DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

BAYESIAN ECONOMISTS...BAYESIAN AGENTS II: EVOLUTION OF BELIEFS IN THE SINGLE SECTOR GROWTH MODEL

Mahmoud A. El-Gamal
California Institute of Technology

Rangarajan K. Sundaram
University of Rochester

SOCIAL SCIENCE WORKING PAPER 736

May 1990

Bayesian Economists ... Bayesian Agents II: Evolution of Beliefs in the Single Sector Growth Model

Mahmoud A. El-Gamal
California Institute of Technology

Rangarajan K. Sundaram
University of Rochester

May 28, 1990

Abstract

In "Bayesian Economists ... Bayesian Agents I" (BBI), we generalized the results on Bayesian learning based on the martingale convergence theorem from the repeated to the sequential framework. In BBI, we showed that the variability introduced by the sequential framework is sufficient under very mild identifiability conditions to circumvent the incomplete learning results that characterize the literature. In this paper, we demonstrate that result in the neo-classical single sector growth model under even weaker identifiability conditions. We study the evolution of agent-beliefs in that model and show that, under reasonable conditions, the dependence of the current capital stock on the previous capital stock induces enough variability for our complete learning results to become relevant. Not only does complete learning take place from the subjective point of view of the agents' priors, but it also takes place from the point of view of an objective observer (modeling economist) who knows the true structure.

Please Send Communications to DHSS 228-77, Caltech, Pasadena, CA 91125.

1 Introduction

The cornerstone of much of contemporary macroeconomics is a behavioral assumption termed the Rational Expectations Hypothesis (REH). Proponents of the REH claim that the hypothesis is a little more than a formalization of the statement that rational economic agents do not commit systematic errors, and use all the information available to them optimally in arriving at decisions. In practice, however, it has become customary to invoke the REH to justify a considerably stronger behavioral assumption: that agents in the economy are aware of, and take expectations with respect to, the "true" distribution of the stochastic process governing their environment. The question naturally arises whether these assumptions are equivalent, at least in a long run sense. In other words, in the context of maximizing agents who are at the outset unaware of the true parameters of their environment but learn optimally (in a Bayesian sense), will the agents fully learn the values of the unknown parameters? This question has been the focus of intense investigations by economic theorists in recent years.

The standard approach in the optimal learning literature can perhaps be best exemplified by the recent paper of Easley and Kiefer (1988). They study a single infinitely lived agent who begins with a prior belief on the set of possible parameter values governing the environment. It is seen that the agent's optimization/learning problem can be reformulated as a dynamic programming problem with beliefs as the state variable. Since the agent cannot expect his beliefs to change in any systematic way, it is then possible to show that the stochastic process of his beliefs, under any solution to the dynamic programming problem, is a martingale. An appeal to the martingale convergence theorem now reveals the existence of limit beliefs to which the sequence of beliefs converges almost surely (with respect to the probability measure induced on the space of sample paths by the agent's initial prior). A confirmation of the REH from the point of view of that literature would now require that that limit be (almost surely) degenerate at the true parameter values. For a detailed discussion of this approach, and the results so far available, we refer the reader to Easley and Kiefer (1988). Here we will rest with the observation that in many examples, full learning does not occur in the limit with positive probability, i.e. the two versions of the REH are not reconciled.

In El-Gamal and Sundaram (1989) [hereafter E-G/S] we question this standard approach on two grounds. Firstly, we argue that the non-learning results that characterize the optimal learning literature are largely the artifact of the use of a repeated statistical decision framework, with beliefs forming the only link between time periods. We demonstrate that in the (economically more realistic context of) presence of a physical link between periods, very mild identifiability conditions suffice to ensure full learning in the limit. Secondly, and more importantly, however, we question the use of the martingale convergence theorem to demonstrate the existence of limit beliefs for the agent, and argue that the evolution of beliefs should be studied from the point of view of an external agent who is informed of the true parameter values. The difference is non-trivial as the examples in Freedman (1963,1965) or Feldman (1989) illustrate. We refer the reader to E-G/S for details, and a justification, of the alternative approach we propose for a theoretical test of the REH in any given model. For purposes of continuity, however, we include a very brief description of our approach here.

Primarily, our approach rests on the notion that economists cannot have, or impose, in their models any restrictions on the prior beliefs of agents. For a confirmation of the REH, we would therefore require that agents who are identical except for their initial beliefs should have the same limit beliefs (which should moreover put full mass at the true parameter value) almost surely with respect to the *true* stochastic process driving their beliefs. We take as our staring point, therefore, a Bayesian economist who is modeling the economic problem in question, and hence knows the true parameter value. Our economist, however, knows that she has to allow for different possible priors of her modelled agents, and she starts with a prior on possible agent priors. The economist prior is assumed to not rule out any agent-priors as long as they have the true parameter in their support, and thus places full support on such agent-priors. No other assumptions are placed on the economist prior, allowing for the distribution of mass on agent-priors to reflect any model-specific information that the economist may have. Knowledge of the true parameter and the agents' decision rules then allows the economist to update her beliefs on agent-priors. This in turn gives rise to a Markov process describing the evolution of economist beliefs.

The conditions required for a confirmation of the REH are now easily stated

1. The economist beliefs converge to a limit belief.

2. The economist's limit belief should be degenerate.

3. The unique agent-belief in the support of the limiting economist-belief should itself be degenerate at the true parameter value.

This clearly shows that the REH is a very special case of a much richer class of models available to the economist. For instance, if 1 and 2 hold but 3 does not, then the modeling economist should use the unique agent-belief in the support of her limit belief instead of using a RE model. Similarly, if 1 holds but 2 does not, then the economist should study the dynamics of her model under all the agent-beliefs in the support of her limit-belief, and then weigh the results of the models accordingly. It is beyond the scope of this paper, however, to demonstrate the more general modeling techniques that our framework suggests.

In this paper, we illustrate the approach of E-G/S using numerical analysis of a simple sequential statistical decision framework. The model we use is the familiar neo-clasical single sector growth model under production uncertainty. The twist we add to this canonical model of dynamic economic theory is that the agent solving this optimization problem (the metaphorical 'planner' of economic theory) is unaware of the true law of motion for the conversion of today's investment into tomorrow's output. The planner begins instead with a prior belief on the set of possible production laws of motion, and uses observations to date to update her beliefs in a Bayesian manner. As E-G/S show, using a modification of the techniques in Easley and Kiefer (1988), the planner's problem has an optimal investment policy as a function of her current belief and current capital stock. Section 2 will describe the model in detail. Section 3 will then describe the numerical techniques that we used to analyze the model. The results of the numerical analysis described in section 3 are shown in Figures 1 through 10 and they clearly show that the REH is justified in the framework of the single sector growth model on the basis that the probability of an agent (the proportion of agents) being outside a small neighborhood of the true belief, even though it never becomes zero as shown theoretically in E-G/S, tends to zero as time goes to infinity.

2 The Framework

2.1 Notation and Definitions

N.1 For any space X, $P(X)$ denotes the set of all probability measures on X. Convergence in $P(X)$ is always with respect to the weak-* topology.

N.2 Given X, $\delta_x \in P(X)$ will denote the probability measure that puts full mass at the point $x \in X$.

2.2 The model

There is a single good which may be consumed or invested. In each period t of an infinite horizon, a *planner* observes the available stock $y_t \geq 0$ of the good and decides on the allocation of y_t between consumption $c_t \geq 0$ and investment $x_t \geq 0$. Consumption of c units yields instantaneous utility of $u(c) = \sqrt{c}$. Conversion of investment to output takes one period, and is the outcome of a stochastic process. Given today's investment of x units, tomorrow's output is randomly determined by the conditional density $q(.|x)$. The planner, however, does not know the form of $q(.|.)$ but knows only that $q(.|.) \in \{q_1(.|.), q_2(.|.)\}$ where

$$q_1(y|x) = \begin{cases} \frac{4}{x}(y - \sqrt{x}) & \text{if } y \in [\sqrt{x}, \frac{3}{2}\sqrt{x}] \\ \frac{4}{x}(2\sqrt{x} - y) & \text{if } y \in [\frac{3}{2}\sqrt{x}, 2\sqrt{x}] \\ 0 & \text{otherwise} \end{cases}$$

and

$$q_2(y|x) = \begin{cases} \frac{1}{\sqrt{x}} & \text{if } y \in [\sqrt{x}, 2\sqrt{x}] \\ 0 & \text{otherwise} \end{cases}$$

The two conditional densities are seen pictorially in Figures 1 and 2 Observe that for all $x \in \Re_+$, $supp.q_1(.|x) = supp.q_2(.|x) = [\sqrt{x}, 2\sqrt{x}]$, and that both q_1 and q_2 have the same expected output, i.e. $\int yq_1(y|x\ dy = \int yq_2(y|x)\ dy = \frac{3}{2}\sqrt{x}$. Note also that if $x = 0$, then tomorrow's stock is 0 with probability 1 (no free production), and that for $x \geq 4$, $supp.q_i(.|x) \subset [0, x], i = 1, 2$. By the last observation, there is no loss of generality in restricting attention to initial stocks in $Y = [0, 4]$.

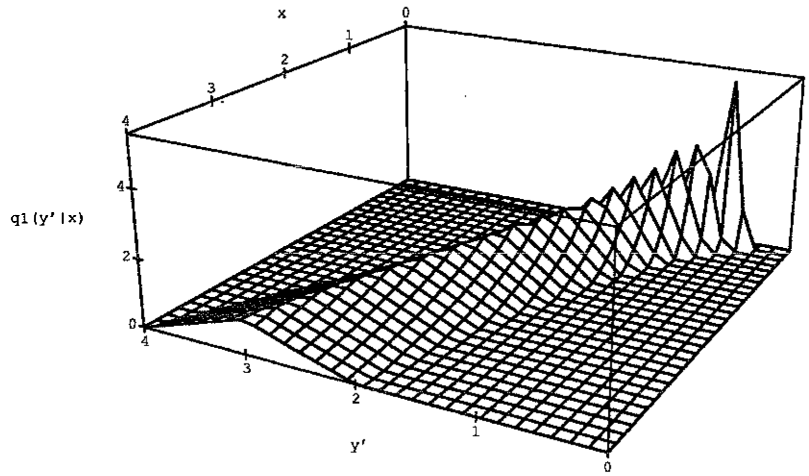

Figure 1: The technology $q1(.|.)$

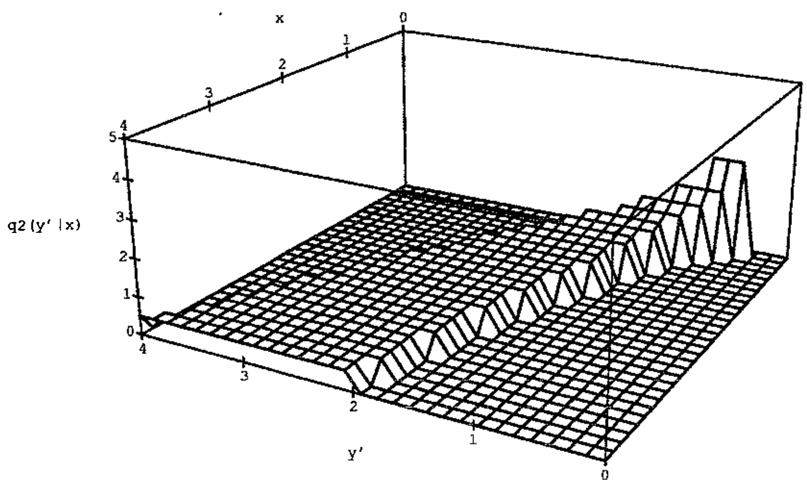

Figure 2: The technology $q2(.|.)$

The planner begins with a prior $p_0 \in [0,1]$, where the prior represents the planner's belief that $q = q_1$ (so $(1-p_0)$ is the belief that $q = q_2$). In each period t, the planner updates her beliefs, after observing the outcome y_{t+1} that resulted from the investment x_t of the prior period, in the usual Bayesian manner:

$$p_{t+1} = B(p_t, x_t, y_{t+1}) = \frac{p_t q_1(y_{t+1}|x_t)}{p_t q_1(y_{t+1}|x_t) + (1-p_t) q_2(y_{t+1}|x_t)} \tag{2.1}$$

The planner discounts the future utilites by $\beta \in (0,1)$ and wishes to maximize total expected discounted utility over the infinite horizon $(E_0 \sum_{t=0}^{\infty} \beta^t u(y_t - x_t))$ given her initial prior p_0 and the initial stock y_0.

As E-G/S demonstrate by extending the techniques of Easley and Kiefer (1988), the planner's decision problem may be reformulated as a standard dynamic programming problem with **state space** $S = Y \times [0,1]$, **action space** $X = Y$, **feasible action correspondence** $\Phi: S \to X$ defined by $\Phi(y,p) = [0,y]$, **payoff function** $r: S \times X \to \Re$ defined by $r(y,p,x) = u(y-x) = \sqrt{y-x}$, **discount factor** $\beta \in (0,1)$, and finally the **transition probabilities** $Q(.|s,x) \in P(S)$ for $(s,x) \in S \times X$, where for A a Borel subset of S,

$$Q(A|s,x) = Pr\{(y',p') \in A|(y,p,x)\} = Pr\{(y', B(p,x,y')) \in A|(y,p,x)\}$$

which is easily calculable using (2.1).

Using standard results, it can be shown, under regularity conditions that are clearly met, that this dynamic programming problem is well defined. The planner's value function $V: S \to \Re_+$ is a continuous function that satisfies at each $s = (y,p) \in S$ the Bellman optimality equation

$$V(s) = \max_{x \in [0,y]} \{\sqrt{y-x} + \beta \int_S V(s')\, Q(ds'|s,x)\} \tag{2.2}$$

Moreover, the correspondence of maximizers in (2.2) is an upper-hemicontinuous correspondence, and admits a measurable selection $g: S \to X$. The function g is a stationary optimal policy for the planner's problem.

We assume that all the agents use the same selection g, and that the economist knows this selection. The economist is further endowed with the knowledge that $q = q_2$. He starts with an initial belief $\mu_0 \in P([0,1])$, where $supp.\mu_0 = [0,1]$. Let A be a Borel subset of $[0,1]$, and $(y_t, x_t = g(y_t,p_t), y_{t+1})$ be the period T capital stock, the period t investment, and the period $t+1$ realization of stock, respectively, and let $\nu_0 \in P(Y)$ be the measure describing the initial distribution of the capital stock. The $\mu_0 \times \nu_0$ defines an initial measure on S which evolves as $\mu_t \times \nu_t$ according to the Markov process defined by the stochastic kernel Q. We observe the margin on beliefs, $\mu_t \in P([0,1])$ and it follows

$$\mu_{t+1}(A) = \int_{[0,1]} \int_Y I_A(y_{t+1}, x_t, y_t, p_t) q_2(y_{t+1}|x_t)\, \nu_t(dy_t)\, \mu_t(dp_t) \tag{2.3}$$

The Markov process $\{\mu_t \times \nu_t\}$, and especially its margin on belifs $\{\mu_t\}$ forms the object of investigation of the numerical analysis described in the next section. To recap the questions raised in the introduction, we are essentially interested in the following questions:

1. Does there exsist a limit economist-belief $\mu^\star$ such that $\mu_t \Rightarrow \mu^\star$ a.s.$[\mu_0]$ as $t \uparrow \infty$?
2. If yes, is $\mu^\star = \delta_{p^\star}$ for some $p^\star \in [0,1]$?
3. Finally, if the answer to 2 is also in the affirmative, is $p^\star = 0$?

It is worthwhile noting at this point that the model specified in this section, despite its apparent simplicity, is general enough to demonstrate the effect of the sequential structure on the answers to questions 1.–3. above. In BBI, we showed that in Kiefer's (1989) repeated framework (which also had only two demand curves corresponding to our two technologies in this model), full learning does not occur and the answers to 2. and 3. above were negative and the limit $\mu^\star$ in 1. depended on the initial prior on priors μ_0. In this paper, by introducing a sequential model, and one where we tried to favor non-learning, the answers to 1.–3. are all in the affirmative, as our theoretical and heuristic arguments in BBI anticipated.

3 The numerical analysis

We proceeded by first solving the dynamic optimization problem, getting numerical solutions for the value function $V(y,p)$ and the optimal investment rule $x = g(y,p)$. We then analyze ensembles of sample paths starting from different initial (y,p) running them through the true law of motion q_2 and then looking at the distribution of p_t's arising from that process. That latter distribution is clearly the sample analog of μ_t. The rest of this section will discuss the details of the procedures that we used to obtain the various components of our model.

3.1 The value function and the optimal investment function

To get a numerical solution for the value function, we proceed in the usual way to obtain a fixed point of equation (2.2), where now, the transition is defined by

$$Q(dy'|p,x) = [p.q_1(y'|x) + (1-p).q_2(y'|x)]dy'$$

the Bayes updating rule is as defined by equation (2.1), and the value function is defined by

$$V(y,p) = \max_{x\in[0,y]}\{\sqrt{y-x} + \beta \int_S V(y', B(p,x,y'))\ Q(dy'|p,x)\}$$

As usual, it is clear that the associated map $T: C(Y \times P) \to C(Y \times P)$ defined by

$$T[V(y,p)] = \max_{x\in[0,y]}\{\sqrt{y-x} + \beta \int_S V(y', B(p,x,y'))\ Q(dy'|p,x)\}$$

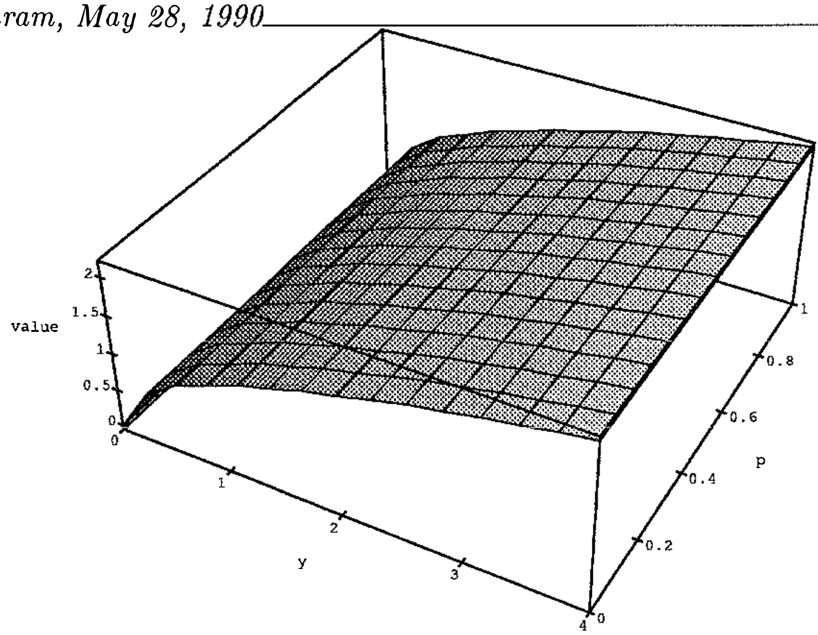

Figure 3: The value function – $\beta = 0.25$

is a contraction mapping. We use the one-shot maximization solution as our initial guess of the value function, i.e. $V_0(y,p) = \sqrt{y}$ and evaluate it on a 50×50 grid of points in $Y \times P = [0,4] \times [0,1]$. We then iterate on the map $V_t(y,p) = T[V_{t-1}(y,p)]$ on the grid, where we evaluate the value of each investment level x on a grid of 100 points on $[0,y]$ and choose the maximum such value. The integration for each iteration was carried out by using the trapezoidal rule on a grid of 100 points. The functions were sufficiently smooth that the experimental use of the significantly more computationally intensive procedure of taking multiple grids in the spots where the integrand varies significantly showed that there was not enough gain in accuracy to justify the increased CPU time. We continue this iteration procedure until the maximum difference over the $Y \times P$ grid between two consequetive value functions is smaller than 0.0005. At each iteration, we keep a matrix representing $V(y,p)$ over the grid, and another representing $x = g(y,p)$, the optimal investment level at each (y,p) corresponding to that value function. When we stop the iterative procedure, these are our value function and optimal investment function. Figures 3 and 4 show the plots of the value function and the optimal investment rule, respectively, for $\beta = 0.25$. Figures 5 and 6 show the same plots for $\beta = 0.75$. In the case of $\beta = 0.25$, it took 7 iterations to bring the maximum difference between two consequetive value functions to within 0.0005, and in the case of $\beta = 0.75$, it took 16 iterations to bring that difference to within 0.0025. Each iteration for the computation of the value function and the optimal investment function took approximately 5 CPU hours (for a total of 115 CPU hours for the computation shown in figures 3-6). All computations for this subsections were done using C code on a Sun4. The computation of the dynamics for the next sections, and all figures, were done using Mathematica on a Sun4.

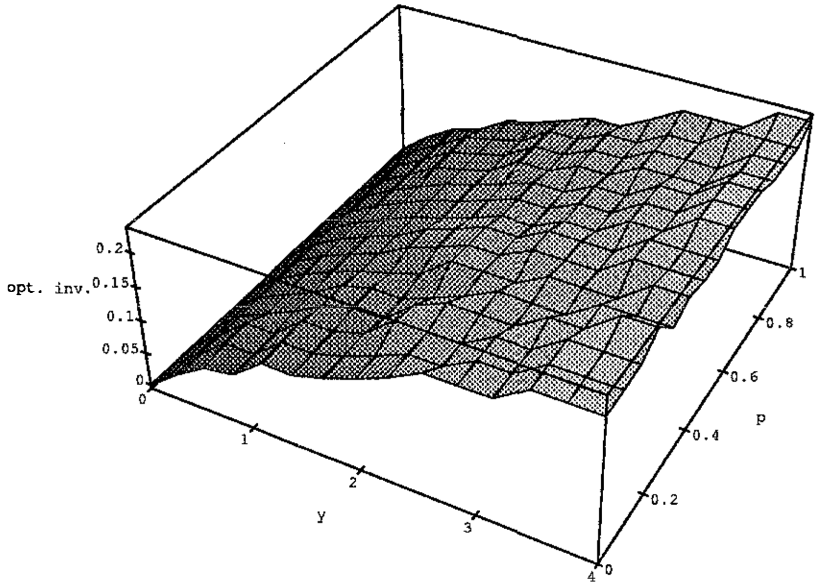

Figure 4: The optimum investment function – $\beta = 0.25$

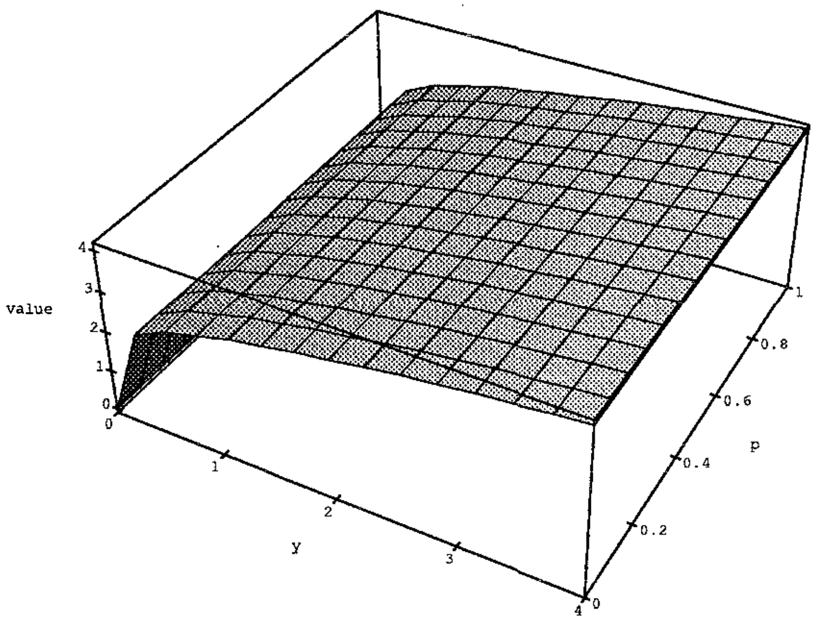

Figure 5: The value function – $\beta = 0.75$

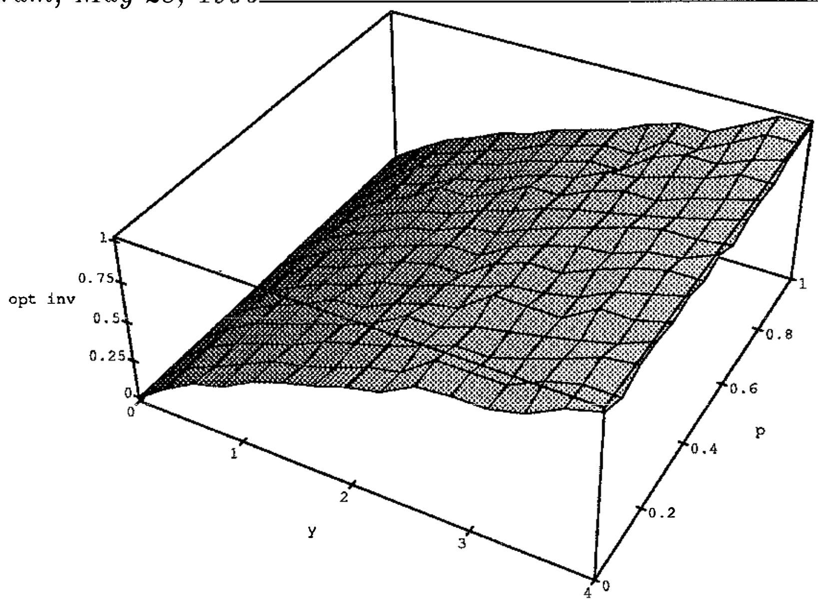

Figure 6: The optimum investment function – $\beta = 0.75$

3.2 Agent-belief, and economist-belief trajectories

Now, to study the dynamics of agent and economist-beliefs, we start by drawing 100 agents' initial capital stocks and priors at random from $Y \times P$. We chose to sample each agent i's initial condition uniformly, i.e. $(y_0^i, p_0^i) \sim (U[0,4], U[0,1])$, $i = 1, ..., 100$. For each of those agents, and for each time period $t = 0, 1, ...T$, we then go through a procedure of choosing the agent's optimal level of investment $x_t^i = g(y_t^i, p_t^i)$ by searching over the grid described in the previous subsection. We then use each agent's chosen level of investment x_t^i to draw her output according to $q_2(.|.)$, i.e. $y_{t+1}^i \sim U[\sqrt{x_t^i}, 2\sqrt{x_t^i}]$. We then use equation (2.1) to find p_{t+1} given the already availble values of p_t, x_t, y_{t+1}. Figures 7 and 9 show the 100 resulting trajectories of $\{p_t^i\}, i = 1, ..., 100$ from that procedure for $\beta = 0.25$ and $\beta = 0.75$ respectively. Next, we look at the histograms of the p_t^i's at each t as the sample analog of μ_t corresponding to the initial agent prior μ_0. The evolution of those economist-priors are shown in Figure 8 for the case with $\beta = 0.25$ and Figure 10 for the case with $\beta = 0.75$. It is clear that the REH is justified in this framework due to full asymptotic learning (in the sense that $\mu_t \Rightarrow \delta_{p*=0}$). It is also clear that the higher value of δ led to a faster rate of learning.

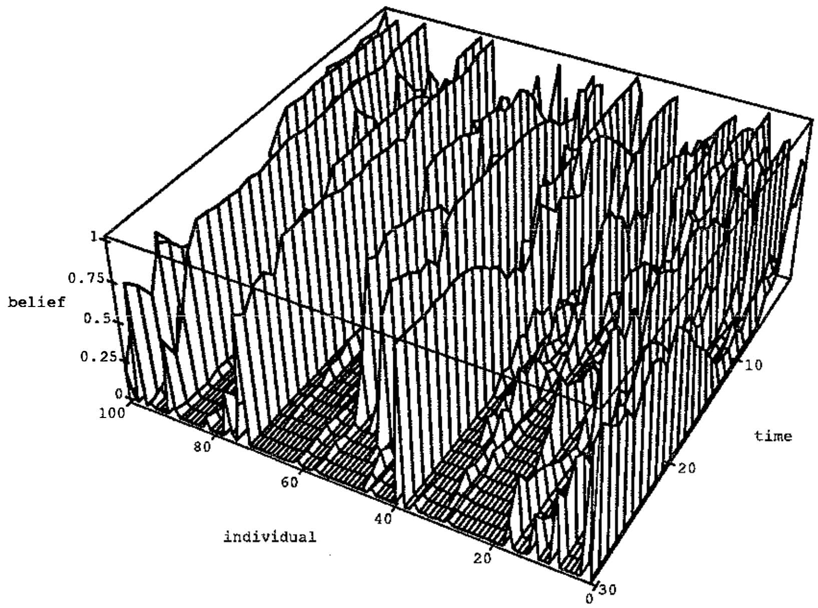

Figure 7: 100 trajectories of agent beliefs – $\beta = 0.25$

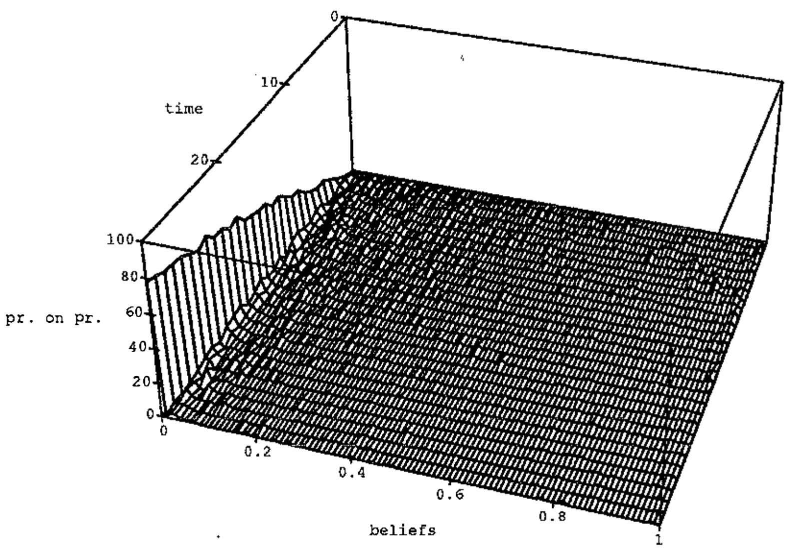

Figure 8: Trajectory of economist beliefs – $\beta = 0.25$

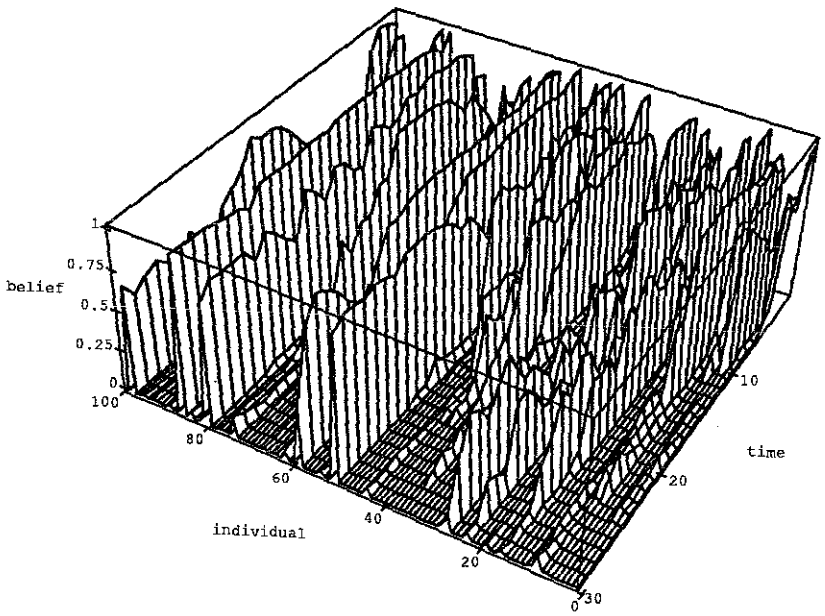

Figure 9: 100 trajectories of agent beliefs - $\beta = 0.75$

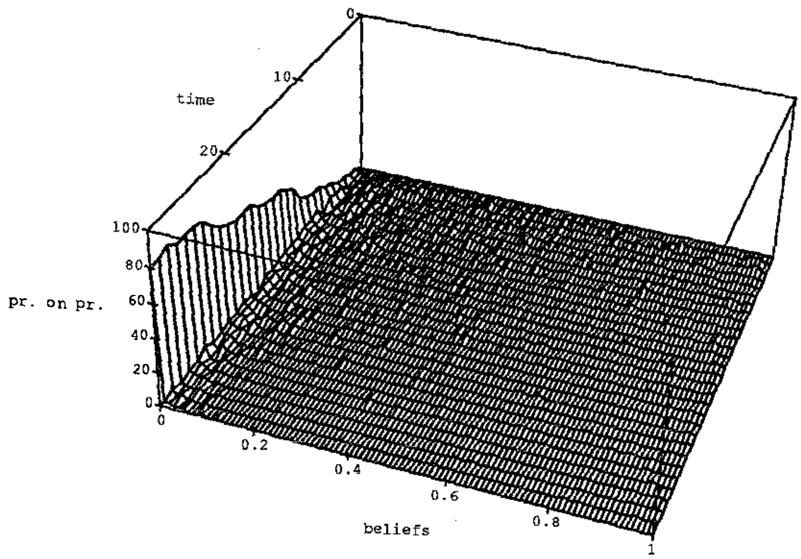

Figure 10: Trajectory of economist beliefs - $\beta = 0.75$

References

Easley, D. and N. Kiefer "Controlling a Stochastic Process with Unknown Parameters", *Econometrica*, 1988.

El-Gamal, E. and R. Sundaram "Bayesian Economists ... Bayesian Agents I: An Alternative Approach to Optimal Learning", Soc. Sc. working paper # 705, Caltech, 1989.

Feldman, M. "On the Generic Nonconvergence of Bayesian Actions and Beliefs", BEBR working paper #89-1528, University of Illinois- Urbana, Champaign, 1989.

Freedman, D. "On The Asymptotic Behavior of Bayes Estimates in the Discrete Case", ***Annals of Mathematical Statistics***, 1963, pp. 1386 -1403.

Freedman, D. " ... II", ***Annals of Mathematical Statistics***, 1965, pp. 454–456.

Kiefer, N. "A Value Function Arising in the Economics of Information", ***Journal of Economic Dynamics and Control***, vol. 13, 1989.