Bayesian Foundations of Constant-Gain Learning *

Anton Nakov
Bank of Spain

Galo Nuño
European Central Bank

November 6, 2012

Abstract

Constant-gain learning (CGL) models replace the assumption of model-consistent expectations by a near-rational expectation formation mechanism in which the agents estimate over time the coefficients of a perceived law of motion for economic variables. A crucial parameter under CGL is the gain coefficient, which governs the rate at which agents discount past information when forming expectations. In this paper, we provide a microfoundation for the use of CGL based on Bayesian theory. In particular, we demonstrate how, under certain conditions, CGL expectations can approximate the average beliefs of heterogeneous finitely-lived agents who "learn from experience". That is, individuals update their expectations by Bayesian methods based on observations from their own lifetimes. In this case, the constant gain coefficient is a function of the average life-expectancy and of the precision of the priors. We study the example of a Lucas-tree asset-pricing model where individuals learn from experience about stock prices and dividends. In this model, the above approximation is valid and the stock price expectations display stochastic fluctuations around the rational expectations solution due to waves of optimism and pessimism loosely related to fundamentals.

Keywords: heterogeneous beliefs, agent-based models, Bayesian learning, OLG, asset pricing

JEL codes: G12, D83, D84

*A previous version of this paper was circulated as "Learning from Experience in the Stock Market." We are grateful for comments and suggestions to George-Marios Angeletos, Gunter Coenen, Alistair Dieppe, Oliver de Groot, Miguel Angel Iraola, 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. Anton Nakov thanks the European Central Bank and the Federal Reserve Board for their 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, European Central Bank, Kaiserstrasse 29, 60311 Frankfurt am Main, Germany. E-mail: galo.nuno@ecb.europa.eu

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.

The idea that people may deviate from the REH has been explored by the literature on constant-gain learning (CGL).[2] This literature typically assumes that a representative agent uses a linear perceived law of motion for economic variables with the same structural form as the solution under the REH and that the agent learns model coefficients over time. CGL is usually motivated based on its ability to produce realistic model features, such as amplification of the persistence of macroeconomic variables in response to aggregate shocks. Rarely is there a discussion of the reasons why all agents should learn in the same suboptimal way. A crucial parameter under CGL is the gain coefficient, which governs the rate at which agents discount past information when forming expectations. The value of the gain parameter is either estimated from the data or calibrated to yield the smallest possible mean-squared forecasting error.

In this paper, we provide an alternative justification for using CGL in a representative-agent context. Namely, we see it as a useful approximation to the aggregate dynamics of an economy populated by many finetely-lived Bayesian learners, each of them using a decreasing gain sequence, under the assumption that they "learn from experience." Specifically, we assume that a small random fraction of individuals retire 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 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

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

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

stock prices and dividends which they observe during their own lifetimes.

We show that, under certain conditions, 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 and of the initial prior precision. The 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."

After deriving the general result, we illustrate it with an example. In particular, we show how replacing the REH with "learning from experience" modifies the results of a simple general equilibrium model of the stock market. 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. 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 analyze the model quantitatively, we propose a computational method that allows us to solve the equilibrium with heterogeneous agents. 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 expectations equilibrium (REE). Two forces create 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).

After discussing the dynamics of the heterogenous agents economy, we analyze numerically the validity of our approximation of the average belief by a CGL algorithm. We find that the approximation is quite good in terms of both R-squared and correlation. We also show how, both in the model and in the approximation, the mean forecast error is of the same order of magnitude as in the REE. Moreover, the forecast errors are unbiased meaning that, despite not using all available information, Bayesian agents do not make systematic errors, maintaining one of the appealing properties of the REH.

Finally, we analyze the asymptotic behavior of our economy in the limit with infinitely lived agents ($\theta \to 1$). We show analytically that, in this case, even if traders do 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 about 30 percent higher than its REE counterpart. This is important because it means that, although the REE is the theoretical limit of a setup where agents use all available information, it may take too long to reach it in a realistic environment.[3]

Our paper is related to several strands of the literature. 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 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.

[3]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.

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.

Finally, this paper relates to the literature on agent-based economic models. An introduction to this literature can be found in Testfatsion and Judd (2006). These models typically simulate the simultaneous operations and interactions of multiple autonomous agents, in an attempt to re-create the appearance of complex phenomena. This paper tries to provide a bridge between this literature and the more *mainstream* approach based on dynamic stochastic general equilibrium (DSGE) models. We do so by presenting an algorithm to simulate an asset-pricing model with autonomous Bayesian agents and demonstrating how the beliefs dynamics are similar to those generated by a representative agent learning from experience.

The rest of our paper is organized as follows. In section 2, we introduce the general framework and the main theoretical results. In section 3, we recast the model of Adam and Marcet (2011) in an OLG setting. In section 4, we calibrate the model, analyze the properties of "learning from experience" and show how the model can be approximated by a representative agent with CGL. Finally, section 6 concludes.

2 Bayesian foundations of constant-gain learning

In this section we show how, under certain conditions, the average belief in an OLG model with Bayesian agents can be approximated by a CGL algorithm. The framework is quite general and can be applied to a variety of models.

First, we assume that the economy is populated by N risk-neutral ex ante identical dynasties. Members of each dynasty have stochastic lifetimes with death 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 $f_j = N(1-\theta)\theta^j$.

Second, we depart from REE by assuming that individuals have only limited information about the world they live in. Let $X_t \in \mathbb{R}^K$ denote the K-vector of state variables that determine the state of the economy at given time t. The perceived law of motion (PLM) of X_t is given by a process

$$X_t = \boldsymbol{\gamma} + \varepsilon_t, \tag{1}$$

where $\varepsilon_t \in \mathbb{R}^K$ is a vector of independent exogenous disturbances of zero mean with diagonal covariance matrix $\Sigma \in \mathbb{R}^{K\times K}$, and $\boldsymbol{\gamma} \in \mathbb{R}^K$ is a vector the values of which are estimated by the individuals.

In order to estimate $\boldsymbol{\gamma}$, we assume that agents employ the available information about the past realization of the state variables, $\omega^t = \{X_\tau\}_{\tau=1}^t$, in an optimal (Bayesian) way. In particular, we assume that the joint prior distribution of $\boldsymbol{\gamma}$ and Σ is a *normal-Wishart*:

$$\Sigma^{-1} \sim W\left(\boldsymbol{\Sigma}_0, n_0\right) \quad \text{and} \quad \boldsymbol{\gamma}|\Sigma^{-1} \sim N\left(\boldsymbol{\gamma}_0, \Sigma/n_0\right), \tag{2}$$

where the Wishart distribution W with precision matrix $\boldsymbol{\Sigma}_0$ and $n_0 > 3$ 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, conditional on the precision matrix Σ^{-1}. The vector $\boldsymbol{\gamma}_0$ denotes the conditional prior mean, while n_0 is the precision of prior beliefs. Individuals are assumed to be born with identical prior beliefs

Third, we introduce the concept of "learning from experience". This means that upon retirement, a successor inherits the assets of the former dynast but not his accumulated knowledge about the economy. Instead, successors embark on their own learning experience "from scratch", starting with the identical initial belief that their predecessors had at birth. Formally, it means that in order to compute the expected value of $\boldsymbol{\gamma}$ and Σ^{-1} a dynast of age j only employs the information subset $\omega_j^t = \{X_\tau\}_{\tau=t-j+1}^t$.

Given these assumptions, we first demonstrate that, for each individual, the expected value of $\boldsymbol{\gamma}$ is computed according to a decreasing-gain recursive least squares (RLS) scheme:

Proposition 1 (Individual learning rule) *The value of* $\boldsymbol{\gamma}_{j,t} \equiv E\left[\boldsymbol{\gamma} \mid \omega_j^t\right]$ *is given by the RLS algorithm*

$$\boldsymbol{\gamma}_{j,t} = \boldsymbol{\gamma}_{j-1,t-1} + \frac{1}{n_0 + j}(X_t - \boldsymbol{\gamma}_{j-1,t-1}), \tag{3}$$

where $\boldsymbol{\gamma}_{0,t} = \boldsymbol{\gamma}_0$.

Proof. See Appendix C. ∎

This proposition provides a Bayesian foundation for the traditional decreasing-gain RLS learning (Evans and Honkapohja, 2003) where each agent is essentially behaving as an econometrician trying to infer the value of $\boldsymbol{\gamma}$ by using the available information. In the case of learning-from-experience, this available information would be heterogeneous across agents as it will only include a subset of all observations (namely, only those realized during the agent's lifetime). In contrast, in the limiting case in which all agents are infinitely-lived ($\theta = 1$), this would correspond to the standard case of a representative agent using all available history of observations.

Now we turn our attention to the average belief. It is defined as $\boldsymbol{\gamma}_t$:

$$\boldsymbol{\gamma}_t \equiv \frac{1}{N}\sum_{j=0}^{\infty} f_j \boldsymbol{\gamma}_{j,t} = (1-\theta)\boldsymbol{\gamma}_0 + \frac{1}{N}\sum_{j=1}^{\infty} f_j \boldsymbol{\gamma}_{j,t} = (1-\theta)\boldsymbol{\gamma}_0 + E_g\left[\boldsymbol{\gamma}_{n,t}\right], \tag{4}$$

where $E_g[\cdot]$ denotes the expectation under the probability distribution function $g(n) \equiv (1-\theta)\theta^{n-1}$, $n \in \mathbb{N}^+$ where $g(n) = \Pr[\text{age} = n \,|\text{age} > 0]$.

We introduce the following condition:

Condition 2 *The model satisfies*

$$E_g\left[\frac{1}{(n_0+n)}\boldsymbol{\gamma}_{n-1,t}\right] \approx E_g\left[\frac{1}{(n_0+n)}\right] E_g\left[\boldsymbol{\gamma}_{n-1,t}\right], \quad \forall t.$$

That is, the model should satisfy independence of $\boldsymbol{\gamma}_{n-1,t}$ across age cohorts. This condition should be checked in each specific application as there is no ex-ante guarantee that it will be satisfied. However, if the condition is satisfied, then the average belief is given by the following proposition

Proposition 3 (Average belief) *Provided that condition 2 is satisfied, the average belief $\boldsymbol{\gamma}$ is given by*

$$\boldsymbol{\gamma}_t \approx (1-\theta)\boldsymbol{\gamma}_0 + \boldsymbol{\gamma}_{t-1} + \Phi(\theta, n_0)(X_t - \boldsymbol{\gamma}_{t-1}), \tag{5}$$

where $\Phi(\theta, n_0) \equiv (1-\theta)\theta^{-n_0-1}\left[-\sum_{j=1}^{n_0}\frac{\theta^j}{j} - \log(1-\theta)\right]$.

Proof. See Appendix C. ■

This proposition shows that, given condition 2, the average belief in a heterogeneous agents OLG model can be approximated by a CGL model. The result is quite intuitive. In a decreasing-gain algorithm all observations have the same weight, whereas in CGL recent observations are weighted more than older ones, that is, there is a tendency to "forget" past information at a constant rate given by the gain parameter. The proposition implies that this is approximately what occurs in a model with Bayesian learners who have finite lives and learn from their own experience.

Finally, in the case of θ close to, but smaller than 1, we have that

$$\boldsymbol{\gamma}_t \approx \boldsymbol{\gamma}_{t-1} + \Phi(\theta, n_0)(X_t - \boldsymbol{\gamma}_{t-1}). \tag{6}$$

Given this finding, in the next section we illustrate how good an approximation (6) is in the context of a simple asset-pricing model.

3 Example: an asset-pricing model

In this section we show the applicability of approximation (6) in a simple asset pricing framework. To that end, we recast the model of Adam and Marcet (2011) in an OLG setup. We make two additional changes to their model as follows. First, we assume Bayesian learning of the means *and* the variances of the stock price and dividends.[4] And second, 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."[5]

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. As discussed above, 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, 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

[4] Adam 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).

[5] A "projection facility" is a technical assumption that mechanically constrains beliefs to a pre-specified neighborhood.

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.

3.1 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 $\beta\theta$, 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

$$E_{i0} \sum_{t=0}^{\infty} (\beta\theta)^t C_{it}, \tag{7}$$

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

Individual i faces the period budget constraint

$$C_{it} + P_t S_{it} \leq (P_t + D_t) S_{it-1} + Y_{it}, \tag{8}$$

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.

$$\underline{E}_t \leq P_t S_{it} \leq \bar{E}_t. \tag{9}$$

Constraints (9) 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 (9) 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.[6] 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

$$\log\left(D_t/D_{t-1}\right) = \varepsilon_t \sim N(\gamma, \sigma^2), \tag{10}$$

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 (7), subject to the budget constraint (8), and the exposure constraints (9).

3.2 Learning from experience

Individuals in this model 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 (10). 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 are assumed to "learn from experience," meaning that 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,

$$\omega_{i,n}^t = \{P_\tau, D_\tau\}_{\tau=t-n+1}^t .$$

Individuals 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 PLM is

$$\begin{bmatrix} \log\left(P_t/P_{t-1}\right) \\ \log\left(D_t/D_{t-1}\right) \end{bmatrix} = \begin{bmatrix} \varepsilon_{it}^P \\ \varepsilon_{it}^D \end{bmatrix} \sim N(\boldsymbol{\gamma}_i, \Sigma_i), \quad \boldsymbol{\gamma}_i = \begin{bmatrix} \gamma_i^P \\ \gamma_i^D \end{bmatrix}, \quad \Sigma_i = \begin{bmatrix} \sigma_{iP}^2 & \sigma_{iPD}^2 \\ \sigma_{iDP}^2 & \sigma_{iD}^2 \end{bmatrix}. \tag{11}$$

This specification allows for beliefs about the growth rates in the share price and dividends

[6] To see this fact, note that the budget constraint (8) 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.

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,

$$\Sigma_i^{-1} \sim W\left(\mathbf{\Sigma}_0, n_{i0}\right) \quad \text{and} \quad \boldsymbol{\gamma}_i | \Sigma_i^{-1} \sim N\left(\boldsymbol{\gamma}_{i0}', \Sigma_i / n_{i0}\right), \tag{12}$$

where 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,[7]

$$\boldsymbol{\gamma}_0' = (\gamma, \gamma), \quad \mathbf{\Sigma}_0 = \sigma^2 \begin{bmatrix} 1 & \delta \\ \delta & 1 \end{bmatrix} (n_0 - 3), \quad \text{where } \delta \to 1,\ \delta < 1. \tag{13}$$

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

$$\mathbf{e}_{it} = \begin{bmatrix} \log\left(P_t/P_{t-1}\right) - \gamma_{it}^P \\ \log\left(D_t/D_{t-1}\right) - \gamma_{it}^D \end{bmatrix}, \tag{14}$$

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

$$\boldsymbol{\gamma}_{it+1} = \boldsymbol{\gamma}_{it} + \frac{\mathbf{e}_{it}}{n_{it}+1}, \quad \mathbf{\Sigma}_{it+1} = \mathbf{\Sigma}_{it} + \frac{n_{it}}{n_{it}+1}\mathbf{e}_{it}\mathbf{e}_{it}', \quad n_{it+1} = n_{it} + 1.^8 \tag{15}$$

3.3 Equilibrium

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). The space of realizations is

$$\Omega \equiv \Omega_P \times \Omega_D, \tag{16}$$

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. Denote by Ω^t the set of histories up to period t, and let $\omega^t \in \Omega^t$. When investor i chooses his stock holding in period t, he takes as given Π_i and his choice is contingent on ω^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

[7]See Appendix A.

[8]The proof is the same as that of Proposition 1, but in this case it also includes the evolution of the covariance matrix

joint probabilities for all possible dividend and stock price realizations on any set of dates.

The definition of equilibrium in this model is:

Definition 4 (Equilibrium) *An internally rational expectations equilibrium (IREE) consists of a sequence of equilibrium price functions* $\{P_t\}_{t=0}^{\infty}$ *where* $P_t : \Omega_D^t \to R_+$ *for each* t*, contingent choices* $\{C_{it}, S_{it}\}_{t=0}^{\infty}$ *where* $(C_{it}, S_{it}) : \Omega^t \to R^2$*, 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 (7) subject to the budget constraint (8), and the exposure constraints (9), taking as given the probability measure* Π_i*.*

2. *When agents evaluate* $\{C_{it}, S_{it}\}$ *at equilibrium prices, markets clear:*

$$N = \sum_{j=0}^{\infty} f_j S_{jt}. \tag{17}$$

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

$$\begin{aligned} \text{if } P_t &< P_{it}, \text{ then } S_{it} = \bar{E}_t / P_t & (18a) \\ \text{if } P_t &= P_{it}, \text{ then } S_{it} \in [\underline{E}_t / P_t, \bar{E}_t / P_t], & (18b) \\ \text{if } P_t &> P_{it}, \text{ then } S_{it} = \underline{E}_t / P_t, & (18c) \end{aligned}$$

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

$$P_{it} = \beta\theta E_{it} \left(P_{t+1} + D_{t+1}\right) \tag{19}$$

is individual i's "reservation price". Because the objective function is linear and the feasible set is closed, a maximum exists.

The FOC can also be written as

$$P_t = \beta\theta E_{it} \left(P_{t+1} + D_{t+1}\right) + \mu_{it} = \beta\theta E \left[P_{t+1} + D_{t+1} \mid \omega_i^t\right] + \mu_{it}, \tag{20}$$

where $\mu_{it} \in \mathbb{R}$ is the sum of the Lagrange multipliers associated with the exposure constraints (9).

Given that for a normally distributed random variable $E \exp(\varepsilon) = \exp\left[E(\varepsilon) + Var(\varepsilon)/2\right]$, we have:

$$\begin{aligned} P_{it} &= \beta\theta E \left[P_{t+1} + D_{t+1} \mid \omega_i^t\right] = \beta\theta E \left[\exp\left(\varepsilon_{it+1}^P\right) P_t + \exp\left(\varepsilon_{it+1}^D\right) D_t \mid \omega_i^t\right] \\ &= \beta\theta \left\{\exp\left[\gamma_{it}^P + \mathbf{\Sigma}_{it}(1,1)/(2n_{it} - 6)\right] P_t + \exp\left[\gamma_{it}^D + \mathbf{\Sigma}_{it}(2,2)/\left(2n_{it} - 6\right)\right] D_t\right\}. \end{aligned}$$

The model is completely characterized by the first-order conditions for individual investors (20), the recursive Bayesian learning scheme (15), the market-clearing condition (17), and the exogenous process for dividends (10). A numerical solution algorithm to solve this model is presented in Appendix B.

4 Numerical results

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. Then we evaluate the goodness of the approximation of the average belief across agents by a CGL scheme. Finally, we analyze the limiting case in which the probability of survival is $\theta = 1$.

4.1 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 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 convergence rounds per period to $K = 2000$.[9] We perform 1000 Monte Carlo simulations of 5000 months each, equivalent to more than four centuries of trading.

[9] See Appendix B. We also report results with $N = 1000$ agents, which are very similar.

4.2 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 (9) 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.[10] 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. Note 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

[10]Note 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.

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.

Table 1 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, explaining it as a consequence of the dynamic coordination of heterogeneous beliefs.

4.3 Approximate average beliefs

In this subsection we show how well the average belief can be approximated by a constant-gain learning scheme. To this end, we take the actual price and dividend sequences from the benchmark heterogeneous-agents OLG model (HA-OLG) and construct series for stock price and dividend growth expectations using the representative-agent CGL approximation given by equation (6) (RA-CGL). To evaluate the accuracy of the approximation, we consider two metrics of fit: the correlation between the HA-OLG and RA-CGL expectations and the R^2, defined as one minus the ratio of the variance of the approximation error to the variance of the HA-OLG expectations.[11] Table 2 reports the two metrics. Correlation and R^2 are higher than 0.97 and 0.88, respectively. By these measures, the approximations of both the stock price and the dividend learning dynamics are very accurate. These results are robust to the number of agents and to alternative calibrations (not shown in the table). Figure 3 confirms the goodness of fit visually.

Note that the value of the gain parameter $\Phi(\theta, n_0)$, which appears in the approximation, is a function of the survival rate θ and of the prior precision n_0. In our baseline calibration, Φ is equal to 0.006, corresponding to an expected life on the market of 20 years and to a prior of 48 months. In quarterly terms, the retirement probability is 0.018, 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.

Finally, Table 3 compares the one-step-ahead forecast errors (14) generated by the HA-OLG model, the RA-CGL approximation, and the rational expectations model (REM). The distribution of forecast errors is quite similar between the HA-OLG and the RA-CGL models.[12] 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.

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

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

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, 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.

4.4 The case with infinitely lived individuals

Finally, 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.

The proof of convergence consists of two steps.[13] 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 $t \to \infty$ to establish the asymptotic convergence. The two steps are summarized by the following proposition and corollary, the proofs of which are in Appendix C.

Proposition 5 (Price dynamics with infinitely lived individuals) *If $\theta = 1$, the equilibrium price P_t is given by*

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

where

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

Proof. See Appendix C. ■

Corollary 6 (Convergence to REE) *The economy converges to the REE with equilibrium price $P_t^{REE} \equiv \frac{\beta\theta \exp\left(\gamma+\sigma^2/2\right)}{1-\beta\theta\exp(\gamma+\sigma^2/2)} D_t$.*

Proof. See Appendix C. ■

Having established asymptotic convergence, it is useful to know how long it takes for the stock price to converge to the REE.[14] Figure 4 plots one randomly drawn path of the ratio of the stock

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

[14] In 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.[15] 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 a 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 (9) 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 5, 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 5 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.[16]

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,

[15] In 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.

[16] The 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 analyze the effects of "learning from experience" in a stochastic OLG setup. We show that the average beliefs in this framework 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 gain parameter is a function of 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 down the gain parameter.

We illustrate the effect of learning from experience in an asset-pricing model. 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. Despite the apparently irrational waves of optimism and pessimism, we show that agents forecasting errors are unbiased, similar to the case of REE.

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, casting some doubts on whether the REE can be attained in reality.

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

[12] 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.

[13] 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.

[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] 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.

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

[22] Lucas Jr., R. E. 1972. "Expectations and the Neutrality of Money," *Journal of Economic Theory*, vol. 4 (2), pp. 103–24.

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

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

[25] 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.

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

[27] 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.

[28] Radner, R. 1979. "Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices," *Econometrica* 47 (3), pp. 655-678.

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

[30] Testfatsion, L. and K. Judd (May 2006). *Handbook of Computational Economics* 2. Elsevier.

[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: Symmetric rational expectations equilibrium

If individuals are identical, and this fact is common knowledge, they can compute the equilibrium asset price by deduction. Namely, dividing (19) 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 (10), yields:

$$\frac{P_t}{D_t} = \beta\theta E_t \left[\frac{D_{t+1}}{D_t} \left(1 + \frac{P_{t+1}}{D_{t+1}} \right) \right] = \sum_{j=1}^{\infty} (\beta\theta)^j \, e^{j\left(\gamma+\sigma^2/2\right)} + \lim_{T\to\infty} E_t \left((\beta\theta)^T \, \frac{D_{t+T}}{D_{t+T-1}} \frac{P_{t+T}}{D_{t+T}} \right). \quad (23)$$

Given that the sum of stock holdings must equal the fixed supply of the stock N, it follows from (9) that the price-dividend ratio is bounded above by λ,

$$N = \sum_{i=1}^{N} S_{it} \leq N\lambda D_t/P_t \implies P_t/D_t \leq \lambda. \quad (24)$$

Hence the last term in (23) is zero, and therefore the equilibrium asset price is given by

$$P_t^{REE} = \frac{\beta\theta \exp\left(\gamma + \sigma^2/2\right)}{1 - \beta\theta \exp\left(\gamma + \sigma^2/2\right)} D_t > 0, \quad (25)$$

where dividends follow the exogenous stochastic process defined in (10). We further impose the parameter restrictions,

$$\beta\theta \exp\left(\gamma + \sigma^2/2\right) < 1 \text{ and } \lambda > \frac{\beta\theta \exp\left(\gamma + \sigma^2/2\right)}{1 - \beta\theta \exp\left(\gamma + \sigma^2/2\right)}, \quad (26)$$

which ensure that the price-dividend ratio is finite and that it is not a corner solution due to binding leverage constraints (24).

Appendix B: Simulation algorithm

We briefly sketch the algorithm used to find the equilibrium price of the agent-based heterogeneous-beliefs model. 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 (17 - 20), we propose a recursive procedure of K rounds aimed at finding a price which is consistent with agents' beliefs and constraints and which guarantees that the market clears. Proposition 5 and Corollary 6 demonstrate that this algorithm converges to the REE in the case of infinitely-lived individuals.

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

1. Generate an exogenous series for dividends D_t following (10) 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 (25).

2. Initialize the prior beliefs, $\boldsymbol{\gamma}_{i0}$, $\boldsymbol{\Sigma}_{i0}$, and n_{i0}, for all agents following (12) and (13).

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

   (a) Compute the one-step-ahead forecast errors $\mathbf{e}_{it}$ using (14)

   (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, $\boldsymbol{\gamma}_{it}$, $\boldsymbol{\Sigma}_{it}$, and n_{it}, using (15). If he retires (he is replaced by a new agent), set $\boldsymbol{\gamma}_{it} = \boldsymbol{\gamma}_{i0}$, $\boldsymbol{\Sigma}_{it} = \boldsymbol{\Sigma}_{i0}$, and $n_{it} = n_{i0}$.

   (d) Set the initial equilibrium price to $P_{t0} = P_{t-1}$.

   (e) Compute the reservation price for each agent in round zero P_{it0} using

   $$P_{it0} = \beta\theta \left\{ \exp\left[2\gamma_{it}^P + \boldsymbol{\Sigma}_{it}(1,1)/\left(n_{it}-3\right)\right] P_{t-1} + \exp\left[2\gamma_{it}^D + \boldsymbol{\Sigma}_{it}(2,2)/\left(n_{it}-3\right)\right] D_{t-1} \right\}.$$

   (f) **Equilibrium price computation**. For each round $k = 1, .., K$:

      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 receiving 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 by P_{itk-1}^*.

      ii. If $\sum_{i=1}^{N} S_{it} \leq N$, then set $P_{tk} = P_{it-1}$.[17] Otherwise, set $P_{tk} = P_{itk-1}^*$.

[17] If several investors have the same reservation price, they receive an equal share of the stock. If there is

iii. The reservation price of each agent in round $k > 1$, P_{itk}, is computed using $P_{itk} = \beta\theta \left\{ \exp\left[\gamma_{it}^{P} + \mathbf{\Sigma}_{it}(1,1)/(2n_{it} - 6)\right] P_{tk} + \exp\left[\gamma_{it}^{D} + \mathbf{\Sigma}_{it}(2,2)/\left(2n_{it} - 6\right)\right] D_t \right\}$.

(g) The price computation is over in round K, and the stock price in period t is $P_t = P_{tK}$.

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

insufficient demand to clear the market because leverage constraints are binding for everyone, then the market closes recording $P_t = P_{t-1}$. This possibility did not arise in any of the simulations, but we consider it here for the sake of completeness.

Appendix C: Proofs

Proposition 1.
Proof. The PLM is

$$X_t = \boldsymbol{\gamma} + \varepsilon_t.$$

The joint prior distribution of $\boldsymbol{\gamma}$ and Σ is a *normal-Wishart*:

$$\Sigma^{-1} \sim W\left(\boldsymbol{\Sigma}_0, n_0\right) \quad \text{and} \quad \boldsymbol{\gamma}|\Sigma^{-1} \sim N\left(\boldsymbol{\gamma}_0, \Sigma/n_0\right), \tag{27}$$

Following DeGroot (1970),the value of $\boldsymbol{\gamma}_{j,t} \equiv E\left[\boldsymbol{\gamma} \mid \omega_j^t\right] = E\left[\boldsymbol{\gamma} \mid \{X_\tau\}_{\tau=t-j+1}^t\right]$ is:

$$\boldsymbol{\gamma}_{j,t} = \frac{n_0\boldsymbol{\gamma}_0 + j\bar{X}_t^j}{n_0 + j},$$

where $\bar{X}_t^j = \frac{1}{j}\sum_{\tau=t-j+1}^{t} X_\tau$ is the sample mean vector. Using that $\bar{X}_t^j = \frac{1}{j}\left[X_t + (j-1)\bar{X}_{t-1}^j\right]$ and substituting above yields

$$\boldsymbol{\gamma}_{j,t} = \boldsymbol{\gamma}_{j-1,t-1} + \frac{1}{n_0 + j}\left(X_t - \boldsymbol{\gamma}_{j-1,t-1}\right),$$

where $\boldsymbol{\gamma}_{0,t} = \boldsymbol{\gamma}_0$. ■

Proposition 3.
Proof. First, we show that**,** given $\theta \in (0,1)$ and $n_0 \in \mathbb{N}^+$:

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

where

$$\int\frac{\theta^{n_0}}{(1-\theta)}d\theta = -\sum_{j=1}^{n_0}\frac{\theta^j}{j} - \log(1-\theta).$$

Then we proceed to compute $\boldsymbol{\gamma}_t$:

$$\boldsymbol{\gamma}_t = (1-\theta)\boldsymbol{\gamma}_0 + E_g\left[\boldsymbol{\gamma}_{n,t}\right] = (1-\theta)\boldsymbol{\gamma}_0 + E_g[\boldsymbol{\gamma}_{n-1,t-1}] + E_g\left[\frac{1}{(n_0+n)}(X_t - \boldsymbol{\gamma}_{n-1,t-1})\right],$$

where $E_g[\boldsymbol{\gamma}_{n-1,t-1}] = \sum_{j=1}^{\infty}(1-\theta)\theta^{j-1}\boldsymbol{\gamma}_{j-1,t-1} = \boldsymbol{\gamma}_{t-1}$ and by condition 2, $E_g\left[\frac{1}{(n_0+n)}\boldsymbol{\gamma}_{n-1,t}\right] \approx E_g\left[\frac{1}{(n_0+n)}\right] E_g\left[\boldsymbol{\gamma}_{n-1,t}\right]$

$$\begin{aligned}
\boldsymbol{\gamma}_t &= (1-\theta)\boldsymbol{\gamma}_0 + \boldsymbol{\gamma}_{t-1} + E_g\left[\frac{1}{n_0+n}\right](X_t - \boldsymbol{\gamma}_{t-1}) \\
&= (1-\theta)\boldsymbol{\gamma}_0 + \boldsymbol{\gamma}_{t-1} + \Phi(\theta, n_0)(X_t - \boldsymbol{\gamma}_{t-1})
\end{aligned}$$

■

Proposition 5.

Proof. Because individuals are identical we can drop index i. In the initial round of the equilibrium price computation in Appendix B at time t, the price is given by

$$P_{t0} = \beta\left(2\pi_t^P P_{t-1} + 2\pi_t^D D_{t-1}\right), \tag{28}$$

where (22) holds. In the subsequent rounds, the price evolves as

$$P_{tk} = \beta\left(\pi_t^P P_{tk-1} + \pi_t^D D_t\right) = \left(\beta\pi_t^P\right)^k P_{t0} + \beta\pi_t^D D_t \sum_{i=1}^{k-1}\left(\beta\pi_t^P\right)^i. \tag{29}$$

If $\beta\pi_t^P \geq 1$, then as $k \to \infty$, P_{tk} would grow unboundedly were it not for constraint (9) 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 $k \to \infty$, the first term in equation (29) tends to zero and the price for period t is

$$P_t \equiv \lim_{k\to\infty} P_{tk} = \frac{\beta\pi_t^D}{1-\beta\pi_t^P} D_t. \tag{30}$$

■

Corollary 6.

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

$$\lim_{t\to\infty}\gamma_t^D = \gamma, \quad \lim_{t\to\infty}\frac{\boldsymbol{\Sigma}_t(2,2)}{n_t-3} = \sigma^2, \quad \text{and} \quad \lim_{t\to\infty}\pi_t^D = \exp\left(\gamma + \sigma^2/2\right). \tag{31}$$

Second, given the equilibrium price (21), the value of $\log\left(P_{t-1}/P_{t-2}\right)$ is bounded as $t \to \infty$.

Therefore, given the Bayesian updating scheme, π_t^P must converge ,

$$\begin{aligned}
\lim_{t\to\infty} \pi_t^P &= \lim_{t\to\infty} \exp\left(\gamma_t^P + \frac{\mathbf{\Sigma}_t(1,1)}{2(n_t-3)}\right) \\
&= \lim_{t\to\infty} \exp\left[\gamma_{t-1}^P + \frac{\log\left(P_{t-1}/P_{t-2}\right) - \gamma_{t-1}^P}{1+n_{t-1}} + \frac{\mathbf{\Sigma}_{t-1}(1,1)}{2(1+n_{t-1}-3)} + \right. \\
&\left. + \frac{n_{t-1}\left(\log\left(P_{t-1}/P_{t-2}\right) - \gamma_{t-1}^P\right)^2}{2(1+n_{t-1}-3)\left(1+n_{t-1}\right)}\right] \\
&= \lim_{t\to\infty} \pi_{t-1}^P \equiv \pi^P. \qquad (32)
\end{aligned}$$

Third, the limit π^P must satisfy

$$\lim_{t\to\infty} \pi_t^P = \pi^P < 1/\beta \qquad (33)$$

This last point can be proved by contradiction: suppose $\pi^P \geq 1/\beta$. Then, all individual constraints (9) must be binding, so that (24) is binding as well, and

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

by (31) and (26); thus, we have reached a contradiction.

Finally, by taking the log-difference of (30),

$$\lim_{t\to\infty} \log\left(P_t/P_{t-1}\right) = \lim_{t\to\infty} \left\{\log\left[\frac{\pi_t^D\left(1-\beta\pi_{t-1}^P\right)}{\pi_{t-1}^D\left(1-\beta\pi_t^P\right)}\right] + \log\left(D_t/D_{t-1}\right)\right\}. \qquad (34)$$

Together, (31) and (33) imply that the first term in the brackets on the right-hand side of (34) converges to zero, and hence the learning parameters for the stock price must also converge to the asymptotic values of the REE,

$$\lim_{t\to\infty} \gamma_t^P = \gamma, \quad \lim_{t\to\infty} \frac{\mathbf{\Sigma}_t(1,1)}{(n_t-3)} = \sigma^2, \quad \text{and} \quad \lim_{t\to\infty} \pi_t^P = \exp\left(\gamma + \sigma^2/2\right). \qquad (35)$$

Substituting the above in equation (30) we obtain $\lim_{t\to\infty} P_t = \frac{\beta\exp\left(\gamma+\sigma^2/2\right)}{1-\beta\exp(\gamma+\sigma^2/2)} D_t = P_t^{REE}$. ■

Appendix D. Tables and figures

Table 1. Moments of the price-dividend ratio

	Data	REM	HA-OLG
Mean	320.3	307.6	316.8
Standard deviation	166.1	0	65.4
Autocorrelation	0.996	–	0.995

Note: REM stands for "rational expectations model"

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

Table 2. Accuracy of the RA-CGL approximation

	1000 agents		100 agents	
	R^2	Correl.	R^2	Correl.
Price learning	0.94	0.97	0.94	0.97
Dividend learning	0.88	0.99	0.88	0.99

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

Table 3. Moments of the forecast errors

	Mean	Std. Dev.	Skewness	Kurtosis
Price forecast errors				
REM	9.2×10^{-5}	0.0114	-0.0148	3.0766
HA-OLG	5.2×10^{-5}	0.0207	0.1346	3.4809
RA-CGL	2.1×10^{-5}	0.0208	0.1327	3.4680
Dividend forecast errors				
REM	9.2×10^{-5}	0.0114	-0.0148	3.0766
HA-OLG	3.4×10^{-5}	0.0114	-0.0149	3.0605
RA-CGL	1.7×10^{-5}	0.0114	-0.0153	3.0576

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."

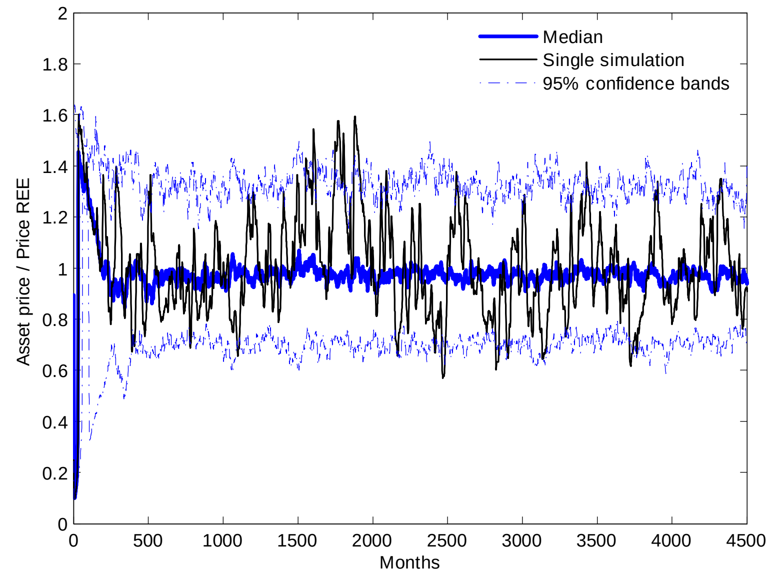

Figure 1: Stock price divided by the rational expectations price

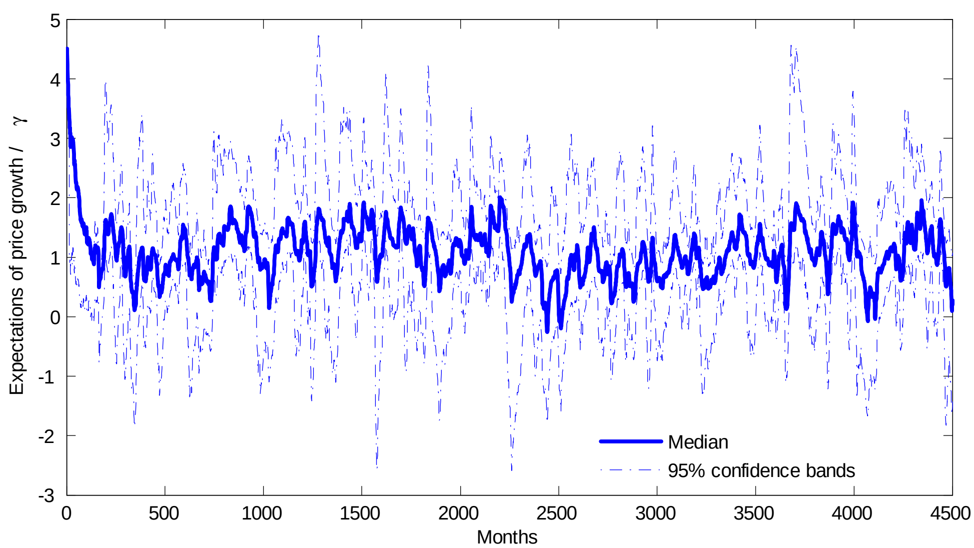

Figure 2: Expectations of stock price growth relative to REE

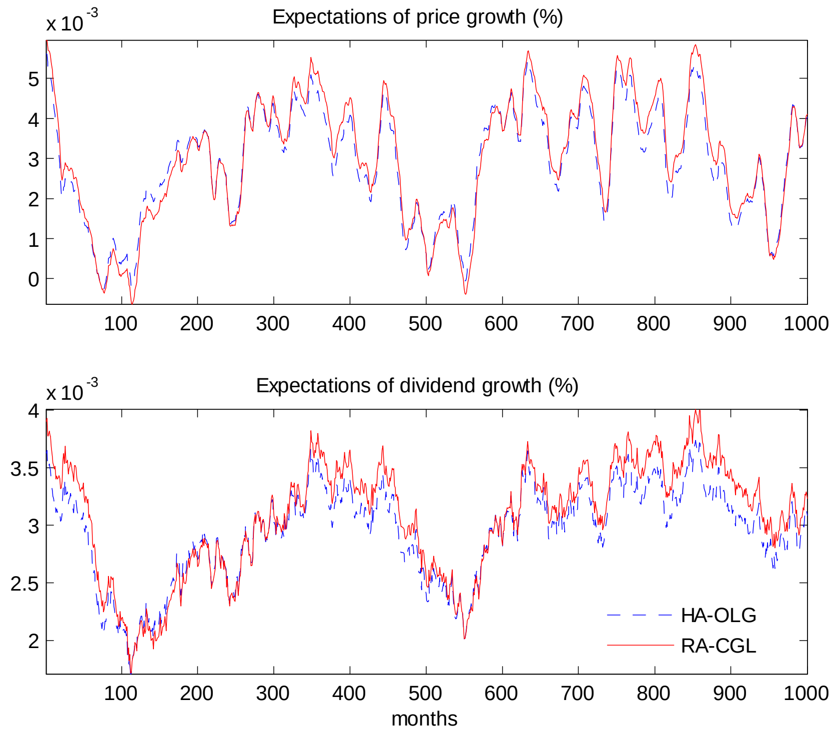

Figure 3: Average of heterogeneous beliefs (HA-OLG) versus representative-agent CGL beliefs (RA-CGL)

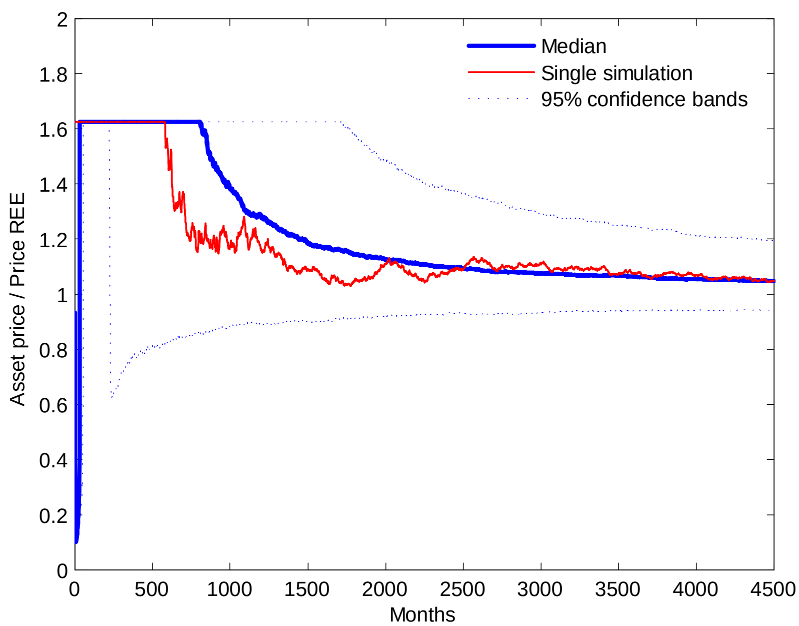

Figure 4: Convergence to rational expectations with infinitely-lived agents

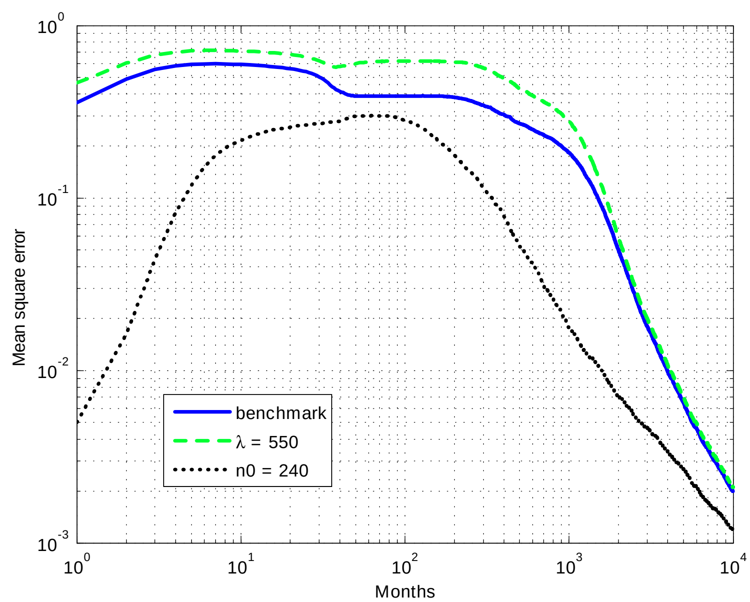

Figure 5: Robustness of the convergence to changes in leverage and in prior precision