

Self-Organization and the Persistence of Noise in Financial Markets

David Goldbaum*
Rutgers University Newark

Abstract

A dynamic model of financial markets with learning is demonstrated to produce a self-organized system that displays critical behavior. The price contains private information that traders learn to extract and employ to forecast future value. Since the price reflects the beliefs of the traders, the learning process is self-referencing. As the market learns to correctly extract information from the price, the market deemphasizes private information. Despite the convergence of the model towards the parameters producing efficiency, pricing deviations remain constant due to the increased sensitivity of the price to small errors in information extraction produced by the model's own convergence.

JEL codes: G14, C62, D82

Keywords: Efficient Markets, Learning, Dynamics, Computational Economics

* Department of Economics, Rutgers University Newark, 360 Dr. Martin Luther King, Jr. Blvd., Newark, NJ 07302, phone 973 353-5164, fax 973 353-5819, goldbaum@andromeda.rutgers.edu. The author is grateful for comments received at the Conference on Computing in Economics and Finance in Amsterdam and the Society for Nonlinear Dynamics and Econometrics in London, healthy discussion at the CeNDEF, and useful input from Cars Hommes, Thomas Lux, and Allan Timmermann. Financial support was provided by the Rutgers University Research Council.

1. Introduction

When Grossman and Stiglitz (1980) find that no equilibrium exists able to produce efficient markets under rational expectations, they create one by introducing a second source of noise that hampers information extraction from the price. The equilibrium is created at the expense of market efficiency. Finding the conditions sufficient to create equilibrium is, of course, the traditional thrust of economic research. This paper explores a dynamic model of information extraction for which no equilibrium exists. Analytical analysis and simulation characterize the asymptotic behavior of a market in which an informationally efficient price is an unachievable attractor to the system. Despite the attraction and continuous improvement in information extraction, the market price does not converge towards the efficient market price.

A fully revealing efficient price is impossible in a rational expectations equilibrium. Grossman and Stiglitz (1980) (hereafter GS) reach this conclusion based on a market in which fundamental information is costly to obtain. They find that traders who are informed with private information have an impact on the price. Uninformed traders use the price to costlessly extract the private information. The GS paradox is that a “long run” equilibrium does not exist. No trader wants to bear the cost of informing the market without reward. Subsequent papers reach the same conclusion after relaxing the assumption of full rationality, for example by replacing rationality with dynamic learning on the part of the uninformed traders.¹ This paper further generalizes the GS model by examining dynamic populations in concert with a dynamic learning process.

Among those papers extending GS is Bray (1982) which introduced learning through repeated realizations of the GS single period model. Bray provides support for the GS

¹ Muendler (2003) demonstrates that some traders will pay to inform the market when the portfolio problem is part of a two period consumption problem.

assumption of full rationality by finding that learning converges to the rational expectations equilibrium (REE) for any fixed population proportion of uninformed to informed traders. Hussman (1992) and Timmermann (1996) also examine learning in a financial market setting. Like Bray, they have a static population of traders, but with a multi-period asset paying a dividend that follows an AR(1) process.

Routledge (1999) examines a dynamic population in a noisy GS model. With two types of traders, a dynamic population process is introduced by allowing traders to switch strategies based on imitating the strategy of the more successful traders they encounter. Thus, the single act of imitation embodies the two dynamic processes of learning and population evolution. Modeled with a random supply of the risky asset, the exogenous noise ensures the existence of a stable REE. The process generally converges to this REE.

Bak, Tang, and Weisenfeld (1988) introduce the notion of self-organized criticality (SOC) by describing what has become the canonical example of a system that attains SOC, the sand pile. The slope of the sand pile increases as grains of sand are added. At the critical slope an additional grain of sand causes an avalanche which locally reduces the slope and temporarily returns the system to stability. The size of the avalanche is determined by local conditions. The distribution of the avalanche size follows a power law. The self-organization refers to the natural convergence of the system to the attracting critical slope. The criticality refers to the chaotic behavior of the avalanches near the critical slope. A number of papers have applied the notion of SOC to economic settings, including Bak *et al* (1992), Bak, Paczuski, and Shubik (1997), Berg *et al* (2001), and Challet and Marsili (2002). The latter three describe financial markets as SOC.

Goldbaum (2005) separates the learning process from the population process in an examination of a market in which dividends follow a random walk. The model produces self-

organized criticality with the market attracted to a critical division in the population separating a stable region of the state space from an unstable region. The system naturally oscillates between the two states producing critical behavior. The market never approaches total adoption of market based analysis as the critical population requires a non-zero mass of informed traders. The market presented in this paper also contains a phase transition, but between the two more traditional stable states of an informed market or an uninformed market. It is the convergence towards the complete abandonment of the fundamental information, without any actual switching between states, that produces the interesting market behavior. The convergence properties under learning as discussed by Marcet and Sargent (1989a, 1989b) are keenly relevant.

The paper proceeds as follows. Section 2 introduces the model of the market. The population is divided into those who are informed with private fundamental information and those who attempt to extract information from the price. The nature of the private information is developed along with the trading strategy of those who obtain it. The method by which traders learn to extract information from the price is also described in this section. Section 3 presents analysis of the model under learning and introduces the dynamic population process. Simulation results characterize the model's asymptotic behavior, displaying the critical behavior produced by the model's self-organizing aspect. The model developed in Section 4 removes the artificial population division, allowing all traders to both receive private information and extract information from the price. This model is also examined through simulation. Conclusions are drawn in Section 5 of the paper.

2. Model

The market setting is similar to Goldbaum (2005) with the modification that the dividend follows an AR(1) process rather than a random walk. This small change produces substantially

different market behavior. The following provides a basic description of the environment and solutions based on the stationary dividend process.

2.1 The market

A large but finite number of agents, indexed by $i = 1, \dots, N$, trade a risky asset and a risk-free bond. The risk-free bond, with a price of one, pays R . The risky asset is purchased at the market determined price, p_t , in period t . In $t+1$, it pays a stochastic dividend d_{t+1} , and sells for the *ex dividend* market determined price p_{t+1} . The market participants are aware that the stochastic dividend follows a commonly known AR(1) process centered around d_0 :

$$d_t = d_0 + \eta_t,$$

$$\text{with } \eta_t = \phi\eta_{t-1} + \varepsilon_t, \varepsilon_t \sim IIDN(0, \sigma_\varepsilon^2). \quad (1)$$

Let $z_t = p_t + d_t$ and $\theta_{it} = 1/\gamma\sigma_{it}^2$ with $\sigma_{it}^2 = \text{var}_{it}(z_{t+1})$ indicating the conditional variance.

The parameter γ is the coefficient of absolute risk aversion. In each period, each myopic trader maximizes a negative exponential utility function on one period ahead wealth conditional on his individual information set (to be developed below). This produces the demand for the risky asset,

$$q_{it}(p_t) = \frac{E_{it}(p_{t+1} + d_{t+1} - Rp_t)}{\gamma \text{Var}_{it}(p_{t+1} + d_{t+1})} = \theta_{it}(E_{it}(z_{t+1}) - Rp_t). \quad (2)$$

Assume K strategies for estimating payoffs, z_{t+1} . In a Walrasian equilibrium, the market price equates supply and demand for the asset. Supply is fixed to avoid the exogenous introduction of noise. For convenience, set fixed net supply of the risky asset to zero. Let N_k be the total number of traders employing information I^k . Let q_t^k be the *per capita* demand for the

risky security among group k traders, $q_t^k(p_t) = \frac{1}{N_k} \sum_{i=1}^{N_k} q_{it}(p_t)$. Let n^k , $0 \leq n^k \leq 1$, be the

proportion of the trader population employing strategy k , $\sum n^k = 1$. The price p_t clears the market by solving

$$0 = \sum_{k=1}^K n^k q_t^k(p_t). \quad (3)$$

2.2 Information

2.2.1 The fundamental trader

The estimate of the future payoff is in the nature of Hellwig (1980). Trader i 's private research conducted at time t produces a noisy signal of the intrinsic value at time $t+1$. This information gathering process is captured by the time t private signal, s_{it} , centered around d_{t+1} , but subject to an idiosyncratic error term,

$$s_{it} = d_{t+1} + e_{it} = d_0 + \eta_{t+1} + e_{it}, \quad (4)$$

$$\text{with } e_{it} \sim IIDN(0, \sigma_e^2).$$

A linear projection of η_{t+1} onto the information set produces the fundamental investor's mean squared error minimizing forecast

$$E_i(\eta_{t+1} | \eta_t, s_{it}) = (1-\beta)\phi\eta_t + \beta s_{it} \quad (5)$$

where the weight β is known based on the traders' knowledge of the dividend and information processes,

$$\beta \equiv \text{cov}(\eta_{t+1}, s_{it}) / \text{var}(s_{it}) = \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \sigma_e^2}.$$

The "fundamental" price prevails in a market populated exclusively by fundamental investors. Derive the fundamental price by using the estimate (5) in (2),

$$p_t^F = p^F(\eta_t, \eta_{t+1}) = b_0 + b_1^F \eta_t + b_2^F \eta_{t+1} + v_t, \quad (6)$$

$v_t = \frac{\beta}{R-\phi} \frac{1}{N} \sum_i e_{it}$. Price is a function of the current private and public information.

Advancing (6) one period, substituting it into the demand (2) and using the market clearing condition (3), the price coefficients solve to $b_0 = d_0 / (R-1)$, $b_1^F = (1-\beta)\phi / (R-\phi)$ and $b_2^F = \beta / (R-\phi)$. The fact that prices, p_t^F here and the general market clearing price p_t elsewhere in the paper, are functions of η_t and η_{t+1} will be implicit, captured by the subscript t . Future discussions of the fixed point price refer to the pricing equation at which the coefficients take on their fixed point values. There is not a time independent value for the price at the fixed point.

For large N , the impact of the idiosyncratic signal noise on the price is negligible. Assume a sufficiently large N such that the v_t term can be dropped.²

For $\sigma_e^2 \rightarrow \infty$, $p_t^F \rightarrow p_t^0 = \frac{d_0}{R-1} + \frac{\phi}{R-\phi} \eta_t$. As the private signal becomes increasingly

noisy, the price converges to reflect just the public information contained in d_t , *i.e.* the price is Semistrong-form efficient according to the Fama (1970) definitions of efficiency. For $\sigma_e^2 = 0$,

$p_t^F = p_t^{EM} = \frac{d_0}{R-1} + \frac{1}{R-\phi} \eta_{t+1}$. When traders receive a perfect noise-free signal on the next

period's dividend, the price fully reflects the d_{t+1} based value producing a Strong-form efficient price. Between the two extremes of signal accuracy, p_t^F is not efficient, reflecting a weighted combination of the public d_t and the private d_{t+1} information.

Fundamental traders rely on (6) in forming demand. Plug (6) back into (2) to solve for the average demand of the group of fundamental traders,

$$q_t^F(p_t) = \theta_t^F \left(\frac{Rd_0}{R-1} + \frac{R}{R-\phi} ((1-\beta)\phi\eta_t + \beta\eta_{t+1}) - Rp_t \right). \quad (7)$$

² Formally, v_t is $o(1)$.

2.2.2 Regression traders

The regression traders model the relationship between the payoff, $z_{t+1} = p_{t+1} + d_{t+1}$, and current market observables. The traders appropriately estimate

$$z_t = c_0 + c_1 p_{t-1} + c_2 \eta_{t-1} + \zeta_t. \quad (8)$$

The traders employ least-squares learning to update the parameters of their model. The learning process is self-referential with an endogenous state variable, p_{t-1} , included as a regressor. The value of c_0 is exogenous to beliefs and can be derived analytically by the traders given their knowledge of the dividend process. Thus traders impose the correct value of c_0 . Let $x_t' = [1 \ p_t \ \eta_t]$. The regression traders update the coefficients, $\mathbf{c}_t = [c_0 \ c_{1t} \ c_{2t}]$, using the standard recursive updating algorithm for least-squares learning of Marcet and Sargent (1989a, 1989b):

$$\mathbf{c}_t = \mathbf{c}_{t-1} + (Q_t^{-1} x_{t-2}' (z_{t-1} - \mathbf{c}_{t-1} x_{t-2}'))' / t, \quad (9)$$

$$Q_t = Q_{t-1} + (x_{t-1}' x_{t-1}' - Q_{t-1}) / t,$$

given (\mathbf{c}_0, Q_0) . The regression traders all rely on the same public information, and thus all employ the same forecast, $E(z_{t+1} | x_t)$. *Per capita* demand among regression traders is thus

$$q_t^R = \theta_t^R (\mathbf{c}_t x_t - R p_t). \quad (10)$$

2.3 Price Formation

With $K = 2$, let $n_t = n_t^F$, and thus $(1-n_t) = n_t^R$. From (3),

$$0 = n_t q_t^F + (1-n_t) q_t^R. \quad (11)$$

Use (7), (10), and (11) to solve for the market clearing price. A consistent price function takes the form

$$p_t = p_t(n_t, \mathbf{c}_t) = b_0 + b_1(n_t, \mathbf{c}_t) \eta_t + b_2(n_t, \mathbf{c}_t) \eta_{t+1} \quad (12)$$

with

$$b_0 = d_0 / (R - 1),$$

$$b_1(n_t, \mathbf{c}_t) = \Psi(n_t, \mathbf{c}_t)^{-1} (n_t \theta_t^F \frac{R}{R-\phi} (1-\beta)\phi + (1-n_t)c_{2t} \theta_t^R), \quad (13)$$

$$b_2(n_t, \mathbf{c}_t) = \Psi(n_t, \mathbf{c}_t)^{-1} (n_t \theta_t^F \frac{R}{R-\phi} \beta), \text{ and}$$

$$\Psi(n_t, \mathbf{c}_t) = n_t R \theta_t^F + (1-n_t)(R - c_{1t}) \theta_t^R.$$

3. Analysis and simulation

3.1 Learning under fixed n

Consider just the learning process by fixing $n_t = n$. A fixed point to the learning process given n is established. Based on this solution, the implicit evolution in the population can be developed and examined before resorting to simulations.

3.1.1 A fixed point solution

Three equations describe the dynamic processes under a fixed n . Equation (1) is the exogenous dividend process. Equation (12) is the endogenous price equation. The coefficients of (12) depend on the beliefs of the regression traders as captured by (8) that evolve according to (9) and upon the fixed proportion of the market relying on fundamental analysis.

The fixed point for the learning process produces a fixed point pricing function $p_t^*(n)$ based on the solution, for $0 < n \leq 1$:

$$b_1^*(n) = \frac{n\theta^F (1-\beta)\phi}{(R-\phi)(n\theta^F + (1-n)\theta^R)}, \quad b_2^*(n) = \frac{n\theta^F \beta + (1-n)\theta^R}{(R-\phi)(n\theta^F + (1-n)\theta^R)}, \quad (14)$$

$$c_1^*(n) = \frac{R}{(R-\phi)b_2^*(n)} = \frac{R(n\theta^F + (1-n)\theta^R)}{n\beta\theta^F + (1-n)\theta^R}, \quad (15)$$

$$c_2^*(n) = \frac{\phi}{R-\phi} (R - c_1^*(n)) = -\frac{n(1-\beta)\phi\theta^F}{(R-\phi)(n\beta\theta^F + (1-n)\theta^R)},$$

$$\sigma_F^*(n)^2 = ((1-\beta)(R/(R-\phi))^2 + b_2^*(n)^2) \sigma_\varepsilon^2, \text{ and} \quad (16)$$

$$\sigma_R^*(n)^2 = b_2^*(n)^2 \sigma_\varepsilon^2.$$

For $n = 0$ the regression captured by (8) is undefined since p_{t-1} is a linear function of η_{t-1} . Note that $n \rightarrow 0$ produces $c_1^*(n) \rightarrow R$, $c_2^*(n) \rightarrow 0$ and for $n = 0$, $c_1^*(0) = R$ and $c_2^*(0) = 0$ is a consistent solution producing $b_1^*(0) = \phi/(R-\phi)$ and $b_2^*(0) = 0$.

Regardless of n , $b_1^*(n) + \phi b_2^*(n) = \phi/(R-\phi)$. This result ensures that the fundamental trader's demand is unaffected by n . Equation (7), derived based on $n = 1$, is correct for the mixed population setting as well, substantiating the use of (7) in deriving the market equilibrium solution.

The solution described by (14) through (16) is implicitly defined since $b_2^*(n)$ in (14) and the conditional variance terms in (16) remain interdependent.³ The fixed point solutions based on the typical exogenous parameter values used in the simulations are plotted in Figure 1 as functions of n .

[Figure 1 about here]

At the fixed point the regression traders correctly deduce the relationship between price and payoff. This is equivalent to knowing the value of d_{t+1} . Essentially, at the fixed point the regression traders know the values of the price coefficients in (12) and thus can use the price to derive the value of η_{t+1} and thus d_{t+1} . That the price reveals d_{t+1} is not equivalent to having the price fully reflect d_{t+1} , the latter requiring $p_t = p_t^{EM}$. In general

$$p_t - p_t^{EM} = (b_{1t} + \phi b_{2t} - \phi/(R-\phi))\eta_t + (\phi b_{2t} - \phi/(R-\phi))\varepsilon_{t+1}. \quad (17)$$

Since $b_1^*(n) + \phi b_2^*(n) = \phi/(R-\phi) \forall n$, the first term of (17) is equal to zero for all $p_t^*(n)$. The first term thus takes a zero value when the regression model is correct, regardless of n . The

³ Solving for $b_2^*(n)$ involves solving a cubic polynomial. The unique real solution is available, but lengthy.

second term only converges to zero when the regression model is correct and $n \rightarrow 0$. With

$$p_t^*(0) = p_t^0 \text{ there is no value of } n \text{ at which } p_t^*(n) = p_t^{EM}.^4$$

3.1.2 Performance

Let the measure of profits earned by each information source be the excess return realized for the risky asset multiplied by the group average demand:

$$\pi_t^k = q_t^k (p_{t+1} + d_{t+1} - Rp_t), k = F, R. \quad (18)$$

Based on the fixed point solution, the modeler with knowledge of n can compute

$$\left. \begin{aligned} E(\pi^F) &= -\theta^{F^2} \theta^R \left(\frac{R}{R-\phi} \right)^2 \left(\frac{n(1-n)(1-\beta)^2}{(n\theta^F + (1-n)\theta^R)^2} \right) \sigma_\varepsilon^2 \\ E(\pi^R) &= \theta^{F^2} \theta^R \left(\frac{R}{R-\phi} \right)^2 \left(\frac{n(1-\beta)}{n\theta^F + (1-n)\theta^R} \right)^2 \sigma_\varepsilon^2 \end{aligned} \right\} \text{for } 0 < n \leq 1 \text{ and} \quad (19)$$

$$E(\pi^F) = \beta/(1-\beta), E(\pi^R) = 0 \text{ at } n = 0. \quad (20)$$

Thus, for all values of $n > 0$, the fixed point expected profits are positive for the regression traders and weakly negative for fundamental traders. The regression traders outperform the fundamental traders because with a correct model of the relationship between price and payoff, the price is a better source of information, revealing d_{t+1} , than is the noisy signal.

A discrete jump in profits occurs at $n = 0$, reflecting the benefits to even noisy information on η_{t+1} , when the price reflects only the public time t information. There is no value of n at which $E(\pi^F) = E(\pi^R)$.

3.2 Evolution in the population

Allow for dynamics in the population proportion. The Replicator Dynamic model of population evolution produces an dynamic population in which the dominant strategy attracts

⁴ Efficient market is in the nature of Radner (1979). The price reflects the aggregation of the fundamental traders knowledge since no individual trader knows d_{t+1} with certainty.

converts from the inferior strategy. The two choice version of the more general K choice replicator dynamic model found in Branch and McGough (2003) results in the transition equation

$$n_{t+1} = \begin{cases} n_t + r(\pi_t^{Fe} - \pi_t^{Re})(1 - n_t) & \text{for } \pi_t^{Fe} \geq \pi_t^{Re} \\ n_t + r(\pi_t^{Fe} - \pi_t^{Re})n_t & \text{for } \pi_t^{Fe} < \pi_t^{Re} \end{cases} \quad (21)$$

where π_t^{Fe} and π_t^{Re} indicate the traders' performance measures of fundamental and market-based approaches, respectively. These are updated according to the process

$$\pi_t^{ke} = \pi_{t-1}^{ke} + (\pi_{t-1}^k - \pi_{t-1}^{ke})/t, \quad (22)$$

$$\pi_0^{ke} = 0, \quad k = F, R. \quad ^5$$

A number of different functional forms for r exist in the literature. The simulations that follow are based upon

$$r(x) = \tanh(\delta x / 2), \quad (23)$$

a choice that ensures $0 < n_t < 1$ for bounded $\pi_t^{Fe} - \pi_t^{Re}$. Thus, by construction, the discontinuity of $n = 0$ will never be realized in the simulation. The parameter δ determines how responsive the population is to differences in expected profits.

From (19), given $\mathbf{c}_t = \mathbf{c}^*(n_t)$ there is no value of n_t that produced equal profits, thus there is no joint fixed point to both the learning and population processes. Without the assumption $\mathbf{c}_t = \mathbf{c}^*(n_t)$, the learning and the population processes must evolve together. In Figure 2, the 3-dimensional phase space in n_t , c_{1t} , and c_{2t} has been collapsed to two dimensions by setting c_{2t} to be consistent with the c_{1t} parameter, a feat accomplished by replacing $c_1^*(n)$ in the $c_2^*(n)$ solution of (15) with c_{1t} . In the phase space plotted in Figure 2, the learning process updates the

⁵ For continuity in estimation, assume all traders maintain and update the regression trader's model which they ignore as fundamental trades.

regression coefficients and thus moves the model vertically in the phase space. The population process creates an evolution in n_t and thus moves the model horizontally. The curve labeled “ c_1^* ” plots the n_t dependent fixed point values $c_1^*(n_t)$ that represent the correct c_1 for the given n_t (the same curve labeled “ c_1^* ” in the second row of Figure 1).

The curves $c_1^-(n)$ and $c_1^+(n)$ are the two outer boundaries between which the regression traders’ parameters are sufficiently accurate to produce positive expected regression trader profits. Above $c_1^*(n)$ the regression traders over-react to price innovations. Along $c_1^+(n)$, expected profits are zero because the over-reaction forces $p_t(\mathbf{c}_t^+, n) = p_t^{EM}$ so that $E(p_{t+1} + d_{t+1} - Rp_t)$ in (18) is equal to zero. Along $c_1^-(n)$, the regression traders’ under-reaction to price innovations produces $p_t(\mathbf{c}_t^-, n) = p_t^F$ and thus $q_t^F = q_t^R = 0$ in (18). Above $c_1^+(n)$ and below $c_1^-(n)$ the regression traders’ beliefs are sufficiently in error to produce expected profits for the fundamental traders.

At $c_{1t} = \bar{c}_1$, the regression traders’ strong positive response to a price increase creates an upward sloping demand curve. The curve thus represents an n dependent upper bound on the value of c_{1t} . By (2) demand is attenuated by large estimates in the conditional variance. The regression traders estimate the conditional variance based on the mean squared regression error. Including the current period’s price in the regression traders’ current estimate of the conditional variance allows the model parameters to exist above \bar{c}_1 but still produce finite market prices.

Analysis of the phase space suggests a convergence path between the curves c_1^* and c_1^+ that, under the RD, leads to the point $n_t = 0$ and $c_{1t} = R$.

[Figure 2 about here]

3.3 Simulations

Simulations are necessary to determine convergence properties when allowing for the interaction between the learning and population processes. Figures 3 through 6 display the typical evolution of endogenous parameters produced by $\delta = 0.01$ and $\delta = 1$. Each frame plots the time progression of one of the endogenous parameters of the model. Across the top row are plotted the price parameters $b_{1t} + \phi b_{2t}$ and ϕb_{2t} , each determining the coefficients of the two terms in (17). The time series are plotted in grey. A solid black line at $\phi/(R - \phi)$ indicates the attracting value that also produces an efficient price. The second row contains the regression coefficients from the regression traders' extraction, c_{1t} and c_{2t} , plotted in grey. The black curve plots the respective element of $\mathbf{c}_t^*(n_t)$, the correct coefficient values for the regression traders given the n_t value. The lines at R and zero in their respective frames represent the attracting values of the learning process as $n_t \rightarrow 0$. The bottom row presents the population parameter n_t and the difference $p_t - p_t^{EM}$.

The data plotted in Figures 3 are from the first 50,000 periods of a simulation based on $\delta = 0.01$ after dropping the first 100 observations. The parameter δ from (23) determines the rate of adjustment in the population towards the superior performing choice. The smooth progression of n_t is the result of a population reluctant to switch approaches. The slow moving n_t produces a slow moving target to the learning process, making it easy for the regression parameters to keep pace with the changing market environment.

[Figure 3 about here]

It seems reasonable to consider the extent to which the results may be driven by the parameters of the model. Given that the price remains bounded, the initial values and the parameters do not substantially change the results of the model. The regression parameters and

the population simultaneously converge, $\mathbf{c}_t \rightarrow \mathbf{c}^*(n_t)$ and $n_t \rightarrow 0$. The state may change substantially from period to period early in the simulation when a single observation can substantially change \mathbf{c}_t or $\delta(\pi_t^{Fe} - \pi_t^{Re})$, the latter producing a large jump in n_t . As t increases, the single period innovations in \mathbf{c}_t , π_t^{Fe} , and π_t^{Re} become small. Figure 4 demonstrates that a large δ introduce cycles into the convergence without undermining the fact of the convergence.

[Figure 4 about here]

The declining cycles seen in Figure 4 (and more pronounced in Figure 6) results from the evolution in n_t outpacing the evolution in the regression coefficients. The higher $\delta = 1$ setting produces greater responsiveness in the trader population to differences in performance. Profits earned by the accuracy of the regression equation cause n_t to drop, but at a rate faster than the regression coefficients can keep pace. When n_t becomes too low for the contemporaneous regression parameters, the pricing error allows the fundamental traders to profit, temporarily reversing the progress in n_t .

For both $\delta = 0.01$ and $\delta = 1$ the system shows convergence towards $n_t = 0$. Figures 5 and 6 display the parameter values from the same simulation during the $t = 1.5 \times 10^6$ to $t = 2.0 \times 10^6$ interval. The most striking characteristic is that despite the continued convergence of the population and learning processes toward values that analytically produce greater efficient markets, the price and the price coefficients do not seem to reflect this advancement.

[Figure 5 about here]

[Figure 6 about here]

The convergence in both the learning and population parameters under the two values of δ demonstrates that the self-organization of the market is the interplay between learning, profits, and the population. Early in the simulation the regression coefficients quickly converged to

reflect the concurrent value of n_t . In (17), errors in beliefs by the regression traders cause the first term to deviate from zero, but for large n_t , the second term dominates, becoming the determining factor for the failure of market efficiency. The accuracy of regression produces $p_t \approx p_t^*(n_t)$, but the regression traders' market impact remains low so that $p_t^*(n_t)$ remains distant from p_t^{EM} .

The accuracy of the regression trader beliefs produce superior performance, driving n_t towards zero and $p_t^*(n_t) \rightarrow p_t^{EM}$. For small values of n_t , $\phi b_2^*(n_t) - \phi/(R - \phi)$ is approximately zero and yet the simulation produces a price consistently different from p_t^{EM} . The errors in $b_{1t} \neq b_1^*$ and $b_{2t} \neq b_2^*$ caused by $\mathbf{c}_t \neq \mathbf{c}^*(n_t)$ mean neither term of (17) is zero. Once n_t is sufficiently small, deviation in b_{2t} from $b_2^*(n_t)$ becomes a greater source of error in $\phi b_{2t} - \phi/(R - \phi)$ than the difference between $b_2^*(n_t)$ and $b_2^*(n)$ as $n \rightarrow 0$. Further convergence in $\mathbf{c}_t \rightarrow \mathbf{c}^*(n_t)$ and n_t towards zero does not noticeably improve the market efficiency.

Intuition might lead one to believe that continued learning by the regression traders would increase the efficiency of the price. Two explanations for why this is not the case deserve discussion. The convergence of $\mathbf{c}_t \rightarrow \mathbf{c}^*(n_t)$ produces progression in n_t towards zero. This changes the target for \mathbf{c}_t in the learning process, impeding further model improvement. Improvement occurs nonetheless and thus this endogenously produced error in learning is not the underlying cause of the failure to observe $p_t \rightarrow p_t^{EM}$. The decreasing n_t also produces instability by increasing the need for accuracy on the part of the regression traders. From (14), assuming $\mathbf{c}_t = \mathbf{c}^*(n)$, $n_t \rightarrow 0$ produces $b_1^*(n) \rightarrow 0$ and $\phi b_2^* \rightarrow \phi/(R - \phi)$. Thus, as $n_t \rightarrow 0$, $p_t^*(n_t) \rightarrow p_t^{EM}$. Without the assumption of $\mathbf{c}_t = \mathbf{c}^*(n_t)$, an examination of (13) reveals that both the numerator

and denominator of both $b_1(n_t, \mathbf{c}_t)$ and $b_2(n_t, \mathbf{c}_t)$ converge to zero as $\mathbf{c}_t \rightarrow \mathbf{c}^*(n_t)$ and $n_t \rightarrow 0$; $b_1(n_t, \mathbf{c}_t) \rightarrow 0/0$ and $b_2(n_t, \mathbf{c}_t) \rightarrow 0/0$. The closer the true market is to efficiency, the greater the impact of a small error in the regression traders' model on the price. The market is self-organized to remain inefficient even as $\mathbf{c}_t \rightarrow \mathbf{c}^*(n_t)$ and $n_t \rightarrow 0$.

Consider again the phase space in Figure 2. Recall that along $c_1^+(n), p_t = p_t^{EM}$. The convergence of $c_1^+(n) - c_1^* \rightarrow 0$ is a reflection of (17), that the price produced by the correct beliefs of the regression traders produces the convergence $p_t \rightarrow p_t^{EM}$. Recall as well that along $c_1^-(n), p_t = p_t^0$. The narrowing of the distance between $c_1^-(n)$ and $c_1^+(n)$ as $n_t \rightarrow 0$ indicates a narrowing in the pricing mechanism's tolerance for error in the regression trader's beliefs. As the market becomes increasingly reliant on the regression traders to set the price, small errors in the regression traders' model produced substantial pricing error.

The self-organization of the system contributes to the stability of the pricing error distribution. A tendency for substantial error in the regression model results in a large n_t , putting the model in a region in which the regression traders have little ability to produce wildly incorrect prices. Only once the regression traders are fairly accurate with their model can n_t converge towards zero, a region in which the market is heavily dependent on their beliefs. Profits regulate n_t . The system only achieves $n_t \rightarrow 0$ when the regression trader model is sufficiently correct and stable such that doing so does not cause an increase in the pricing error. An increasingly correct regression trader model produces a decline in the amount of self produced error in the price so that $|p_t - p_t^*(n_t)|$ and $|p_t - p_t^{EM}|$ should both fall. It is the resulting accuracy induced decline in n_t that increases the requirement for model accuracy, reversing the gains in $|p_t - p_t^*(n_t)|$. The net result is zero improvement in the market

efficiency. The stability of the distribution of the pricing error despite the improving regression trader model is consistent with the GS finding that information content of the price is unaffected by the variance of their model's exogenous noise.

4. Complete Information

4.1 Learning

In this section traders are allowed to employ both the public price and the private signal in determining a forecast of the value of the risky security. Traders are modeled as learning agents individually estimating the relationship⁶

$$z_t = c_0 + c_{1it}p_{t-1} + c_{2it}\eta_{t-1} + c_{3it}s_{it} + \zeta_{it}. \quad (24)$$

Let $\mathbf{c}_{it} = [c_0 \ c_{1it} \ c_{2it} \ c_{3it}]$ be the estimated individual regression coefficients, where $x_t' = [1 \ p_t \ \eta_t \ (\eta_{t+1} + e_{i,t})]$. Individual demand is

$$q_{it} = (\mathbf{c}_{it}x_t - Rp_t) / \gamma\sigma_{it}^2. \quad (25)$$

The equilibrium price solution that sets $\sum_i q_{it} = 0$ takes the same linear structure as in (12),

$$p_t = b_0 + b_1(\mathbf{c}_t)\eta_t + b_2(\mathbf{c}_t)\eta_{t+1}, \quad (26)$$

but with

$$b_1(\mathbf{c}_t) = \Psi_t^{-1} \sum_i \theta_{it} c_{2it}, \quad b_2(\mathbf{c}_t) = \Psi_t^{-1} \sum_i \theta_{it} c_{3it}, \quad (27)$$

$$\Psi_t = \sum_i \theta_{it} (R - c_{1it}).$$

All traders share the same time consistent beliefs at the fixed point. The fixed point for the learning process produces the regression coefficients

$$c_1^* = R, \quad c_2^* = 0, \quad c_3^* = 0.$$

⁶ The model closely resembles that of Hellwig (1982) in its structure, except that the traders employ the contemporaneous price rather than the lagged price.

The price and the price coefficients are indeterminate at the fixed point,

$$b_1^* = c_2^* / (R - c_1^*), \quad b_2^* = c_3^* / (R - c_1^*).$$

As the system approaches the fixed point, traders are able to rely heavily on the price for information on the value of d_{t+1} . They place positive, but negligible, weight on their noisy private signal since the price is an almost perfect indicator of d_{t+1} .

At the fixed point, both price and private information are removed from the agents' demand function in (25) and demand for the risky asset becomes zero. The No Trade solution is attained. The price is removed because the traders believe the market has converged to the point at which p_t fully reveals d_{t+1} , removing the ability of the price to reveal profitable trades. At the same time, the noisy private information is completely dominated by p_t , and thus receives zero weight. In a GS type paradox, the reality of the fixed point is just the opposite of the traders' belief. The price contains no information, both because no one trades based on the d_{t+1} information and because any price clears the market.

4.2 Simulations

Simulations produce market behavior much like that produced by the Replicator Dynamics. The typical progression can be seen in Figures 7 and 8. The positive value of the coefficient c_{3t} indicates that the traders choose to rely on the fundamental signal. The convergence of c_{3t} towards zero suggests that the traders are able to reduce their reliance on the private signal as the price becomes an increasingly reliant indicator of d_{t+1} . The greater reliance on the price is reflected in the increasing value of c_{1t} towards its fixed point value of R . As previously obtained, the convergence of the regression coefficients towards the efficient market values does not produce convergence in p_t towards p_t^{EM} .

[Figure 7 about here]

[Figure 8 about here]

5. Conclusion

The analysis explores a dynamic financial market model that lacks a fixed point. The unattainable point of attraction is market efficiency in which the price fully reflects the aggregate of private information. Though the market is constructed to preclude the existence of an equilibrium with an efficient price, this absence is not the driving force behind the failure to achieve market efficiency. The failure to produce an efficient price is the result of the interaction between the two dynamic processes, accuracy of the model extracting information from the price and the extent to which the market relies on this model to set the price.

Traders maintain and update a model of the relationship between observed price and future payoff. By construction, if the traders have an accurate model of the pricing equation, then the price is a better source of information than the noisy signal received through individual fundamental based research. The market serves as a filter on individual traders' errors, leaving the price as an error-free indicator of the underlying information. Improvements in this pricing model should reasonably produce both increased accuracy of the market price at reflecting value and convergence towards an efficient market. This does not turn out to be the case.

Improvement in the ability to extract information from the price elicits greater reliance on market-based information by the traders. As the population of traders decreases its dependence on the fundamental information the tolerance for error in the perceived model also declines. As the traders become increasingly accurate in their extraction of information from the price the market requires an increasingly accurate model.

The rate at which the traders' pricing model improves and the rate of increasing reliance on the market-based information co-evolve. Each modulates the other to maintain balance so that

the magnitude of the market's error in pricing remains stable as the other parameters in the model continue to converge.

References

- Bak, P., Tang, C., Wiesenfeld, K., 1988. Self-organized criticality. *Physical Review A*, 364-374.
- Bak, P., Chen, K., Scheinkman, J., Woodford, M., 1992. Aggregate fluctuations from independent sectoral shocks: self-organized criticality. NBER Working Paper Series No. 4241.
- Bak, P., Paczuski, M., Shubik, M., 1997. Price variations in a stock market with many agents. *Physica A* 245, 430-453.
- Berg, J., Marsili, M., Rustichini, A., Zecchina, R., 2001. Statistical mechanics of asset markets with private information. *Quantitative Finance* 1, 203-211.
- Branch, W., McGough, B., 2003. Replicator dynamics in a cobweb model with rationally heterogeneous expectations. working paper.
- Bray, M., 1982. Learning, estimation, and the stability of rational expectations. *Journal of Economic Theory* 26, 318-339.
- Challet, D., Marsili, M., 2002. Criticality and finite size effects in a simple realistic model of stock market. cond-mat/0210549 v2.
- Goldbaum, D.H., 2005. Market efficiency and learning in an endogenously unstable environment, *Journal of Economics Dynamics and Control* 29(5), 953-978.
- Grossman, S.J., Stiglitz, J.E., 1980. On the impossibility of informationally efficient markets. *The American Economic Review* 70(3), 393-408.
- Fama, E.F., 1970. Efficient capital markets: a review of theory and empirical work. *The Journal of Finance* 25(2), 383-417.
- Hellwig, M.F., 1980. On the aggregation of information in competitive markets. *Journal of*

- Economic Theory 22, 477-498.
- Hellwig, M.F., 1982. Rational expectations equilibrium with conditioning on past prices: a mean-variance example. *Journal of Economic Theory* 26, 279-312.
- Hussman, J.P., 1992. Market efficiency and inefficiency in rational expectations equilibria. *Journal of Economic Dynamics and Control* 16, 655-680.
- Manski, C.F., McFadden, D., 1981. *Structural Analysis of Discrete Data with Econometric Applications* (MIT Press, Cambridge, MA).
- Marcet, A., Sargent, T.J., 1989a. Convergence of least squares learning mechanisms in self-referential linear stochastic models. *Journal of Economic Theory* 48, 337-368.
- Marcet, A., Sargent, T.J., 1989b. Convergence of least-squares learning in environments with hidden state variables and private information. *The Journal of Political Economy* 97(6), 1306-1322.
- Muendler, M.A., 2003. Demand for information on assets with Gaussian returns. working paper.
- Radner, R., 1979. Rational expectations equilibrium: generic existence and the information revealed by prices. *Econometrica* 47(3), 655-678.
- Routledge, B.R., 1999. Adaptive learning in financial markets. *The Review of Financial Studies* 12(5), 1165-1202.
- Sethi, R., Franke, R., 1995. Behavioral heterogeneity under evolutionary pressure: Macroeconomic implications of costly optimization. *The Economic Journal* 105, 583-600.
- Timmermann, A., 1996. Excess volatility and predictability of stock prices in autoregressive dividend models with learning. *Review of Economic Studies* 63, 523-557.

Figure 1 Fixed point values of $b_1^*(n)$, $b_2^*(n)$, $c_1^*(n)$, $c_2^*(n)$, $\sigma_F^*(n)$, $\sigma_R^*(n)$ plotted against n .

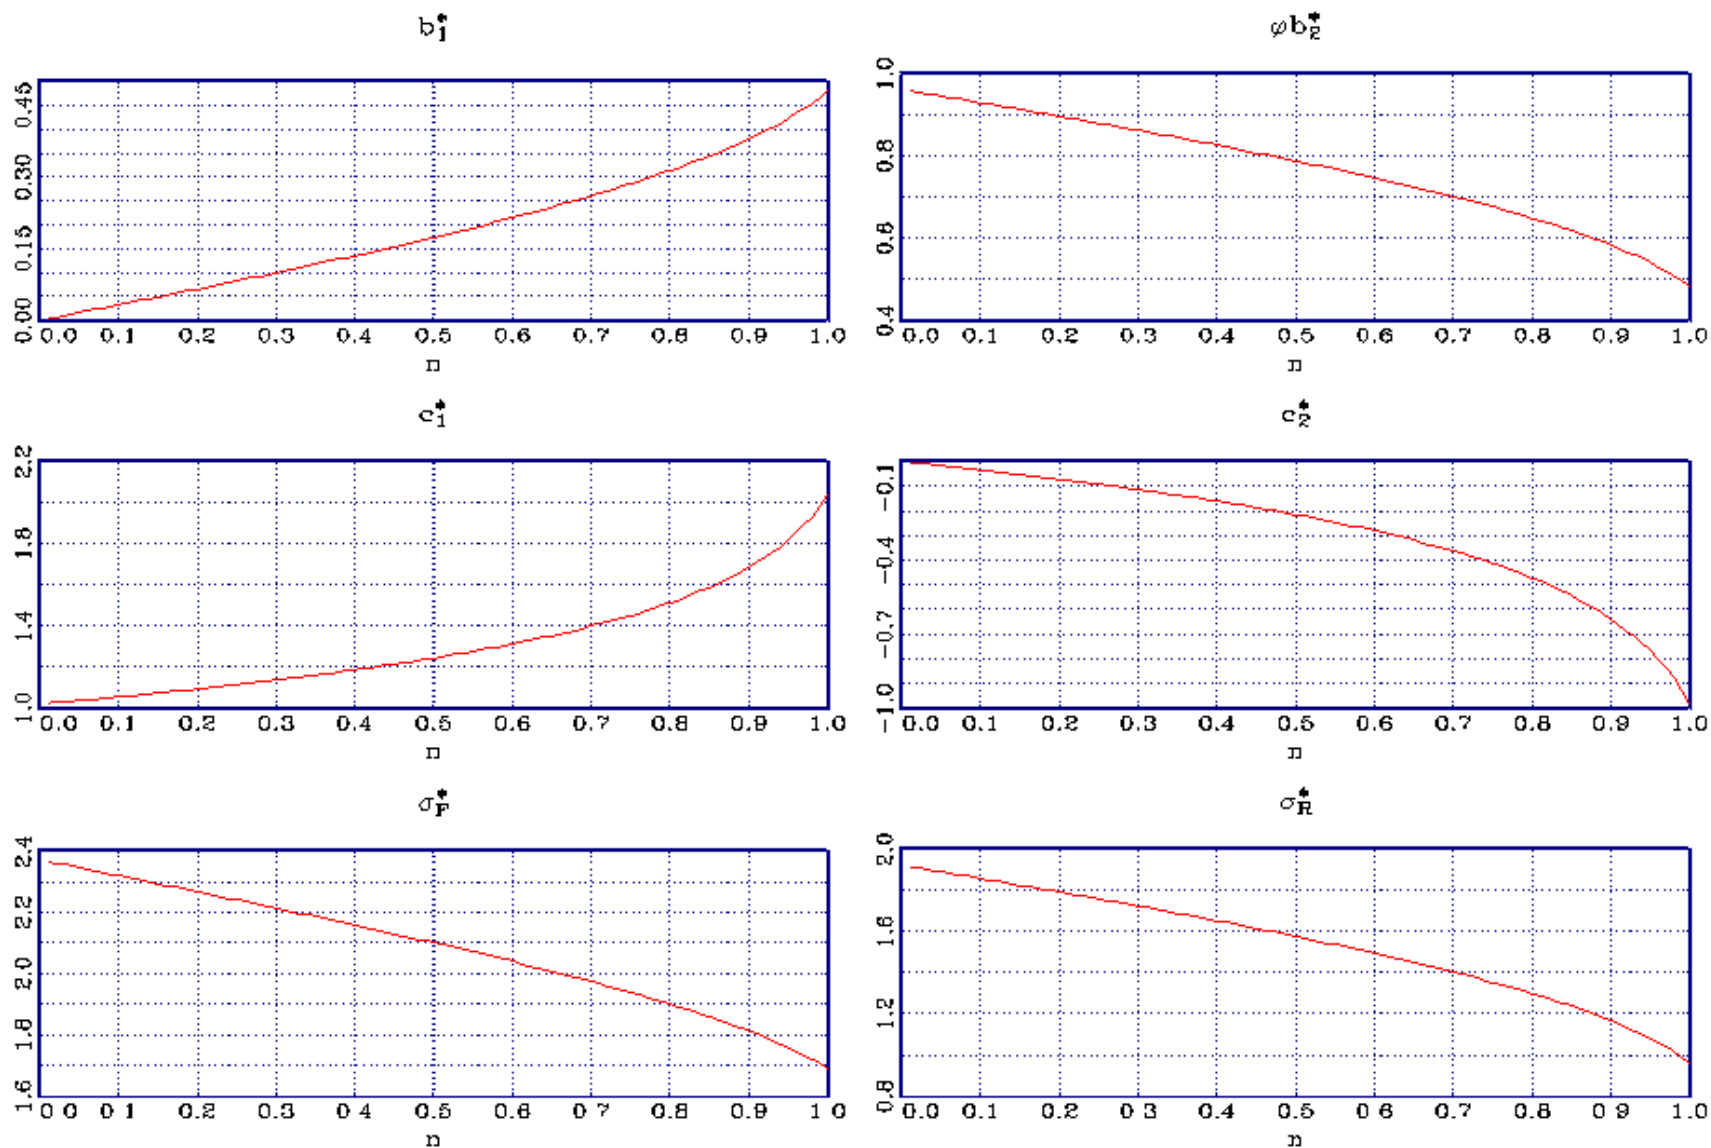


Figure 2: Phase-space in n_t and c_t^1 . The curve labeled c_1^* is the correct n dependent regression coefficient. The curve labeled \bar{c}_1 is an upper bound on c_{1t} for the market price to remain finite. For c_{1t} between c_1^+ and c_1^- the regression model is sufficiently accurate to generate profits for the regression traders. Outside of this region, the fundamental traders are expected to profit.




Figure 3: Evolution of endogenous parameters. RD dynamics, $\delta = 0.01$, beginning of sample ($t = 100 - 50,100$). Row one left: $b_{1t} + \phi b_{2t}$ with the solid black line at $\phi/(R - \phi)$. $b_{1t} + \phi b_{2t} = \phi/(R - \phi)$ implies the regression traders have the correct model. Row one right: ϕb_{2t} with the solid black line at $\phi/(R - \phi)$. $\phi b_{2t} = \phi/(R - \phi)$ implies an efficient market. Row two left: c_{1t} (lighter shade or green) and $c_1^*(n_t)$ (darker shade or red). Row two right: c_{2t} (lighter shade or green) and $c_2^*(n_t)$ (darker shade or red). $c_1^*(n_t) \rightarrow R$ and $c_2^*(n_t) \rightarrow 0$ as $n_t \rightarrow 0$. Row three left: n_t . Row three right: $p_t - p_t^{EM}$ with a solid black line a zero.

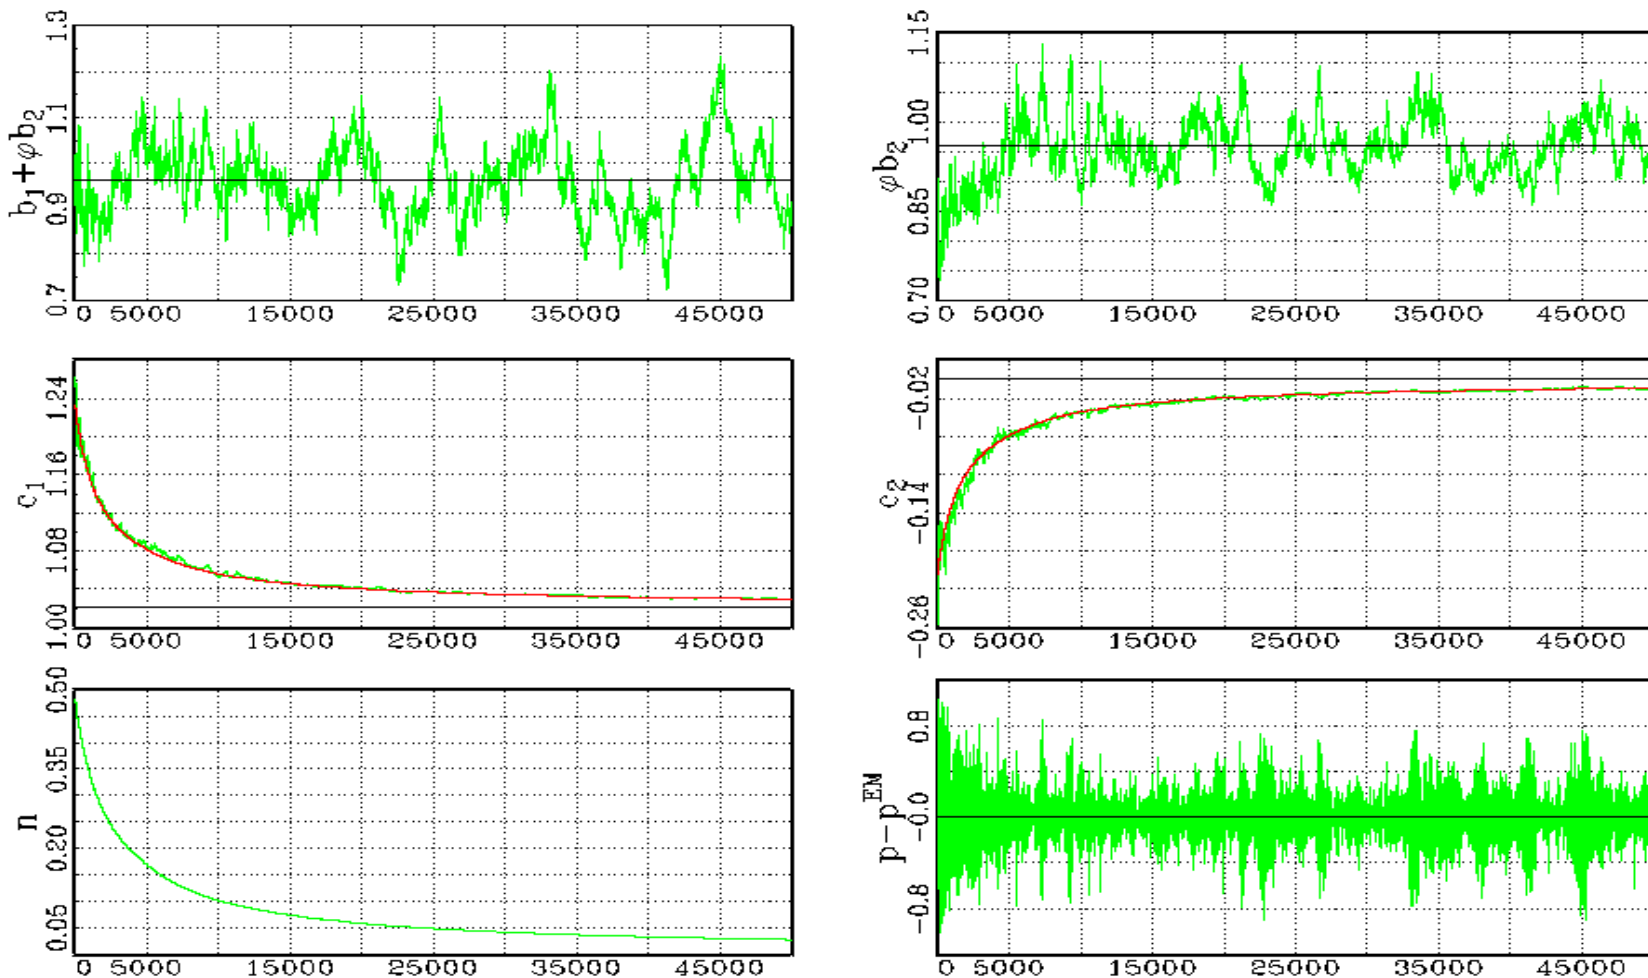


Figure 4: Evolution of endogenous parameters. RD dynamics, $\delta = 1$, beginning of sample ($t = 100 - 50,100$). Row one left: $b_{1t} + \phi b_{2t}$ with the solid black line at $\phi/(R - \phi)$. $b_{1t} + \phi b_{2t} = \phi/(R - \phi)$ implies the regression traders have the correct model. Row one right: ϕb_{2t} with the solid black line at $\phi/(R - \phi)$. $\phi b_{2t} = \phi/(R - \phi)$ implies an efficient market. Row two left: c_{1t} (lighter shade or green) and $c_1^*(n_t)$ (darker shade or red). Row two right: c_{2t} (lighter shade or green) and $c_2^*(n_t)$ (darker shade or red). $c_1^*(n_t) \rightarrow R$ and $c_2^*(n_t) \rightarrow 0$ as $n_t \rightarrow 0$. Row three left: n_t . Row three right: $p_t - p_t^{EM}$ with a solid black line a zero.

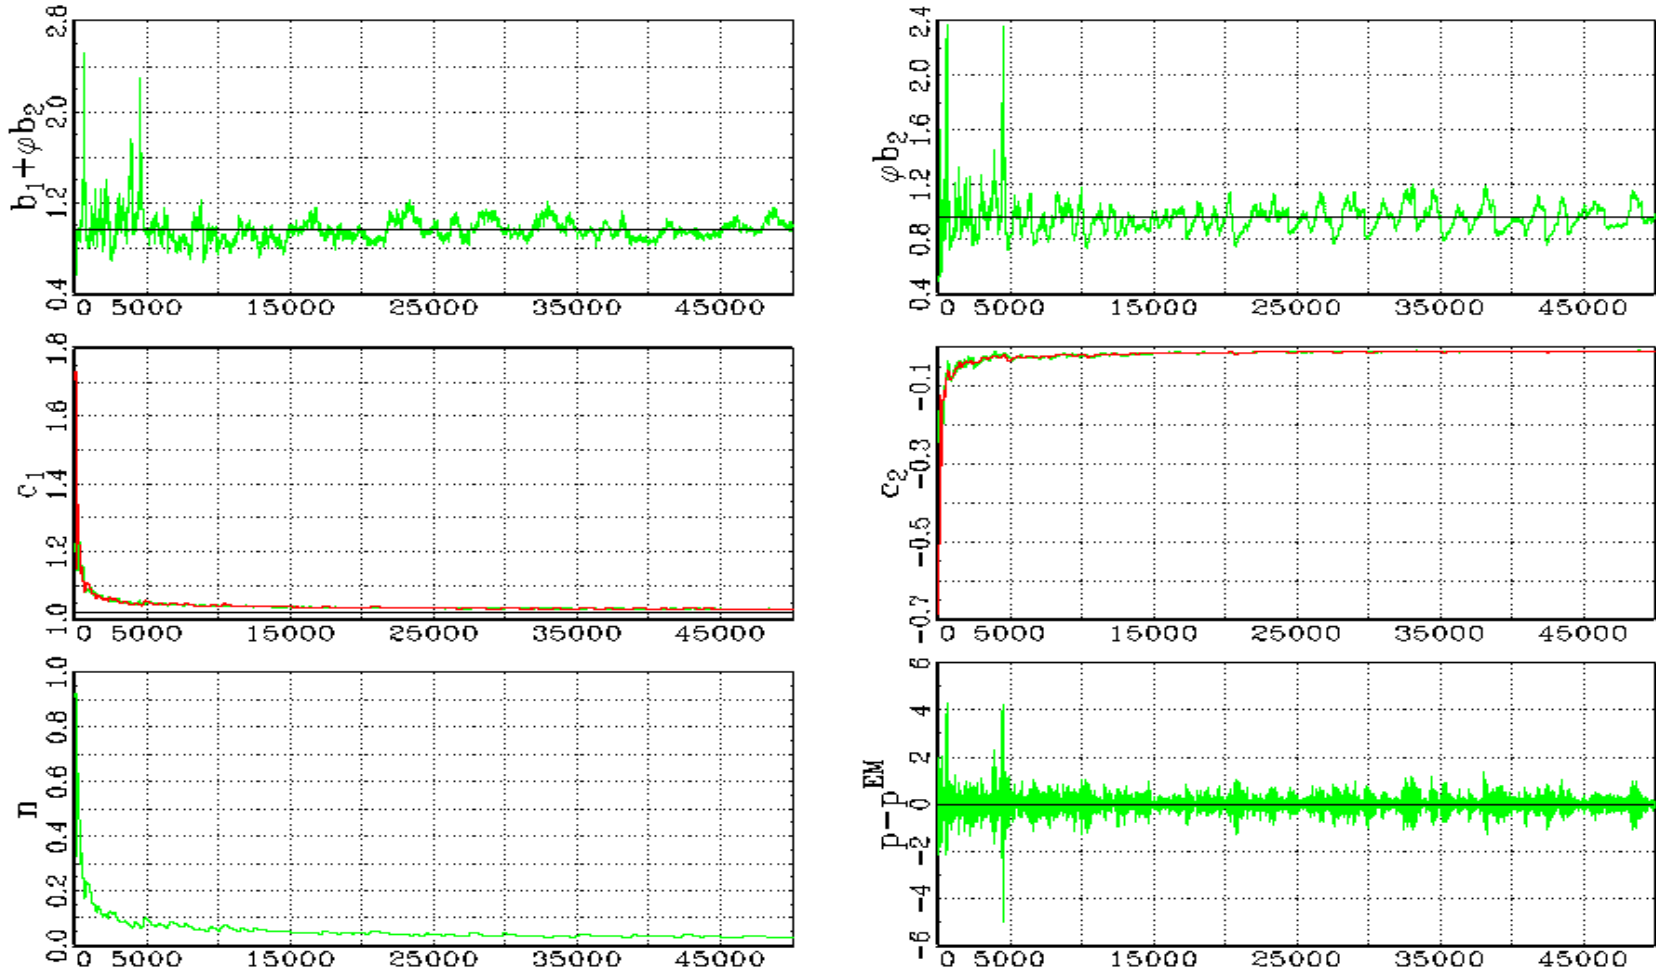


Figure 5: Evolution of endogenous parameters. RD dynamics, $\delta = 0.01$, end of sample ($t = 1.5 - 2.0 \times 10^6$). Row one left: $b_{1t} + \phi b_{2t}$ with the solid black line at $\phi/(R - \phi)$. $b_{1t} + \phi b_{2t} = \phi/(R - \phi)$ implies the regression traders have the correct model. Row one right: ϕb_{2t} with the solid black line at $\phi/(R - \phi)$. $\phi b_{2t} = \phi/(R - \phi)$ implies an efficient market. Row two left: c_{1t} (lighter shade or green) and $c_1^*(n_t)$ (darker shade or red). Row two right: c_{2t} (lighter shade or green) and $c_2^*(n_t)$ (darker shade or red). $c_1^*(n_t) \rightarrow R$ and $c_2^*(n_t) \rightarrow 0$ as $n_t \rightarrow 0$. Row three left: π . Row three right: $p_t - p_t^{EM}$ with a solid black line a zero.

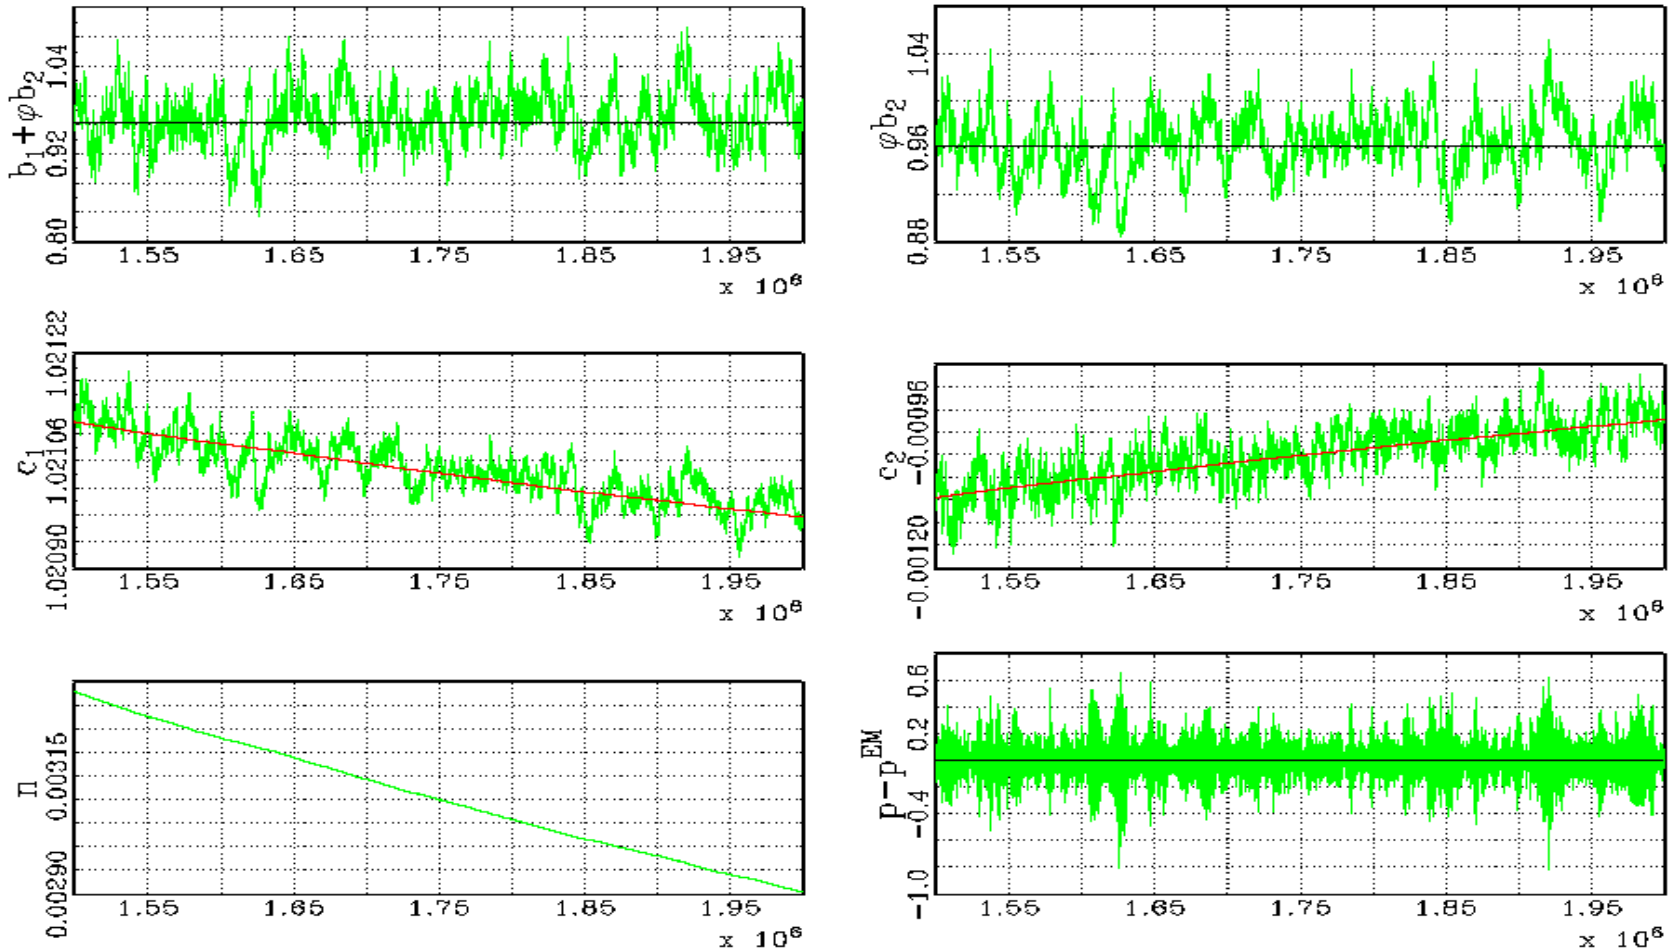


Figure 6: Evolution of endogenous parameters. RD dynamics, $\delta = 1$, end of sample ($t = 1.5 - 2.0 \times 10^6$). Row one left: $b_{1t} + \phi b_{2t}$ with the solid black line at $\phi/(R - \phi)$. $b_{1t} + \phi b_{2t} = \phi/(R - \phi)$ implies the regression traders have the correct model. Row one right: ϕb_{2t} with the solid black line at $\phi/(R - \phi)$. $\phi b_{2t} = \phi/(R - \phi)$ implies an efficient market. Row two left: c_{1t} (lighter shade or green) and $c_1^*(n_t)$ (darker shade or red). Row two right: c_{2t} (lighter shade or green) and $c_2^*(n_t)$ (darker shade or red). $c_1^*(n_t) \rightarrow R$ and $c_2^*(n_t) \rightarrow 0$ as $n_t \rightarrow 0$. Row three left: n_t . Row three right: $p_t - p_t^{EM}$ with a solid black line a zero.

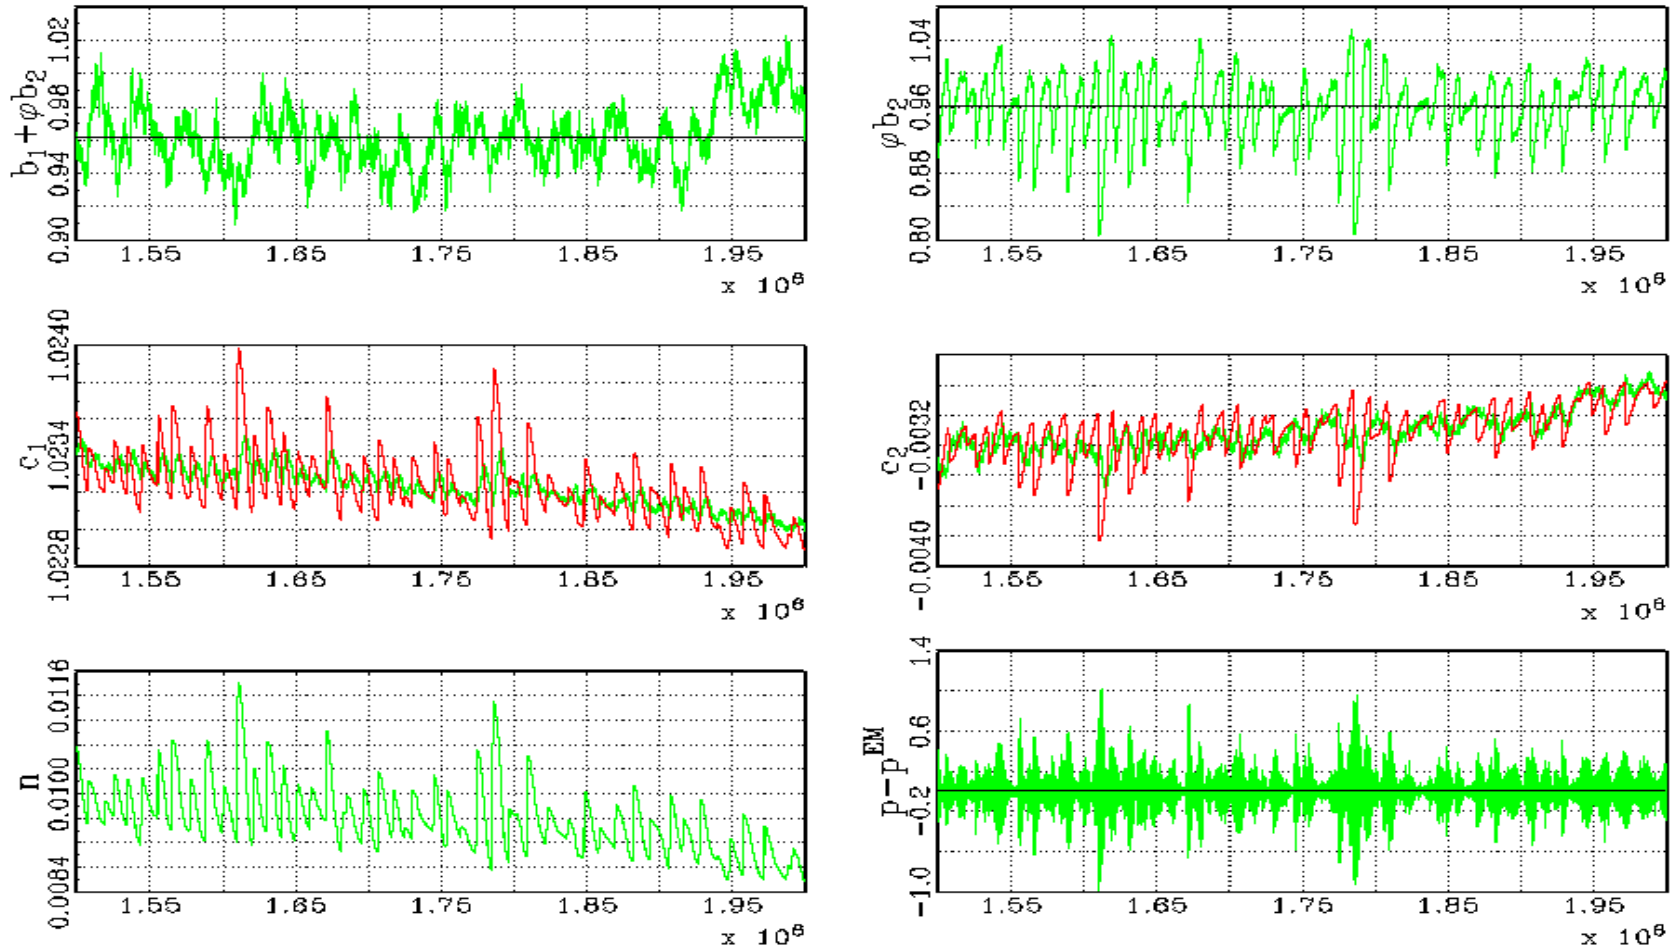


Figure 7: Evolution of endogenous parameters. Complete Information, beginning of sample ($t = 100 - 50,100$). Row one left: $b_{1t} + \phi b_{2t}$ with the solid black line at $\phi/(R - \phi)$. $b_{1t} + \phi b_{2t} = \phi/(R - \phi)$ implies the regression traders have the correct model. Row one right: ϕb_{2t} with the solid black line at $\phi/(R - \phi)$. $\phi b_{2t} = \phi/(R - \phi)$ implies an efficient market. Row two left: c_{1t} with a solid black line at R . Row two right: c_{2t} with a solid black line at zero. Row three left: c_{3t} . Row three right: $p_t - p_t^{EM}$ with a solid black line a zero.

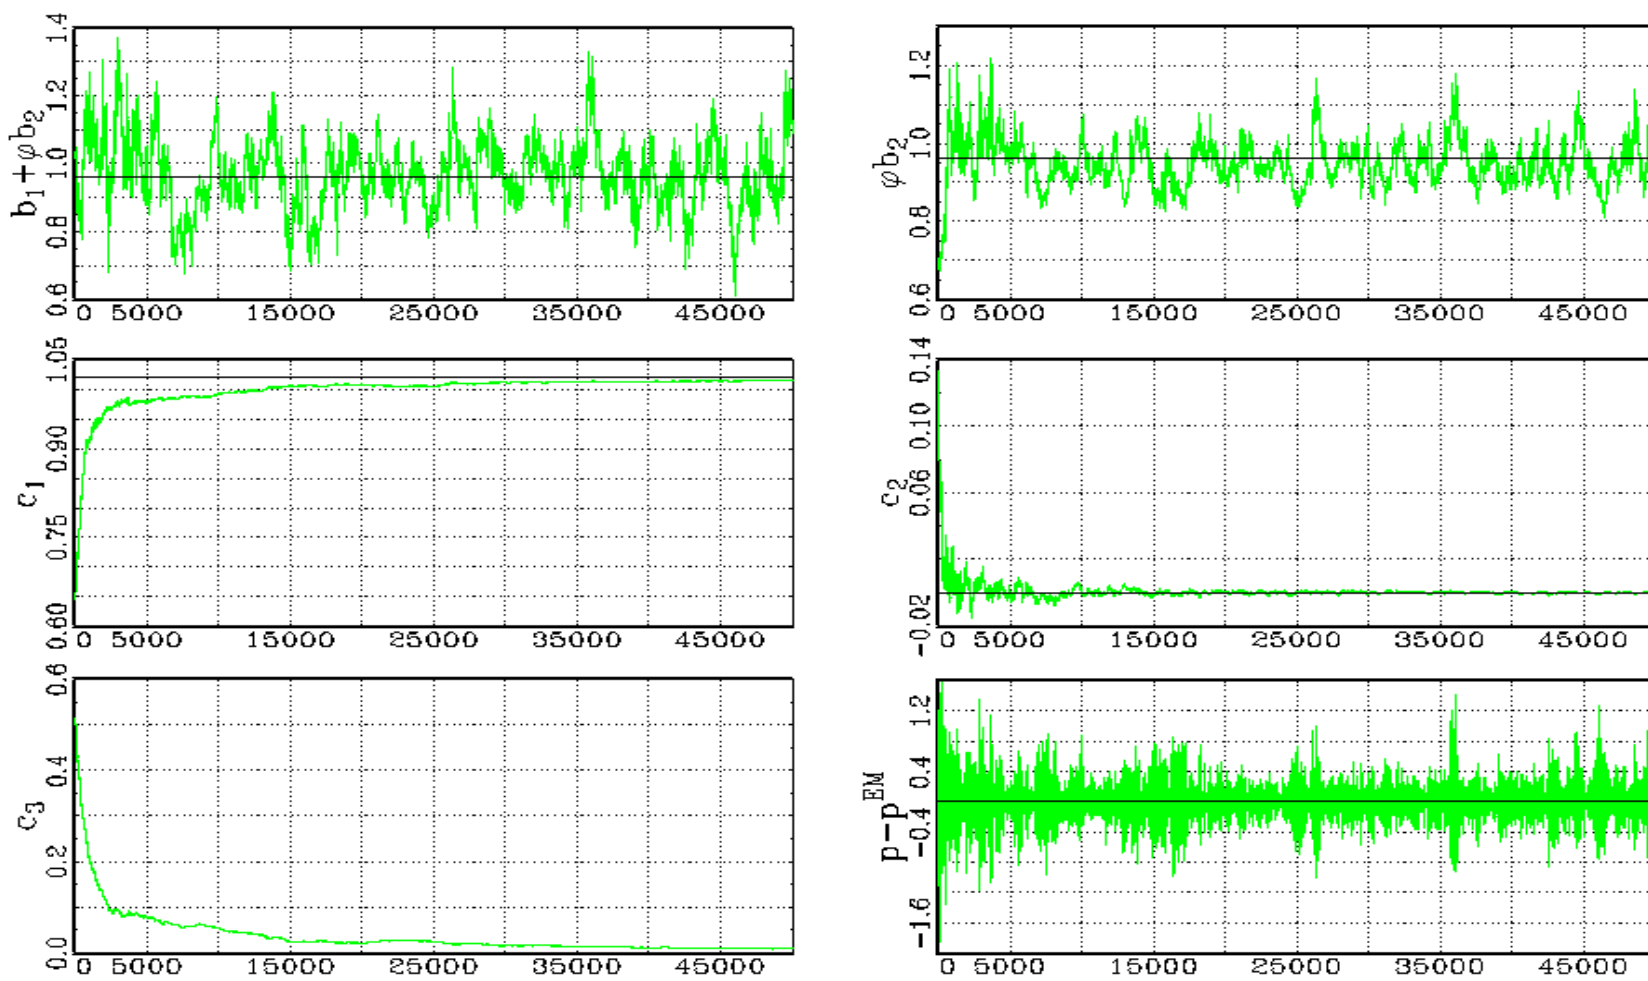


Figure 8: Evolution of endogenous parameters. Complete Information, end of sample ($t = 1.5 - 2.0 \times 10^6$). Row one left: $b_{1t} + \phi b_{2t}$ with the solid black line at $\phi/(R - \phi)$. $b_{1t} + \phi b_{2t} = \phi/(R - \phi)$ implies the regression traders have the correct model. Row one right: ϕb_{2t} with the solid black line at $\phi/(R - \phi)$. $\phi b_{2t} = \phi/(R - \phi)$ implies an efficient market. Row two left: c_{1t} with a solid black line at R . Row two right: c_{2t} with a solid black line at zero. Row three left: c_{3t} . Row three right: $p_t - p_t^{EM}$ with a solid black line a zero.

