi}$ the speed of convergence is asymptotically $t^{\chi/2}$.

price to its REE counterpart, the median across simulations, and the 95 percent confidence band.¹⁸ Remarkably, after 100 years of trading, the median stock price is still about 30 percent above the REE price. That is, even though there is asymptotic convergence, it takes a very long time for the rational expectations model to become a good approximation to the short-run dynamics generated by our model.

The convergence path is characterized by an initial “overshooting” of the stock price above the REE. Because individual learning begins with the REE as a prior belief, initially agents overestimate the growth rate of the stock price. This overestimation occurs because individuals observe greater stock price volatility than their prior belief suggests. Thus, the initial rise in the price-dividend ratio is self-fulfilling: The stock price rises because agents expect it to rise, which generates an further increase in the stock price until the constraint $P_t/D_t \leq \lambda$ is reached. The stock price remains at this level for some time, as agents progressively revise down their beliefs, eventually pulling the price back toward the REE.

The individual exposure constraints (3) are therefore central for the convergence process. They amount to a practical implementation of the standard transversality condition, which rules out asset price bubbles in infinite horizon models. The looser the constraint is (the larger is λ), the larger the initial overshooting and the longer it takes for the market to converge back to the REE. Another way to see this outcome is illustrated in Figure 9, which plots the convergence in mean squared error (MSE) of the ratio of the stock price to the REE price over time. MSE is consistently higher than in the baseline calibration when the exposure constraint is relaxed by 10 percent ($\lambda = 550$).

Figure 9 also illustrates how prior uncertainty affects the convergence. In particular, we set the confidence parameter to $n_0 = 240$, equivalent to 20 years of prior observations of the REE outcome. Qualitatively, the convergence is similar to the baseline calibration with $n_0 = 48$, with initial price overshooting followed by progressive convergence to the REE price. However, the convergence is now faster so that after 40 years, the median stock price is less than 10 percent away from the REE.¹⁹

6 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,

¹⁸In this exercise with $\theta = 1$, we need to recalibrate the time preference parameter to $\beta = 0.994$ to make the model’s output consistent with Shiller’s evidence.

¹⁹The initial beliefs are assumed to be centered on the REE. As a robustness check, we simulated the model with biased prior beliefs. The results (not reported here) are qualitatively similar to the benchmark case.

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 data realized during their own lifetimes than to earlier historical information.

We extend the model of Adam and Marcet to a 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. Even a tiny degree of “learning from experience” is sufficient to generate chaotic dynamics, which roughly resemble what we find in the data.

We show that 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. To a first-order approximation the gain parameter equals the survival rate, reflecting the fact that historical data is lost when successive generations “learn from experience.” This result provides a plausible 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 needed discipline by tying the gain parameter to the survival rate.

Finally, we show that in the limiting case with infinitely lived agents, individuals can coordinate through a centralized market, and, eventually, achieve convergence to the REE. The only requirement for the equilibrium to be stationary are bounds on asset exposure that prevent coordination to an explosive path. This requirement is akin to the way transversality conditions are imposed in standard representative-agent models. We show that, for a plausible parameterization, the market converges very slowly to rational expectations. Moreover, the speed of convergence is strongly affected not only by the prior beliefs but also by the tightness of the exposure constraints.

References

- [1] Adam, K. and A. Marcet. 2011. “Internal Rationality, Imperfect Market Knowledge and Asset Prices,” *Journal of Economic Theory*, In Press.
- [2] Adam, K., A. Marcet, and J. P. Nicolini. 2008. “Stock Market Volatility and Learning,” European Central Bank Working Paper No. 862.
- [3] 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.
- [4] Angeletos, G. M. and J. La’O. 2011. “Decentralization, Communication, and the Origins of Fluctuations,” Working paper, MIT
- [5] Arrow, K.J. 1986. “Rationality of Self and Others in an Economic System,” *The Journal of Business*, vol. 59 (4), S385–99.
- [6] 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.
- [7] Branch, W. A., and G. W. Evans. 2006. “Intrinsic Heterogeneity in Expectation Formation,” *Journal of Economic Theory*, vol. 127, pp. 264–95.
- [8] 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.
- [9] 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.
- [10] Brown, A., and L. Rogers. 2009. “Heterogeneous Beliefs with Finite-Lived Agents,” University of Cambridge, Statistical Laboratory, Working Paper.
- [11] Brown, A., and L. Rogers. 2009. “Diverse Beliefs,” University of Cambridge, Statistical Laboratory, Working Paper.
- [12] Carceles-Poveda, E., and C. Giannitsarou. 2007. “Adaptive Learning in Practice,” *Journal of Economic Dynamics and Control*, vol. 31 (8), pp. 2659-97.
- [13] 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.
- [14] DeGroot, M. H. 1970. *Optimal Statistical Decisions*. McGraw-Hill Press.

- [15] Eusepi, S., and B. Preston. 2011. “Expectations, Learning and Business Cycle Fluctuations,” *American Economic Review*, forthcoming.
- [16] Evans, G. W. and S. Honkapohja. 2003. *Learning and Expectations in Macroeconomics*, Princeton University Press.
- [17] Ferrero, G. 2007. “Monetary Policy, Learning and the Speed of Convergence,” *Journal of Economic Dynamics and Control*, vol. 31 (9), pp. 3006–41.
- [18] Giannitsarou, C. 2003. “Heterogeneous Learning,” *Review of Economic Dynamics*, vol. 6 (4), pp. 885–906.
- [19] Graham, L. 2011. “Learning, Information and Heterogeneity,” Working paper.
- [20] Hamilton, J. D., 1994. *Time Series Analysis*, Princeton University Press.
- [21] 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.
- [22] Honkapohja, S., and K. Mitra. 2006. “Learning Stability in Economies with Heterogeneous Agents,” *Review of Economic Dynamics*, vol. 9 (2), pp. 284–309.
- [23] Lucas Jr., R. E. 1972. “Expectations and the Neutrality of Money,” *Journal of Economic Theory*, vol. 4 (2), pp. 103–24.
- [24] Malmendier, U. and S. Nagel. 2011. “Depression Babies: Do Macroeconomic Experiences Affect Risk-Taking?,” *Quarterly Journal of Economics*, vol. 126 (1), pp. 373–416.
- [25] Malmendier, U. and S. Nagel. 2009. “Learning from Inflation Experiences,” Unpublished manuscript, U.C. Berkley
- [26] 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.
- [27] Milani, F., 2007. “Expectations, Learning and Macroeconomic Persistence,” *Journal of Monetary Economics*, vol. 54, pp. 2065–82.
- [28] 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.

- [29] Radner, R. 1979. "Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices," *Econometrica* 47 (3), pp. 655-678.
- [30] Shiller, R. 2005. *Irrational Exuberance*, 2nd edition, Princeton University Press.
- [31] 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.
- [32] Vives, X. 1993. "How Fast do Rational Agents Learn?," *Review of Economic Studies*, vol. 60 (2), pp. 329-47.
- [33] Weitzman, M. L. 2007. "Subjective Expectations and Asset-Return Puzzles," *American Economic Review*, vol. 97 (4), 1102-30.

Appendix A: The model in more detail

Stock holding decision

The first-order optimality conditions of the individual's problem are:

$$\text{if } P_t < P_{it}, \text{ then } S_{it} = \bar{E}_t/P_t \quad (23a)$$

$$\text{if } P_t = P_{it}, \text{ then } S_{it} \in [\underline{E}_t/P_t, \bar{E}_t/P_t], \quad (23b)$$

$$\text{if } P_t > P_{it}, \text{ then } S_{it} = \underline{E}_t/P_t, \quad (23c)$$

$\forall t$, and $\forall \omega^t \in \Omega^t$, where

$$P_{it} = \beta\theta E_{it} (P_{t+1} + D_{t+1}) \quad (24)$$

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).

Symmetric rational expectations equilibrium

If individuals were identical, and this fact were common knowledge, they would be able to compute the equilibrium asset price by deduction. Namely, dividing (24) by the current dividend, dropping the i subscript, and iterating the resulting equation forward while applying the law of iterated expectations and taking into account the known process for dividends (4), yields:

$$\frac{P_t}{D_t} = \beta\theta E_t \left[\frac{D_{t+1}}{D_t} \left(1 + \frac{P_{t+1}}{D_{t+1}} \right) \right] = \sum_{j=1}^{\infty} (\beta\theta)^j e^{j(\gamma+\sigma^2/2)} + \lim_{T \rightarrow \infty} E_t \left((\beta\theta)^T \frac{D_{t+T}}{D_{t+T-1}} \frac{P_{t+T}}{D_{t+T}} \right). \quad (25)$$

Given that the sum of stock holdings must equal the fixed supply of the stock N , it follows from (3) that the price-dividend ratio is bounded above by λ ,

$$N = \sum_{i=1}^N S_{it} \leq N\lambda D_t/P_t \implies P_t/D_t \leq \lambda. \quad (26)$$

Hence the last term in (25) is zero, and therefore the equilibrium asset price is given by

$$P_t^{REE} = \frac{\beta\theta \exp(\gamma + \sigma^2/2)}{1 - \beta\theta \exp(\gamma + \sigma^2/2)} D_t > 0, \quad (27)$$

where dividends follow the exogenous stochastic process defined in (4). We further impose the

parameter restrictions,

$$\beta\theta \exp(\gamma + \sigma^2/2) < 1 \text{ and } \lambda > \frac{\beta\theta \exp(\gamma + \sigma^2/2)}{1 - \beta\theta \exp(\gamma + \sigma^2/2)}, \quad (28)$$

which ensure that the price-dividend ratio is finite and that it is not a corner solution due to binding leverage constraints (26).

Informational limitations

We depart from REE 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 γ 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.

Formally, following Adam and Marcet (2011), denote by the operator E_{i0} investor i 's subjective expectation defined in 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) . Denote by Ω_i^t the set of histories *during the lifetime* of agent i up to period t , and let $\omega_i^t \in \Omega_i^t$. When investor i chooses his stock holding in period t , he takes as given Π_i and his choice is contingent on ω_i^t . The space of realizations is

$$\Omega \equiv \Omega_P \times \Omega_D, \quad (29)$$

where Ω_P contains all possible sequences of stock prices and where Ω_D contains all possible dividend sequences. Individuals can thus condition their investment decision on all possible combinations of dividend and stock price histories. 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.

Appendix B: Simulation algorithm

We briefly sketch the algorithm used to perform a single Monte Carlo simulation of the model:

1. Generate an exogenous series for dividends D_t following (4) and assuming that $D_{-1} = 1$. Set $P_{-1} = P_{-1}^{REE}$ and $P_0 = P_0^{REE}$, where P_t^{REE} is given by (27).
2. Initialize the prior beliefs, γ_{i0} , Σ_{i0} , and n_{i0} , for all agents following (7) and (8).
3. **Main loop.** At each point in time $t = 1, \dots, T$, for all N agents:
 - (a) Compute the one-step-ahead forecast errors \mathbf{e}_{it} using (9)
 - (b) Draw a vector of random numbers from a uniform distribution between 0 and 1. For values greater than θ , the agent retires; otherwise he survives to the following period (the case of infinitely lived agents is nested by setting $\theta = 1$).
 - (c) If an agent survives, update his beliefs, γ_{it} , Σ_{it} , and n_{it} , using (10). If he retires (he is replaced by a new agent), set $\gamma_{it} = \gamma_{i0}$, $\Sigma_{it} = \Sigma_{i0}$, and $n_{it} = n_{i0}$.
 - (d) Set the initial auction price to $P_{t0} = P_{t-1}$.
 - (e) Compute the reservation price for each agent in auction round zero P_{it0} using (11).
 - (f) **Auction.** For each auction round $k = 1, \dots, K$:
 - i. Sort the reservation prices P_{itk-1} in decreasing order and notionally allocate the amount $S_{itk} = \lambda \frac{D_t}{P_{itk-1}}$ to each agent until the entire stock N of the asset gets allocated. To ensure that the total does not exceed N , the marginal agent to receive a share of the asset may receive $S_{itk} < \lambda \frac{D_t}{P_{itk-1}}$. The reservation price of the marginal agent is denoted as P_{itk-1}^* .
 - ii. If $\sum_{i=1}^N S_{it} \leq N$, then set $P_{tk} = P_{t-1}$. Otherwise, set $P_{tk} = P_{itk-1}^*$.
 - iii. The reservation price of each agent in round k , P_{itk} , is computed using (12).
 - (g) The auction is over in round K , and the stock price in period t is $P_t = P_{itK}$.
4. Repeat the main loop (3) for periods $t = 1, \dots, T$.

Appendix C: Proofs

Proposition 1. In the limit as $\theta \rightarrow 1$,

$$\frac{1}{N} \sum_{j=1}^{\infty} f_j \frac{1}{n_0 + j} \approx (1 - \theta).$$

Proof. First we compute the series

$$\begin{aligned} \sum_{j=1}^{\infty} \frac{(1 - \theta)\theta^j}{n_0 + j} &= (1 - \theta)\theta^{-n_0} \sum_{j=n_0+1}^{\infty} \frac{\theta^j}{j} = (1 - \theta)\theta^{-n_0} \sum_{j=n_0+1}^{\infty} \int \theta^{j-1} d\theta \\ &= (1 - \theta)\theta^{-n_0} \int \left(\sum_{j=n_0+1}^{\infty} \theta^{j-1} \right) d\theta = (1 - \theta)\theta^{-n_0} \int \frac{\theta^{n_0}}{(1 - \theta)} d\theta, \end{aligned}$$

as $n_0 \in \mathbb{N}^+$, the integral $\int \frac{\theta^{n_0}}{(1 - \theta)} d\theta$ can be expressed as

$$\int \frac{\theta^{n_0}}{(1 - \theta)} d\theta = Q_{n_0}(\theta) - \log(1 - \theta),$$

where $Q_{n_0}(\theta)$ is a polynomial of order n_0 . Therefore, the limit as $\theta \rightarrow 1$ is

$$\lim_{\theta \rightarrow 1} \sum_{j=1}^{\infty} \frac{(1 - \theta)\theta^j}{n_0 + j} = \lim_{\theta \rightarrow 1} [(1 - \theta)\theta^{-n_0}] [Q_{n_0}(\theta) - \log(1 - \theta)],$$

which can be solved by applying L'Hôpital's rule,

$$\begin{aligned} \frac{d}{d\theta} \frac{\theta^{n_0}}{(1 - \theta)} &= \frac{n_0\theta^{n_0-1}(1 - \theta) + \theta^{n_0}}{(1 - \theta)^2}, \\ \frac{d}{d\theta} [Q_{n_0}(\theta) - \log(1 - \theta)] &= \left[Q'_{n_0}(\theta) + \frac{1}{1 - \theta} \right], \end{aligned}$$

where $Q'_{n_0}(\theta) = \frac{d}{d\theta} Q_{n_0}(\theta)$. Then, we take the limit as $\theta \rightarrow 1$

$$\lim_{\theta \rightarrow 1} [(1 - \theta)\theta^{-n_0}] [Q_{n_0}(\theta) - \log(1 - \theta)] = \lim_{\theta \rightarrow 1} \frac{(1 - \theta)^2 \left[Q'_{n_0}(\theta) + \frac{1}{1 - \theta} \right]}{n_0\theta^{n_0-1}(1 - \theta) + \theta^{n_0}} = \lim_{\theta \rightarrow 1} \frac{1 - \theta}{n_0\theta^{n_0-1}(1 - \theta) + \theta^{n_0}}.$$

■

Proposition 2. The stock price that clears the market at time t is given by

$$P_t = \begin{cases} \frac{\beta\pi_t^D}{1 - \beta\pi_t^P} D_t, & \text{if } \pi_t^P < 1/\beta \\ \lambda D_t, & \text{if } \pi_t^P \geq 1/\beta \end{cases} \quad (30)$$

where

$$\pi_t^P = \exp\left(\gamma_t^P + \frac{\Sigma_t(1,1)}{2(n_t-3)}\right) \quad \text{and} \quad \pi_t^D = \exp\left(\gamma_t^D + \frac{\Sigma_t(2,2)}{2(n_t-3)}\right). \quad (31)$$

Proof. Because individuals are identical we can drop index i . In the initial round of the auction at time t , then, the price is given by

$$P_{t0} = \beta (2\pi_t^P P_{t-1} + 2\pi_t^D D_{t-1}), \quad (32)$$

where (31) holds. In the subsequent rounds, the price evolves as

$$P_{tk} = \beta (\pi_t^P P_{tk-1} + \pi_t^D D_t) = (\beta\pi_t^P)^k P_{t0} + \beta\pi_t^D D_t \sum_{i=1}^{k-1} (\beta\pi_t^P)^i. \quad (33)$$

If $\beta\pi_t^P \geq 1$, then as $k \rightarrow \infty$, P_{tk} would grow unboundedly were it not for constraint (3) that prevents explosive beliefs by effectively setting an upper (and a lower) limit on the price-to-dividend ratio, and hence $P_{tk} = \lambda D_t$. If $\beta\pi_t^P < 1$, then in the limit as $k \rightarrow \infty$, the first term in equation (33) tends to zero and the price for period t is

$$P_t \equiv \lim_{k \rightarrow \infty} P_{tk} = \frac{\beta\pi_t^D}{1 - \beta\pi_t^P} D_t. \quad (34)$$

■

Proposition 3. The economy converges to the REE with stock price P_t^{REE} defined in (27).

Proof. First, because dividends follow an exogenous process, the Bayesian learning algorithm for dividends must converge asymptotically to the true value of the parameters

$$\lim_{t \rightarrow \infty} \gamma_t^D = \gamma, \quad \lim_{t \rightarrow \infty} \frac{\Sigma_t(2,2)}{n_t-3} = \sigma^2, \quad \text{and} \quad \lim_{t \rightarrow \infty} \pi_t^D = \exp(\gamma + \sigma^2/2). \quad (35)$$

Second, given the equilibrium price (21), the value of $\log(P_{t-1}/P_{t-2})$ is bounded as $t \rightarrow \infty$. Therefore, given the Bayesian updating scheme, π_t^P must converge ,

$$\begin{aligned} \lim_{t \rightarrow \infty} \pi_t^P &= \lim_{t \rightarrow \infty} \exp\left(\gamma_t^P + \frac{\Sigma_t(1,1)}{2(n_t-3)}\right) \\ &= \lim_{t \rightarrow \infty} \exp\left[\gamma_{t-1}^P + \frac{\log(P_{t-1}/P_{t-2}) - \gamma_{t-1}^P}{1+n_{t-1}} + \frac{\Sigma_{t-1}(1,1)}{2(1+n_{t-1}-3)} + \right. \\ &\quad \left. + \frac{n_{t-1}(\log(P_{t-1}/P_{t-2}) - \gamma_{t-1}^P)^2}{2(1+n_{t-1}-3)(1+n_{t-1})}\right] \\ &= \lim_{t \rightarrow \infty} \pi_{t-1}^P \equiv \pi^P. \end{aligned} \quad (36)$$

Third, the limit π^P must satisfy

$$\lim_{t \rightarrow \infty} \pi_t^P = \pi^P < 1/\beta \quad (37)$$

This last point can be proved by contradiction: suppose $\pi^P \geq 1/\beta$. Then, all individual constraints (3) must be binding, so that (26) is binding as well, and

$$\lim_{t \rightarrow \infty} \log(P_t/P_{t-1}) = \lim_{t \rightarrow \infty} \log(\lambda D_t/(\lambda D_{t-1})) = \exp(\gamma + \sigma^2/2) < 1/\beta$$

by (35) and (28); thus, we have reached a contradiction.

Finally, by taking the log-difference of (34),

$$\lim_{t \rightarrow \infty} \log(P_t/P_{t-1}) = \lim_{t \rightarrow \infty} \left\{ \log \left[\frac{\pi_t^D (1 - \beta \pi_{t-1}^P)}{\pi_{t-1}^D (1 - \beta \pi_t^P)} \right] + \log(D_t/D_{t-1}) \right\}. \quad (38)$$

Together, (35) and (37) imply that the first term in the brackets on the right-hand side of (38) converges to zero, and hence the learning parameters for the stock price must also converge to the asymptotic values of the REE,

$$\lim_{t \rightarrow \infty} \gamma_t^P = \gamma, \quad \lim_{t \rightarrow \infty} \frac{\Sigma_t(1, 1)}{(n_t - 3)} = \sigma^2, \quad \text{and} \quad \lim_{t \rightarrow \infty} \pi_t^P = \exp(\gamma + \sigma^2/2). \quad (39)$$

Substituting the above in equation (34) we obtain $\lim_{t \rightarrow \infty} P_t = \frac{\beta \exp(\gamma + \sigma^2/2)}{1 - \beta \exp(\gamma + \sigma^2/2)} D_t = P_t^{REE}$. ■

Appendix D. Tables and figures

Table 1. Average Lagrange multiplier

Coeff.	1000 agents	100 agents
a	0.9367	0.9221
s.e.	(0.009)	(0.010)
b	-0.0035	-0.0033
s.e.	(0.0001)	(0.0001)
R^2	0.83	0.77
Obs.	3000	3000

Note: Standard errors are in parenthesis.

The table reports the coefficients a and b from an OLS regression $\mu_t = a + b \frac{P_t}{D_t} + \varepsilon_t$.

Table 2. Approximation accuracy under different assumptions

	1000 agents		100 agents	
	R^2	Correl.	R^2	Correl.
Single approximation step				
Price approximation				
With actual μ_t	0.76	0.89	0.86	0.95
With approx. μ_t	0.92	0.96	0.91	0.96
Price learning	0.88	0.94	0.88	0.94
Dividend learning	0.88	0.94	0.89	0.94
Complete RA-CGL model				
Price approximation	0.67	0.83	0.59	0.80
Price learning	0.58	0.75	0.51	0.72
Dividend learning	0.88	0.94	0.88	0.94

Note: The sample consists of 3000 simulated observations of the benchmark model.

RA-CGL stands for “representative agent constant-gain learning.”

Table 3. Moments of the price-dividend ratio

	Data	REM	HA-OLG	RA-CGL
Mean	320.3	307.6	316.8	309.5
Standard deviation	166.1	0	65.4	56.9
Autocorrelation	0.996	—	0.995	0.989

Note: REM stands for “rational expectations model.”

HA-OLG stands for “heterogeneous agents overlapping generations.”

RA-CGL stands for “representative agent constant-gain learning.”

Table 4. Moments of the forecast errors

	Mean	Std. Dev.	Skewness	Kurtosis
Price forecast errors				
REM	9.0×10^{-5}	0.0114	-0.0148	3.0753
HA-OLG	5.2×10^{-5}	0.0207	0.1346	3.4809
RA-CGL	0.8×10^{-5}	0.0208	0.1300	3.4934
Dividend forecast errors				
REM	9.0×10^{-5}	0.0114	-0.0148	3.0753
HA-OLG	3.4×10^{-5}	0.0114	-0.0149	3.0605
RA-CGL	1.9×10^{-5}	0.0114	-0.0153	3.0640

Note: The sample consists of 3000 simulated observations of the benchmark model.

REM stands for “rational expectations model.”

HA-OLG stands for “heterogeneous agents overlapping generations.”

RA-CGL stands for “representative agent constant-gain learning.”

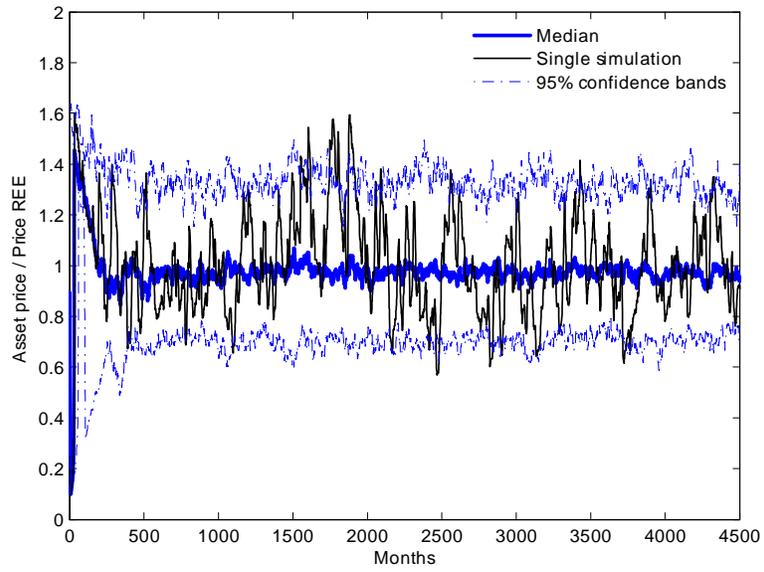


Figure 1: Stock price divided by the rational expectations price

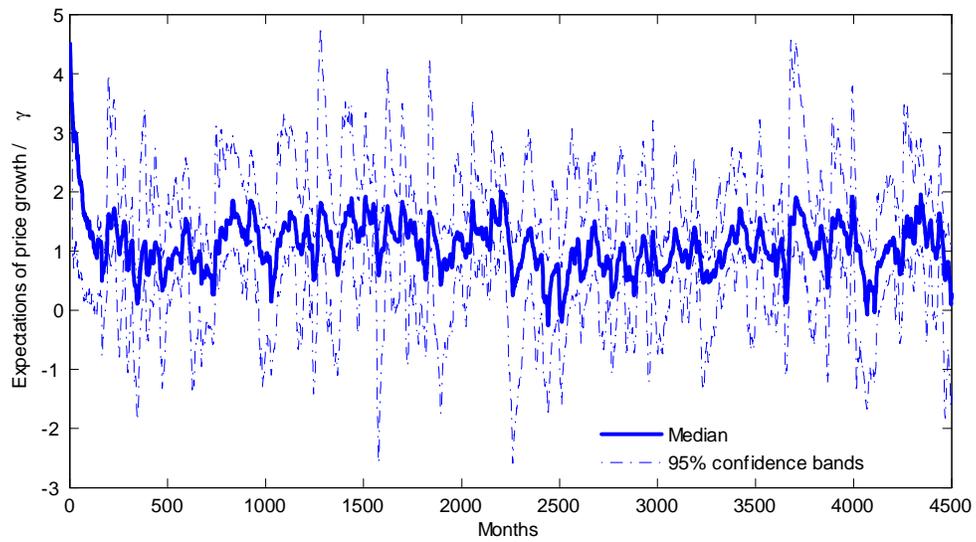


Figure 2: Expectations of stock price growth relative to REE

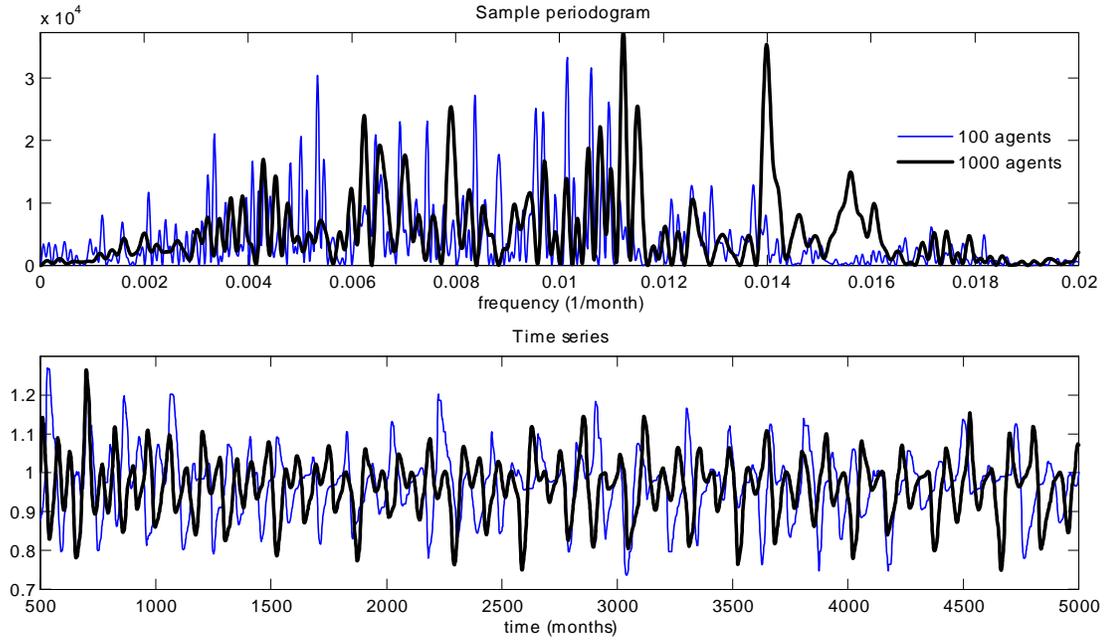


Figure 3: Price-dividend ratio fluctuations in the absence of dividend shocks as a function of the number of traders

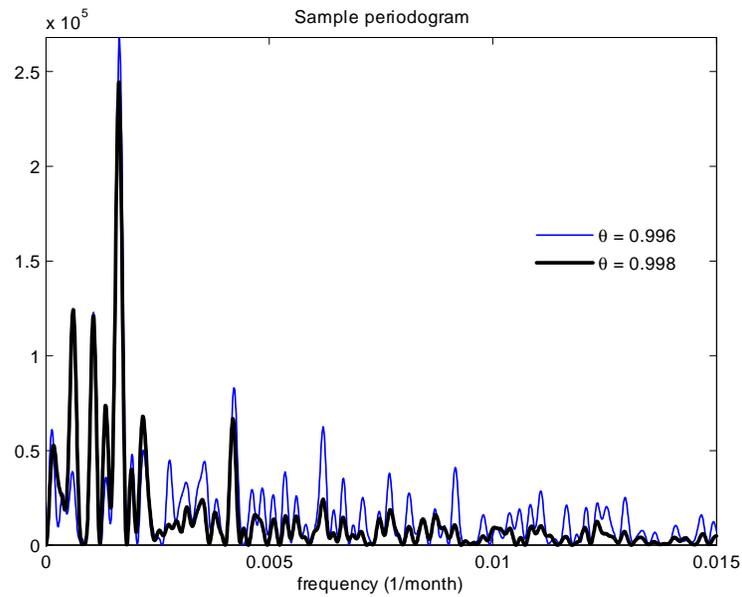


Figure 4: Periodogram of the price-dividend ratio as a function of the survival rate

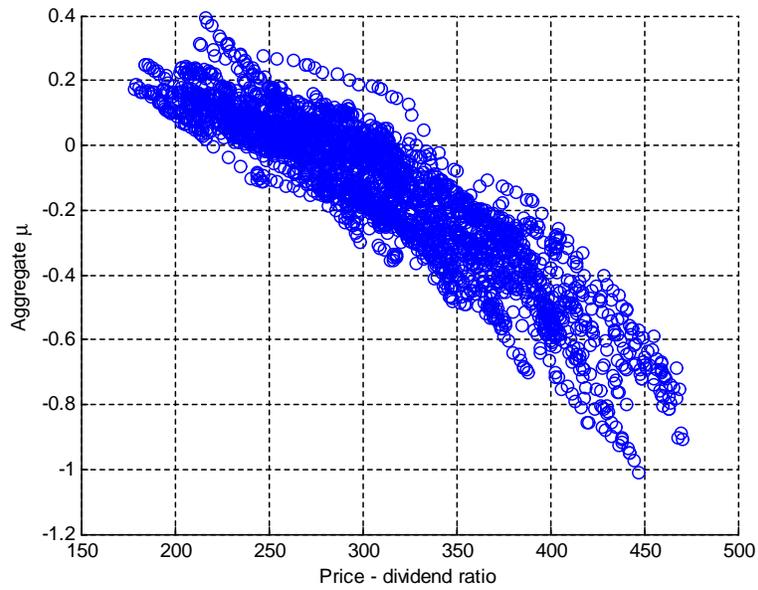


Figure 5: Relationship between the average Lagrange multiplier μ_t and the P_t/D_t ratio

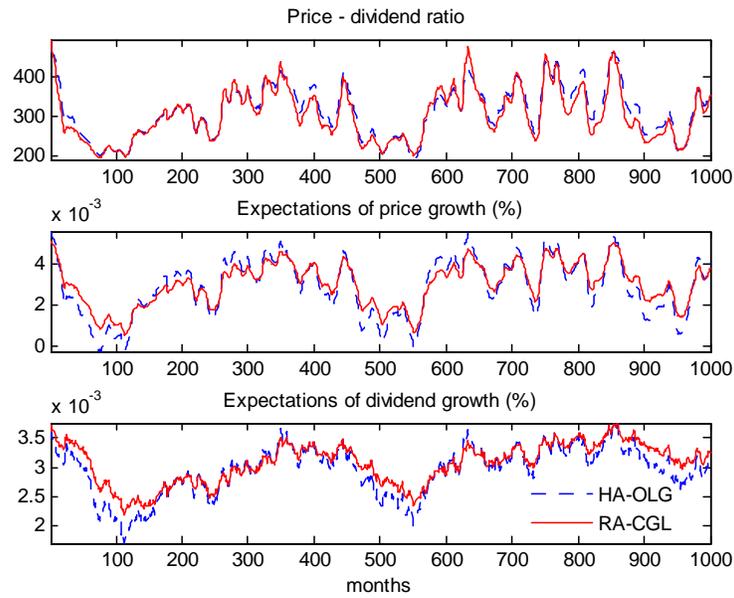


Figure 6: Heterogeneous-agents model vs. single approximation step

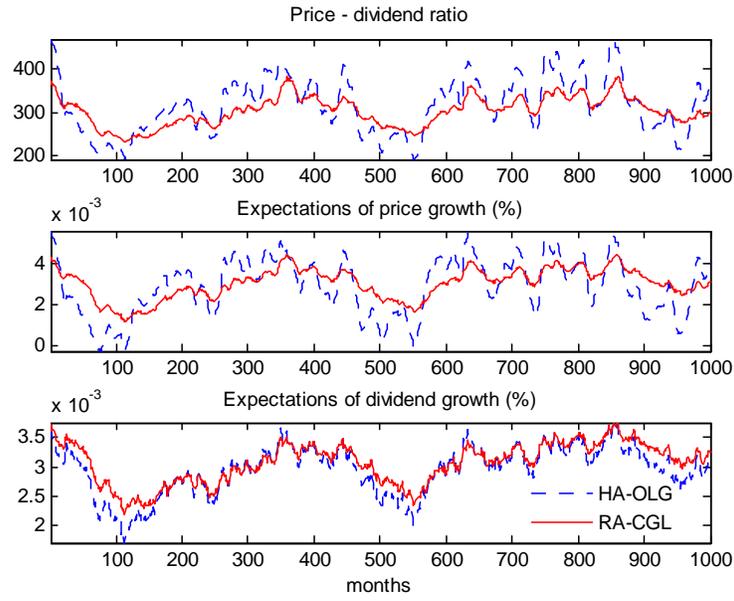


Figure 7: Heterogeneous-agents model vs. full representative-agent approximation

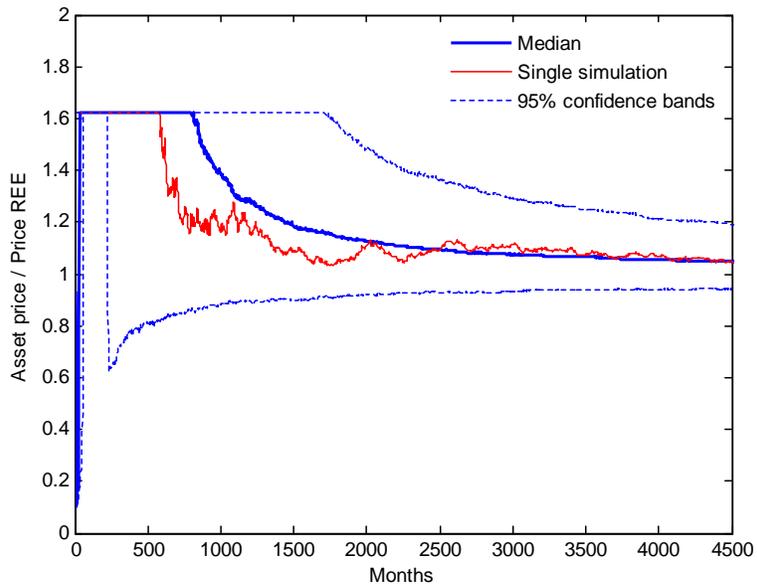


Figure 8: Convergence to rational expectations with infinitely-lived agents

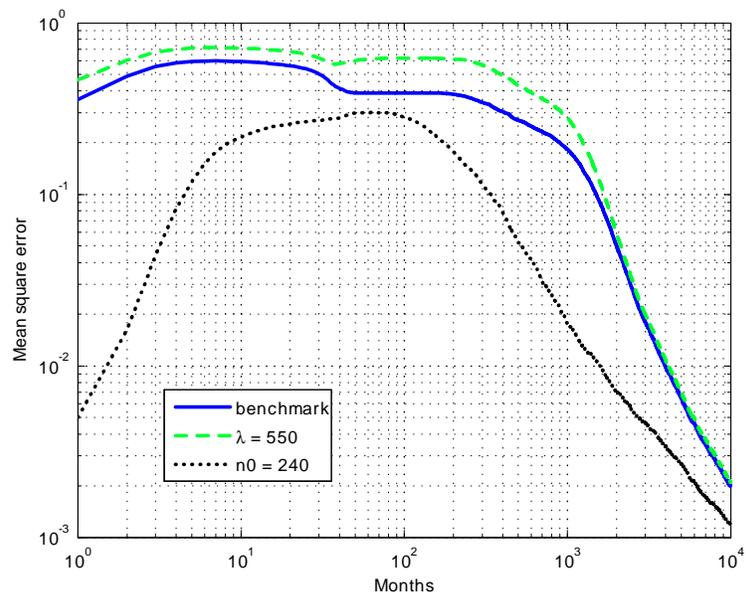


Figure 9: Robustness of the convergence to changes in leverage and in prior precision