

Market Efficiency and Learning in an Endogenously Unstable Environment^{*}

David Goldbaum[†]

An informationally inefficiency market is produced without an exogenous source of noise in the price. Fundamental traders acquire private information directly through research. Regression traders employ a learning process to extract the private fundamental information from the public price. The relative popularity between these two strategies evolves based on performance. The model converges towards adoption of regression analysis to the point of creating instability, endogenously producing a noisy price. The lack of a revealing price in the coupled learning and population processes reflects the Grossman and Stiglitz (1980) impossibility of informationally efficient markets.

Keywords: Efficient Markets, Least-Square Learning

JEL codes: G14, C62, D82

^{*} The author wishes to thank Spyros Skouras for helpful feedback on this project, and Doyne Farmer and the Santa Fe Institute for providing the environment under which this project originated. The author also appreciates the feedback received at the Meeting of the Society of Computational Economics and from two anonymous referees. Financial support was provided by the Rutgers Research Council.

[†] Department of Economics, Rutgers University, Newark, 360 Dr. MLK Jr. Blvd., Newark, NJ 07102, goldbaum@andromeda.rutgers.edu

Market Efficiency and Learning in an Endogenously Unstable Environment

1. Introduction

The noise trader is a construct that has become a mainstay of financial market modeling. Features of the rational expectations and efficient markets equilibria, such as no-trade and zero profit solutions, often stand in contrast to a market phenomena that a researcher is attempting to explain. Inserting noise traders into the market is the standard tool for disrupting the equilibria. Grossman and Stiglitz (1980) first introduced the noise trader to hamper uninformed traders using the price to free-ride on the efforts of informed traders engaging in fundamental research. This paper disrupts the efficient markets equilibria in a Grossman and Stiglitz (hereafter GS) based environment without resorting to noise traders.

The GS paradox of the impossibility of informationally efficient markets is created in a dynamic setting in which the asset is infinitely lived and has a dividend that follows a random walk. GS assumes two types of traders in fixed proportions. One group receives a common private signal on the value of a risky terminal asset. The other group is uninformed about the private information but attempts to extract it indirectly through observation of the market price. GS also assumes a rational expectations equilibrium in which the uninformed traders know the correct relationship between the observed price and the private information. GS find that there is no population proportion at which profits are equal. A population proportion at which traders are content with their strategy only exists if the price imperfectly reflects value, thus the impossibility of an informationally efficient market.

One of the introduced dynamics requires that the uninformed traders learn the relationship between price and value. The analysis shows that the rational expectations

equilibrium (REE) to the modeled learning process exists, but it can be unstable. Another dynamic process allows traders to switch between the fundamental and market-based information producing a fluid population proportion. The two dynamic processes ensure that the market remains in the region producing the unstable REE. Simulations characterize the behavior of the market in the unstable region. The instability of the REE endogenously produces a noisy non-revealing price.

A number of useful precedents have been established for modeling a dynamic version of GS. Bray (1982) and Routledge (1999) create a multiperiod version of GS with Bray employing IID terminal values at the end of each period and Routledge using an asset with an IID dividend process. Bray (1982) proposes a least-squares learning process for a fixed population of uninformed traders learning to extract information from the price, and examines its convergence properties. In general, Bray finds convergence in the learning process to the correct model. The learning traders' strategy matches that employed by the GS rational uninformed traders. The convergence occurs with and without a random supply. Bray's analysis supports the GS assertion that traders can learn the relationship between price and payoff through observation and that the assumption of full rationality on the part of the uninformed is reasonable.

In Routledge (1999), the dual acts of learning and of choosing between being informed and uninformed are embodied in the single act of imitation. Modeled with noise introduced through a random supply, a long-run stable REE exists. The process generally converges to the REE.

For this paper, the learning process is based on least-squares learning in the nature of Marcet and Sargent (1989a, 1989b). Both the Hussman (1992) and Timmerman (1996)

papers employ least-squares learning in a financial market setting, but differ from this examination by considering a stationary AR(1) dividend process.

The population dynamic in the choice of information is based on replicator dynamics. The replicator dynamic has its origins in the study of population evolution based on reproduction as employed by Taylor and Jonker (1978). Adaptations to learning in a game setting include Börgers and Sarin (1997), Cheung and Friedman (1998), Hopkins (2002), and Droste *et al* (2002). For this paper, replicator dynamics are used to describe the dynamics of a population choice. Sethi and Franke (1995) and Branch and McGough (2003) employ the replicator dynamics for the same purpose. Within a population governed by replicator dynamics, strategies that produce superior payoffs attract adherents away from inferior strategies. An interior fixed point exists at a proportion producing equal payoffs. At the boundaries, an inferior strategy has been driven out by a strictly superior strategy.

Other relevant papers that examined the information content of market-based information include Grundy and McNichols (1989), Brown and Jennings (1989), Blume, Easley and O'Hara (1994), and de Fontnouvelle (2000). Brock and LeBaron (1996), and Brock and Hommes (1998) provide important foundations for considering population dynamics in a financial market setting.

Section 2 establishes the market and the participants. Section 3 finds the solution for the fixed point to the learning process. Instability in the fixed point provides the basis for the endogenously noisy price. The analysis of the population process demonstrates that if the market is constrained to the fixed point of the learning process then there is no fixed point to the population process, reflecting the same paradox highlighted by GS. Simulations examined the interaction between learning and the population processes. In Section 4 the

division in the trader population is removed as traders are allowed to employ fundamental and market information simultaneously rather than having to choose between the two.

Section 5 concludes.

2. Model

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 market determined price p_{t+1} . The market participants are aware that the stochastic dividend follows a random walk:

$$d_{t+1} = d_t + \varepsilon_{t+1}, \text{ with } \varepsilon_{t+1} \sim IIDN(0, \sigma_\varepsilon^2). \quad (1)$$

While a natural choice, modeling dividends as a random walk is an important departure from the standard multiperiod adaptations of the GS model that tend to use an IID end of period payoff, IID dividends, or a stationary AR process to model dividends over time. The traders' accommodation of the nonstationarity affects the stability of the fixed point under learning, as the developed model reveals.

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 now familiar demand for the risky asset,

$$q_{it}(p_t) = \theta_{it}(E_{it}(z_{t+1}) - Rp_t) \quad (2)$$

with $z_t = p_t + d_t$, $\theta_{it} = 1/\gamma\sigma_{it}^2$, $\sigma_{it}^2 = \text{var}_{it}(z_{t+1})$ indicating the conditional variance, and γ is the coefficient of absolute risk aversion.

Assume K strategies for estimating payoffs. In a Walrasian equilibrium, the market

price equates supply and demand for the asset. Without loss of generality, set fixed net supply of the risky asset to zero.³ 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

For the convenience of discussion and analysis, consider a division of information into three non-overlapping categories: fundamental information, market information, and public information. Fundamental information consists of exogenous information that is useful for estimating the value of an asset. Fundamental information is not generally observable but traders can access the information through research. In the context of the model, fundamental analysis provides the individual trader with a noisy private signal of the value of next period's dividend.

Market information is based on endogenous market generated data. In the context of this model, the current and past prices are the only pertinent market information, but it could also include other market generated information such as volume and volatility.

Public information is exogenous to the model and observable to all market participants. In the context of this model the public information includes current and past dividends.

In Sections 2 and 3 traders supplement public information with either fundamental or

³ $S > 0$ lowers the market clearing price by a constant term that is a function of the risk aversion parameter and the conditional variance terms. $S = 0$ removes the impact of the risk aversion coefficient from price determination, but does not change the model dynamics.

market information, but not both. A common argument in support of this division is to claim constraints on resources or bounded rationality on the part of the agents (see Hong and Stein; 1999, and Barberis and Shleifer; 2003). The constrained traders are able to process one of the two information sources in a given period. An argument can also be made in support of the division on philosophical grounds. The traders are not so much constrained or bounded to using just one type of information, but rather decide which of two mutually exclusive states, one favoring fundamental research and the other favoring the market information, currently describes the market. Such a characterization of the market will, in fact, be loosely consistent with the asymptotic behavior of the model. Finally, the division is also convenient for direct comparison to GS.

In GS, the market information is weakly inferior to fundamental information. The objective of the uninformed traders is to use the market information to obtain the knowledge possessed by the informed traders without having to pay the associated cost of acquisition. In Brock and Hommes (1998), the traders employing market-based inferior information benefit from the (near) zero profit condition at (or near) the REE produced by the superior perfect foresight traders, again benefiting from the savings of not performing costly research.

The noisy private signal used in the present model creates the possibility that market information is superior to fundamental research, a situation that does not arise in GS or Brock and Hommes (1998). The fact that the market provides superior information at (and near) the REE relative to fundamental research removes the necessity of including a cost on conducting research to induce traders to employ market data. Thus, while it is intuitively reasonable to model fundamental research as a costly process, imposing such a cost serves no purpose and does not alter the characteristics of the solution. In practice, many traders who

employ market-based analysis do so based on the belief that it offers superior returns, and not necessarily to reduce the cost of analysis.

2.2.1 The Fundamental Trader

The estimate of the future payoff is in the nature of Hellwig (1980). Trader i 's private research indicates the fundamental value of the risky security captured by a signal based on d_{t+1} that is subject to idiosyncratic error,

$$s_{it} = d_{t+1} + e_{it}, \text{ with } e_{it} \sim IIDN(0, \sigma_e^2). \quad (4)$$

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

$$E_i(d_{t+1} | d_t, s_{it}) = (1-\beta)d_t + \beta s_{it} \quad (5)$$

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

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

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

$$p_t^F = b_1^F d_t + b_2^F d_{t+1} + v_t, \quad (6)$$

$v_t = \frac{\beta}{R-1} \frac{1}{N} \sum_i e_{it}$. 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 $b_1^F = \frac{1-\beta}{R-1}$ and $b_2^F = \frac{\beta}{R-1}$.

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.⁴ The price reflects the next period's dividend, d_{t+1} , because it is a common component of each trader's private information. The price fails to fully reflect the private information ($\beta < 1$) due to the limited extent that each individual trader incorporates his private signal into his own demand function.

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,

$$q_t^F = \theta_t^F \left(\frac{R}{R-1} ((1-\beta)d_t + \beta d_{t+1}) - R p_t \right). \quad (7)$$

2.2.2 Regression traders

The regression traders model the relationship between the payoff, $z_t = p_t + d_t$, and the market observables, the current and past prices and dividends. To accommodate the random walk in dividends the analysis is based on innovations rather than levels. Let Δ indicate the difference operator such that $\Delta z_t = z_t - z_{t-1}$. A derivation demonstrating that the traders appropriately estimate

$$\Delta z_t = c_1 \Delta p_{t-1} + c_2 \Delta d_{t-1} + \zeta_t \quad (8)$$

is included in the appendix.

The traders employ least-squares learning to update the parameters of their model. The learning process is self-referential with an endogenous state variable, Δp_{t-1} , included as a regressor. Let $x_t' = [\Delta p_t \ \Delta d_t]$. The regression traders update the coefficients, $C_t = [c_{1t} \ c_{2t}]$, using the standard recursive updating algorithm for least-squares learning:

$$C_t = C_{t-1} + (Q_t^{-1} x_{t-2} (\Delta z_{t-1} - C_{t-1} x_{t-2}))' / t, \quad (9)$$

⁴ Formally, v_t is $o(1)$.

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

given (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

$$q_t^R = \theta_t^R (C_t x_t + (p_t + d_t) - R p_t). \quad (10)$$

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

$$p_t = b_{1t} d_t + b_{2t} d_{t+1} + b_{3t} p_{t-1} + b_{4t} \Delta d_t. \quad (12)$$

See (A.4) in the appendix for the coefficients of (12) expressed as a function of the agents' demand parameters.

Through repeated substitution of the lagged price, p_t can be expressed as a function of the exogenous dividend stream:

$$p_t = B_{1t} d_t + B_{2t} d_{t+1} + B_{3t} d_{t-1} + \sum_{k=1}^t \left(\prod_{j=1}^k b_{3,t-j-1} \right) (B_{3,t-k}) d_{t-k-1}. \quad (13)$$

See (A.5) in the appendix for the coefficients B_{1t} through B_{3t} as functions of b_{1t} through b_{4t} .

3. Analysis and simulation

3.1 Learning under fixed n

Consider a fixed $n_t = n$ for all t . 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).

This section develops the solution for the fixed point to the learning process and explores issues related to its stability. Simulations are employed to characterize the convergence properties of the learning model where the fixed point is unstable. The failure of the learning process to converge is the basis for the endogenously noisy price.

3.1.1 A fixed point solution

At the fixed point to the learning process the regression trader's perceived relationship captured by equation (8) is consistent with the actual relationship described by (1) and (12). Employing L to indicate the lag operator, from (13)

$$p_t = B_1(n)d_t + B_2(n)d_{t+1} + B_3(n)(1 - b_3(n)L)^{-1}d_{t-1}. \quad (14)$$

See (A.6) in the appendix for the coefficients B_1 through B_3 as functions of b_1 through b_4 .

Use the notation $FP(n) = \{c_1(n), c_2(n), b_1(n), b_2(n), b_3(n), b_4(n), B_1(n), B_2(n), B_3(n)\}$ to denote the fixed point values for the endogenous regression coefficients and price parameters. For notational convenience, the dependence of the fixed point values on n will be implied and only included when needed for emphasis or clarification.

For $B_3 \neq 0$, the presence of the regression traders introduces lagged dependence in the price that is not present in the absence of the regression traders. For convenience, establish **Condition A.** $B_3 = 0$.

Proposition 1 below characterizes the fixed point solution to the learning process. For consistency with GS, the fixed point will be a REE if the regression traders correctly extract the private information available from the price.

Proposition 1. *Under Condition A, a fixed point to the learning process exists for $0 \leq n \leq 1$.*

(a) *The fixed point is a REE.*

(b) *The parameters $FP(n)$ are smooth continuous functions of n for $0 < n \leq 1$.*

(c) *There is a discrete jump in the fixed point parameter values at $n = 0$.*

Part (a) of Proposition 1 is a feature of the least-squared learning process. At the fixed point the regression traders have correctly identified the model, which allows them to extract d_{t+1} from the observation of x_t . Parts (b) and (c) follow directly from the fixed point coefficient solutions (A.7) and (A.11) in the appendix. From (A.7), as $n \rightarrow 0$, $B_1 \rightarrow 0$ and $B_2 \rightarrow 1/(R-1)$. From (A.11), for $n = 0$, $B_1 = 1/(R-1)$, and $B_2 = 0$. Thus, as in GS, at the fixed point, the regression traders are rational and there is a discrete jump in the price solution at $n = 0$. As $n \rightarrow 0$, p_t converges to fully reflect d_{t+1} but at $n = 0$ p_t can only reflect the public information up to time t .

3.1.2 Stability of the REE fixed point

Condition A need not hold when the market deviates from the fixed point, thus introducing lagged dividend terms into the price equation. A bounded price according to (14) requires that the impact of the past dividends be diminishing, $-1 < b_3(n) < 1$. The fixed point to a least-squares learning process is unstable if the regression coefficients in the neighborhood of the fixed point solutions fail to converge. Stationarity in the transition equations is a requirement for convergence in learning to the fixed point (see Marcet and Sargent; 1989b). If $b_3(n)$ is outside the unit interval, then price innovations become nonstationary so the regression traders' $E(x_t x_t')$ is no longer well defined and the process fails to converge.

Proposition 2. *There exists a critical \underline{n} , $0 < \underline{n} < 1$, for which $n > \underline{n}$ produces a stable REE fixed point. For $n \leq \underline{n}$, the REE fixed point is unstable.*

Proof. From (A.10), $b_3(n)$ is a continuous smooth monotonically increasing function of n with $b_3(n) \rightarrow -\infty$ as $n \rightarrow 0$ and $b_3(1) = 0$. Let \underline{n} indicate the value of n such that $b_3(\underline{n}) = -1$. For $n \leq \underline{n}$, b_3 is on or outside the unit interval, violating a requirement for convergence in least-squares learning.

The lack of stability in the fixed point at lower n results from the presence of the lagged price in (12). Were the dividend process stationary, the procedure to forecasts $E(z_{t+1})$ from the observed data would involve a regression based on levels rather than in differences. In this case, the fixed point would be stable for all values $0 < n \leq 1$, consistent with the convergence found by Bray (1982).

3.1.3 Simulation of the least-squares learning process

Beliefs deviate from the actual model governing the system of price determination during the learning process. Equations (A.8) and (A.9) provide the variance terms for the forecast errors at the fixed point, but the fixed point solution is not known by the agents, nor are they correct when the model deviates from the fixed point. In simulation, the traders need to estimate the conditional variance of the error associated with each strategy. The traders use estimates of the forecast errors based on observed data. Allow the traders to recursively estimate σ_F^2 and σ_R^2 using the least-squares algorithm to update. The aggregate estimates are

$$\hat{\sigma}_{Ft}^2 = \hat{\sigma}_{Ft-1}^2 + ((z_t - z_t^F)^2 - \hat{\sigma}_{Ft-1}^2)/t + \xi/t, \quad (15)$$

$$z_t^F = \frac{R}{R-1}((1-\beta)d_t + \beta d_{t+1}),$$

$$\xi = \left(\frac{R\beta}{R-1}\right)^2 \sigma_e^2 \text{ and}$$

$$\hat{\sigma}_{Rt}^2 = \hat{\sigma}_{Rt-1}^2 + ((z_t - z_t^R)^2 - \hat{\sigma}_{Rt-1}^2) / t \quad (16)$$

$$z_t^R = C_t x_t + z_{t-1}$$

with σ_{F0}^2 and σ_{R0}^2 given. The first adjustment term in (15) is the estimate of the group level error of the fundamental signal. The ξ/t term captures the individual fundamental trader's additional uncertainty due to the idiosyncratic noise associated with the private signal.

The solid line in each frame of Figure 1 plots the REE fixed point value of the endogenous parameter as a function of n . Each point represents simulation output. The simulation data is based on 20 independent draws of a $T = 10,000$ period dividend stream. For each dividend stream, the simulation is run with n ranging from $n = 0.01$ to $n = 1$ in increments of 0.01. The exogenous parameters are set to $R = 1.02$, $\gamma = 1$, and $\sigma_\varepsilon = \sigma_e = 1$ so that $\beta = 0.5$. The chosen parameters produce $\underline{n} = 0.25$. The time $t = 0$ values of the endogenous parameters are the fixed point values given n . The end of sample values of the endogenous parameters are used to compute B_{1T} , B_{2T} , and B_{3T} , according to the appendix equation (A.6). Figure 1 plots B_{1T} , B_{2T} , B_{3T} , b_{3T} , c_{1T} , and c_{2T} against n .

[Figure 1 about here]

Consistent with Proposition 2, for $n > \underline{n}$ the model converges to the REE fixed point but for $n < \underline{n}$ the system fails to penetrate into the region $b_{3t} < -1$. Instead, the process remains trapped in the neighborhood containing $b_3 = -1$, unable to continue its convergence.

There are no constraints placed on any of the model parameters. Whenever b_{3t} drops below -1, the system becomes unstable. During the learning process, the error in the regression traders' coefficients introduces error into the price. Thus for b_{3t} inside the unit interval, the errors dampen, allowing the learning process to converge toward correct beliefs. Once b_{3t} drops below -1, pricing errors become explosive. With the regression traders

themselves introducing error, the learning process cannot continue. Figure 2 plots a 1,000 period sample of the time series of b_{3t} , c_{1t} , and $p_t - p_t^F$ for $n = 0.05 < \underline{n} = 0.25$. The value of b_{3t} oscillates across the boundary separating the stable and unstable price states. For $b_{3t} \leq -1$, the nonstationarity in the price reduces the information content of the price innovations. The new observations reduce c_{1t} , causing b_{3t} to increase, moving the model back into the stable region.

[Figure 2 about here]

Let b_3^* reflect the asymptotic value of b_{3t} as a function of n ,

$$b_3^*(n) = \begin{cases} b_3, n \geq \underline{n} \\ -1, n < \underline{n} \end{cases}.$$

For the given n , b_3^* defines the basin of attraction for the system. For $n > \underline{n}$, the basin contains the fixed point as its attractor. For $n < \underline{n}$, the basin does not contain the fixed point for the system. Instead, $b_{3t} = -1$ acts as an attractor though it is not a fixed point. The deviation of b_{3t} in simulation from b_3 for $n < \underline{n}$ causes the remaining parameters to deviate from the REE fixed point solution as well. Refer to the C^* , b^* and B^* associated with $b_3 = b_3^*$ as the attractor of the model.

According to (A.7), the fixed point values of $\{B_1, B_2\}$ converge towards $\{0, 1/(R-1)\}$ as $n \rightarrow 0$. Alternately, Figure 1 shows that as $n \rightarrow 0$ $\{B_1^*, B_2^*\} \rightarrow \{1/(R-1), 0\} = \{B_1(0), B_2(0)\}$. Thus, as $n \rightarrow 0$, the parameters that serve as the attractor for the system converge to reflect the $n = 0$ solution, producing a smooth transition for the price coefficients at $n = 0$ rather than the discrete jump produced by the REE solution. This eliminates the discontinuity found in the REE fixed point.

3.1.4 Profits

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:

$$\pi_t^k = q_t^k (p_{t+1} + d_{t+1} - Rp_t), k = F, R. \quad (17)$$

Proposition 3. *At the REE fixed point*

a) *For $0 < n \leq 1$, $E(\pi^F) < E(\pi^R)$.*

b) *For $n = 0$, $E(\pi^F) > E(\pi^R)$.*

Proof. From (A.12), $E(\pi^R) > 0$ for $n > 0$. Substitute B_2 from (A.7) into (A.12) to obtain $E(\pi^F) = 0$ at $n = 1$ and $E(\pi^F) < 0$ for $0 < n < 1$.⁵ Part (b) follows directly from (A.13).

Proposition 3 embodies the GS impossibility of informationally efficient markets. The success of the regression traders for $n > 0$ results from the fact that at the REE fixed point the learned regression parameters correctly extract d_{t+1} . At the REE the market filters out the idiosyncratic noise revealing the underlying value, providing better information than the noisy private signal obtained through fundamental analysis. For $n = 0$, as with GS, in the absence of any trader conducting research, the price has nothing to reveal.

Expected profit are a function of n and of the regression coefficient c_{1t} from (8),

$D(n, c_{1t}) = E(\pi^F(n, c_{1t})) - E(\pi^R(n, c_{1t}))$.⁶ At the REE fixed point $c_{1t} = c_1(n)$, allowing a reduced form $D(n) = E(\pi^F(n)) - E(\pi^R(n))$. For $0 < n \leq 1$, $D(n) < 0$ and $n = 0$ produces $D(n) > 0$. Regression traders' beliefs are in error if $c_{1t} \neq c_1(n)$. Proposition 4 below characterizes the impact of this error on $D(n, c_{1t})$. Figure 3 includes the boundaries discussed in

⁵ Figure 4 below includes a plot of the REE profits.

⁶ As discussed in the appendix, since the second variable in (8) is only present to control for the spurious noise produced by the ε_t term in the price innovation, c_{1t} is the heart of the regression traders' forecast of future value.

Proposition 4.

[Figure 3 about here]

Proposition 4. *Assume Condition A, $B_1(c_{1t}) + B_2(c_{1t}) = 1/(R-1)$, and $c_{2t} = -B_2(c_{1t}) * c_{1t}$, then for*

$0 < n \leq 1$ there exists $c_1^- = (R-1) < c_1(n) < c_1^+(n) < \bar{c}_1(n)$ such that

- a) $c_1^- < c_{1t} < c_1^+(n)$ produces $D(n, c_{1t}) < 0$,*
- b) the boundaries $c_{1t} = c_1^-$ or $c_{1t} = c_1^+(n)$ produce $D(n, c_{1t}) = 0$, and*
- c) $c_{1t} < c_1^-$ or $c_1^+(n) < c_{1t} < \bar{c}_1(n)$ produces $D(n, c_{1t}) > 0$.*
- d) There is an upper bound on the value of c_{1t} such that $D(n, c_{1t}) \rightarrow \infty$ as $c_{1t} \rightarrow \bar{c}_1(n)$ from below.*

The proof for Proposition 4 is included in the appendix. The three assumptions limit pricing error to reflect an incorrect estimate of the relationship between price and payoff. *Condition A* ensures that lagged dividends do not introduce bias. The condition $B_1(c_{1t}) + B_2(c_{1t}) = 1/(R-1)$ ensures the price reflects a proper discounting of the traders' expected future dividends (even if the expectation is in error). The condition $c_{2t}(c_{1t}) = -B_2(c_{1t})c_{1t}$ ensures that the second term in (8) correctly cancels out the spurious ε_t component of the price innovation. These three conditions are all features of the fixed point solution. Of course, deviations from the fixed point may include violations of these assumptions, but the behavior of the simulations seems to be approximately governed by the guidelines found here.

At c_1^- the regression traders share the belief of the average fundamental trader despite the different information source. Net holdings for each group are zero. For $c_{1t} = c_1^+$, the resulting market price reflects fundamentals based on d_{t+1} so that no profits are realized despite differences in beliefs and holdings. Between these two boundaries, the regression

traders have sufficiently accurate forecast to generate profits so that $D(n, c_{1t}) < 0$. Outside these two boundaries, the regression traders' error produces profits for the fundamental analysis, $D(n, c_{1t}) > 0$ with $D(n, c_{1t}) \rightarrow \infty$ as $c_{1t} \rightarrow \bar{c}_1(n)$. For $c_{1t} \geq \bar{c}_1(n)$, the regression traders have so overestimated the dividend innovation that their positive response to the price outweighs the negative response from the demand of the fundamental traders, creating an upward sloping aggregate demand curve for the model and an infinite price. In simulation regression trader demand is curtailed by the inclusion of the current error in the variance estimate. This is reasonable since in the dynamic setting, the $D(n, c_{1t}) \rightarrow \infty$ ensures that n_t increases before \bar{c}_1 is encountered. Both \bar{c}_1 and c_1^+ converge towards infinity as $n \rightarrow 1$.

Figure 4 plots the realized average profits from the same simulations used to create Figure 1. Each point in Figure 4 is the average of the profits realized in the last 1,000 periods of the simulation. The “o” points are the profits realized by the fundamental traders. The “+” points are the profits realized by the regression traders. The solid lines are the expected profits earn at the REE fixed point.

[Figure 4 about here]

As displayed in Figure 4, the realized simulated profits diverge from those implied by the REE solution for $n < \underline{n}$. Let n^{**} indicate the fixed point for n at which $D(n, c_1) = 0$. Median profits as plotted in Figure 4 suggest an n^{**} producing $D(n^{**}, C^*(n^{**})) = 0$ at approximately $n^{**} = 0.12$.

3.2 Population Dynamics

Allow for a dynamic population proportion that responds to the difference in the traders' expected profits for the two information approaches. For Section 3.2 remove the learning dynamics of (9) from the system and introduce a dynamic population process ((18)

below). The population reflects expectations of profits, which evolve according to (20). The analysis demonstrates that if the market is constrained to the REE of the learning process there is no fixed point to the population process, reflecting the same paradox highlighted by GS.

3.2.1 Replicator dynamics

Sethi and Franke (1995) propose an evolutionary dynamic population in which "each period the group performing better in the recent past wins some converts" (1995, p584). Allow that traders base their selection of strategy on a relative measure of the expected profits of the two options. A two choice version of the more general K choice replicator dynamic model of Branch and McGough (2003) results in the transition equation

$$n_{t+1} = \begin{cases} n_t + r(\pi_t^{Fe} - \pi_t^{Re})(1 - n_t) & \text{for } \pi_t^{Fe} \geq \pi_t^{Re} \\ n_t + r(\pi_t^{Fe} - \pi_t^{Re})n_t & \text{for } \pi_t^{Fe} < \pi_t^{Re} \end{cases} \quad (18)$$

Here, π_t^{Fe} and π_t^{Re} are the traders' estimates of the profits earned by the fundamental and market-based approaches, respectively. Their computation will be discussed below. A number of different functional forms for r exist in the literature sharing the features, $-1 \leq r(x) \leq 1$, $r'(x) > 0$, and $r(0) = 0$. The increasing function ensures that the population flows towards the perceived superior strategy. Passing through the origin ensures $n_{t+1} - n_t = 0$ if

$\pi_t^{Fe} = \pi_t^{Re}$.⁸ The simulations that follow are based upon

$$r(x) = \tanh(\rho x / 2). \quad (19)$$

⁷ Formally, as a population proportion, n_t is restricted to be an element of the set of rational numbers. Allow that (18) produces a near approximation of n_t .

⁸ This is consistent with how GS envisions evolution in the population. "If the [expected utility of the informed] is greater than the [expected utility of the uninformed], some individuals switch from being uninformed to being informed (and conversely). An overall equilibrium requires the two to have the same expected utility." p394.

The results discussed in this paper should be robust to the selection of r .⁹

The traders rely on observed profits to indicated the expected profits for each strategy,

$$\pi_t^{ke} = \pi_{t-1}^{ke} + \mu_t (\pi_{t-1}^k - \pi_{t-1}^{ke}) \quad (20)$$

$$\pi_0^k = 0, k = F, R.$$

In its reduced form, (18) implies a mapping $n_{t+1} = f(n_t)$. Local stability requires $f'(n_t)|_{n^{**}} < 1$. This can be accomplished asymptotically by employing a sufficiently long memory in the computation of π^{ke} . Setting $\mu_t = 1/t$ is the equivalent of producing the least-squares estimate of the mean by placing equal weights on observed profits. A degenerative function for μ_t such as $\mu_t = 1/t$ ensures $f'(n_t)|_{n^{**}} \rightarrow 0$ as $t \rightarrow \infty$.

3.2.2 Convergence in the population

Consider the evolution of the population under three behavior settings for the regression traders: c_{1t} fixed, $c_{1t} = c_1(n_t)$, and $c_{1t} = c_1^*(n_t)$. It may be useful to refer to Figure 3. For fixed $c_{1t} = c > R-1$, n_t converges to place the system on c_1^+ at which profits are equal. If fixed $c_{1t} = c_1^- = R-1$, then profits are equal regardless of the value of n_t . A fixed $c_{1t} = c < R-1$ leads to a convergence of the system to $n_t = +1$ since the error is sufficient to produce positive profits for fundamental analysis regardless of n_t .

For $c_{1t} = c_1(n_t)$ the regression traders are assigned the correct n_t dependent regression coefficients. The market is constrained to be at a point on the curve labeled $c_1(n)$ in Figure 3. With the regression traders earning superior profits at all points along $c_1(n)$, n_t converges towards the origin. The discontinuity arises at $n_t = 0$ at which the fundamental analysis

⁹ Also examined was $r(x) = \min(e^{px/2} - 1, 1)$.

outperforms the regression analysis.

For $c_{1t} = c_1^*(n_t)$ the market travels down along the $c_1(n)$ curve for $n_t > \underline{n}$. For $n_t < \underline{n}$, Figure 1 reveals that $c_1^*(n)$ is biased downward relative to $c_1(n)$. The bias results in $c_1^*(n)$ crossing c_1^- at $n = 0.02$, imply that $n_t = n^{**}$ at $n_t = 0.02$ produces zero profits along $c_1^*(n_t)$. This conflicts with the earlier finding of $\pi^{Fe}(n) = \pi^{Re}(n)$ at $n = 0.12$ despite both finding stemming from the same set of simulation output. The difference suggests that violations of the assumption of Proposition 4 benefit the fundamental analysis. The violations hampers the regression traders' information extraction such that they fail to earn a profit despite a relatively small coefficient error. Since simulations suggest that $B_1(c_1^*) + B_2(c_1^*) = 1/(R-1)$ and $c_2^*(c_1^*) = -B_2(c_1^*)c_1^*$ both hold, the realization of $B_3 \neq 0$ is the likely culprit.

3.3 Simulation

Interactions between the learning and the population processes require examination of the full dynamic system by simulation. The full dynamic system is captured by the exogenous dividend process, (1), the endogenous price process, (12), the coefficients of which are determined by the regression coefficients from (8) that evolve following (9), and by the population proportion that follows (18) based on the expected profits following (20).

3.3.1 Convergence

The first two columns of data in Table 1 report summary statistics on the asymptotic values of the endogenous parameters, computing means and standard deviations based on the last 10,000 periods of a representative $T = 200,000$ simulation. Parameters are set to $n_0 = 0.75$, $\rho = 0.1$, and $\mu_t = 1/t$. The simulation converges with $n_t \rightarrow n^{**} = 0.12$. The parameters C_t , b_t , and B_t , converge to a neighborhood containing $C^*(n^{**})$, $b^*(n^{**})$, and

$B^*(n^{**})$.

[Table 1 about here]

Figure 5 displays the convergence of the first 20,000 periods of simulation output in a phase space based on n_t plotted horizontally and deviations from $c_1(n)$ plotted vertically. The asymptotic pair $(n^{**}, c_1^*(n^{**})-c_1)$ based on the data presented in Table 1 is indicated by the arrow pointing to the circle centered at (0.1236,-0.006).

[Figure 5 about here]

3.3.2 Constant gains in expected profits

A sufficiently high constant gains in the updating of π_t^{ke} may produce $f'(n_t)|_{n^{**}} \geq 1$, generating local instability in the population dynamic. The values in the second pair of columns of Table 1 derive from simulations with $\mu = 1$ implying adaptive expectations based on the most recently realized profits, $\pi_t^{ke} = \pi_{t-1}^k$. For $B_1 + B_2 = 1/(R-1)$, the realization of a greater B_{1T} as reported in Table 1 reflects the greater reliance on the fundamental information in the more volatile $\mu = 1$ simulation.

As can be seen in the plot of the endogenous parameters and price deviations in Figure 6, the $\mu = 1$ model take substantial incursions into the unstable parameter region, producing sizable pricing errors. The coefficients are subsequently thrown deep into the stable region. This generates noise in the price that detracts from the information extraction process.

[Figure 6 about here]

4. Uniform population

Rather than forcing traders to choose between fundamental information and market-based information, this section allows all traders full access to all information. The

developed model is a pure learning model. Traders share a uniform approach to analysis thus there is no population proportion to track.

4.1 Learning

For convenience, refer to the original model based on separate groups of traders engaged in either fundamental or market-based analysis examined in Sections 3 as the "divided information" model (DI model). Refer to the version developed below based on a uniform population as the "full information" model (FI model). In the FI model, all traders receive an idiosyncratic signal on next period's dividend, s_{it} . They also attempt to incorporate the current price into their estimation of next period's value. Traders employ the same regression analysis approach as in (8), adding the signal s_{it} to the individual's x_{it} . With individual private signals trader i estimates the learning model,

$$\Delta z_t = c_{1it} \Delta p_{t-1} + c_{2it} \Delta d_{t-1} + c_{3it} s_{it} + \zeta_{it} \text{ where} \quad (21)$$

$$s_{it} = \varepsilon_{t+1} + e_{it} \text{ with } e_{it} \sim IIDN(0, \sigma_e^2).^{10} \quad (22)$$

Let $C_{it} = [c_{1it} \ c_{2it} \ c_{3it}]$ be the average of the individual regression coefficients, where $x_t' = [\Delta p_t \ \Delta d_t \ \varepsilon_{t+1}]$. Individual demand is

$$q_{it} = (C_{it} x_t + (p_t + d_t) - R p_t) / \gamma \sigma_{it}^2. \quad (23)$$

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

$$p_t = b_{1t} d_t + b_{2t} d_{t+1} + b_{3t} p_{t-1} + b_{4t} \Delta d_t, \quad (24)$$

but with

$$b_{1t} = \Psi_t^{-1} \sum_i \theta_{it} (1 - c_{3it}), \quad b_{2t} = \Psi_t^{-1} \sum_i \theta_{it} c_{3it},$$

¹⁰ Note that s_{it} is a signal on the innovation in d_t to d_{t+1} rather than the value of d_{t+1} as used in sections 2 and 3. The signal provides the trader with the same information, but is more consistent with its use in regression (21).

$$b_{3t} = \Psi_t^{-1} \sum_i \theta_{it} (-c_{1it}), \quad b_{4t} = \Psi_t^{-1} \sum_i \theta_{it} c_{2it}, \quad \text{and} \quad \Psi_t = \sum_i \theta_{it} (R - (1 + c_{1it})).$$

At the fixed point solution all traders share the same time consistent beliefs. Repeated substitution of the lagged price is of the same linear structure as found in (14),

$$p_t = B_1 d_t + B_2 d_{t+1} + B_3 (1 - b_3 L)^{-1} d_{t-1}. \quad (25)$$

The full system of dynamic equations for the FI model is thus (1), (24), and (9) based on the beliefs of the traders captured in (21). This system of equations does not have a fixed

point since $B_3 = 0$ reduces $c_3 = \frac{(1 + B_1 + B_2 b_3) B_3^2 \sigma_\varepsilon^2}{B_3^2 \sigma_\varepsilon^2 + (B_2^2 + B_3^2 (1 - B_2^2)) \sigma_\varepsilon^2}$ to zero. This is another

reflection of the GS impossibility of informationally efficient price. If the learning process converges (as reflected in $B_3 = 0$), then the traders will be able to perfectly extract d_{t+1} from the price. The private signal becomes redundant information that is inferior due to the idiosyncratic noise. Once $c_3 = 0$, the price no longer contains information about d_{t+1} .

4.2 Simulation

Simulations of the FI model generate three outcomes. The most common (about 79% of the simulations)¹¹ produces a divergent price based on $c_{1t} = R$, which also produces $b_{3t} = R$. This outcome does not represent a bubble solution as the coefficients are not fixed points to the learning process.

Both of the remaining two outcomes produce a bounded price. About 12% of the simulations converge towards coefficients in a market similar to that produced by the DI model. Summary statistics from a representative simulation based on $N = 2$ are presented in the last two columns of Table 1. They show coefficients that converged to a stable small neighborhood containing the values $\{B_1, B_2, B_3, b_3\} \rightarrow \{0, 1/(R-1), 0, -1\}$. Though there is

¹¹ The proportions appear consistent for $N = 2$, $N = 10$, and $N = 100$ traders based simulations.

no longer a fixed point (stable or unstable) to track, it is clear that $b_3 = -1$ is a boundary defining the asymptotic behavior of these simulations in which the model switches between the stable and unstable conditions on either side of $b_3 = -1$. The value of $c_{1t} < 0.02$ in the FI model ensures that the aggregate demand curve is downward sloping. The realization of $c_{3t} > 0$ means that the traders incorporate their private signal into their forecast. The traders in Hussman (1992) also learn to use both fundamental and market information.

Figure 7 reveals a consistent pattern of high and low price dispersion as the model moves in and out of the unstable region.

[Figure 7 about here]

A notable difference between the DI model and this particular outcome of the FI model is the extent to which the price reflects d_{t+1} . The required presence of the fundamental traders in the DI model makes d_t a component of the price. The DI model produces $B_{1T} > 0$, $B_{2T} < 1/(R-1)$. The failure to fully reflect d_{t+1} leaves an exploitable pattern in the price that traders cannot address without creating instability. By contrast, the FI model produces $B_{1T} = 0$ and $B_{2T} = 1/(R-1)$, so that the fundamental component of the price fully reflects aggregate knowledge of the market. Incorporation of the fundamental information by fully informed traders produces an unbiased though still noisy price.

The remaining 9% of the simulations terminate (at $T = 200,000$) with $c_{1t} < 0$ and $0 < b_{3t} < 1$, suggesting a contrarian strategy. However, this outcome does not appear to be a long run asymptotic solution. Allowing one simulation to run for one million periods, the parameters continue to make slow but steady progress in the direction of the $b_3 = -1$ solution.

5. Conclusion

Least-squared learning is used to model the learning dynamics of investors trying to extract information from the market price of a financial asset. The replicator dynamic is introduced to explore the dynamics of the population as individual traders choose between using fundamental and market information. The model offers a new examination of the Grossman and Stiglitz (1980) observation that no interior solution to the population proportion exists unless there is noise in the price that hinders the information extraction efforts of the market informed traders. This model produces results that are consistent with the GS observation, but are obtained without the exogenous insertion of noise. The endogenous noise is assured as a result of the learning environment. The presence of the noise ensures a stable interior proportion of fundamental and market-based strategies.

Allowing all traders to access both fundamental and market information eliminates the population division, leaving just the learning dynamic. No fixed point exists in this learning environment. One of the three possible outcomes generated in simulations produces a market that behaves much like that produced by the divided information dynamic model. Traders make use of both the market and fundamental information. The price is useful because it reflects the aggregate belief of the trader population with the market serving to filter out idiosyncratic biases. The private signal remains important because of the noise that remains present in the price.

The observed "volatility clustering" suggests that further analysis may be warranted to determine whether the model produces a market time-series consistent with observed market patterns.

Appendix A

A.1 Demonstration that equation (8) is the full information regression.

Borrowing from Hussman (1992), a "full information" regression is one in which the learning agents employ all of the useful information to which they have access and no unnecessary variables. With $B_3 = 0$ at the fixed point, from (14),

$$p_t + d_t = (1 + B_1)d_t + B_2d_{t+1} \quad (\text{A.1})$$

so that

$$\Delta z_{t+1} = (1 + B_1)(d_{t+1} - d_t) + B_2(d_{t+2} - d_{t+1}) = (1 + B_1)\varepsilon_{t+1} + B_2\varepsilon_{t+2}. \quad (\text{A.2})$$

At time t , d_{t+1} and d_{t+2} (and thus ε_{t+1} and ε_{t+2}) are yet to be observed directly, however,

$$\Delta p_t = B_1(d_t - d_{t-1}) + B_2(d_{t+1} - d_t) = B_1\varepsilon_t + B_2\varepsilon_{t+1}, \quad (\text{A.3})$$

which is correlated with Δz_{t+1} through the ε_{t+1} term. The innovation Δp_t is not a perfect signal of Δz_{t+1} because it is also a function of ε_t . The traders have observed ε_t indirectly as $d_t - d_{t-1}$ so that including this term in the regression controls for the source of noise.

A.2 Coefficients of (12) expressed as a function of the trader demand parameters.

$$b_{1t} = \Psi_t^{-1} \left(n_t \theta_t^F \frac{R}{R-1} (1-\beta) + (1-n_t) \theta_t^R \right), \quad (\text{A.4})$$

$$b_{2t} = \Psi_t^{-1} \left(n_t \theta_t^F \frac{R}{R-1} \beta \right),$$

$$b_{3t} = \Psi_t^{-1} (1-n_t) \theta_t^R (-c_{1t}),$$

$$b_{4t} = \Psi_t^{-1} (1-n_t) \theta_t^R c_{2t}, \text{ and}$$

$$\Psi_t = n_t \theta_t^F R + (1-n_t) \theta_t^R (R - (1+c_{1t})).$$

A.3 Coefficients B_{1t} , B_{2t} , and B_{3t} in (13) expressed as functions of b_{1t} through b_{4t} .

$$B_{1t} = b_{1t} + b_{4t} + b_{3t}b_{2t-1}, \quad (\text{A.5})$$

$$B_{2t} = b_{2t}, \text{ and}$$

$$B_{3t} = b_{3t}b_{3t-1}b_{2t-2} + (b_{1t-1} + b_{4t-1})b_{3t} - b_{4t}.$$

A.4 Fixed point coefficients B_1 , B_2 , and B_3 in (14) expressed as functions of b_1 through b_4 .

$$B_1(n) = b_1(n) + b_4(n) + b_3(n)b_2(n), \quad (\text{A.6})$$

$$B_2(n) = b_2(n), \text{ and}$$

$$B_3(n) = b_2(n)b_3(n)^2 + (b_1(n) + b_4(n))b_3(n) - b_4(n).$$

A.5 Model's REE fixed point solution supporting Propositions 1 and 2.

For $0 < n \leq 1$:

$$B_1 = \frac{n(1-\beta)\theta^F}{(R-1)(n\theta^F + (1-n)\theta^R)}, \quad B_2 = \frac{n\beta\theta^F + (1-n)\theta^R}{(R-1)(n\theta^F + (1-n)\theta^R)}$$

$$c_1 = (1 + B_1) / B_2 = \frac{n(R-\beta)\theta^F + (1-n)(R-1)\theta^R}{n\beta\theta^F + (1-n)\theta^R}, \quad (\text{A.7})$$

$$c_2 = -B_1(1 + B_1) / B_2 = -\frac{(n(R-\beta)\theta^F + (1-n)(R-1)\theta^R)n(1-\beta)\theta^F}{(R-1)(n\theta^F + (1-n)\theta^R)(n\beta\theta^F + (1-n)\theta^R)},$$

$$\sigma_F^2 = ((1-\beta)(R/(R-1))^2 + b_2^2)\sigma_\varepsilon^2, \text{ and} \quad (\text{A.8})$$

$$\sigma_R^2 = b_2^2\sigma_\varepsilon^2 \quad (\text{A.9})$$

The fixed point solution for b_1 through b_4 follows from plugging the regression coefficient solutions in (A.7) into the price coefficients in (A.4). In particular,

$$b_3 = -\left(\frac{(1-n)\theta^R}{nR\beta\theta^F}\right)\left(\frac{n\theta^F(R-\beta) + (1-n)\theta^R(R-1)}{n\theta^F + (1-n)\theta^R}\right). \quad (\text{A.10})$$

Both terms within the parenthesis are non-negative. The first term is equal to zero at $n = 1$ and is infinite at $n = 0$. The second term is positive. $\partial b_3 / \partial n |_{n \rightarrow 0} \rightarrow \infty$, $\partial b_3 / \partial n |_{n=1} > 0$, and $\partial^2 b_3 / \partial n^2 < 0$, thus b_3 is monotonically increasing in n .

For $n = 0$:

$$c_1 = c_2 = 0, B_1 = b_1 = 1/(R-1), \text{ and } b_2 = b_3 = b_4 = 0. \quad (\text{A.11})$$

A.6 Expected profits at the REE fixed point.

$$\left. \begin{aligned} E(\pi^F) &= \theta^F R^2 B_1 \left(\frac{\beta}{R-1} - B_2 \right) \sigma_\varepsilon^2 \\ E(\pi^R) &= \theta^R R^2 B_1^2 \sigma_\varepsilon^2 \end{aligned} \right\} \text{for } 0 < n \leq 1 \text{ and} \quad (\text{A.12})$$

$$E(\pi^F) = \frac{\beta}{\gamma(1-\beta)}, E(\pi^R) = 0 \text{ at } n = 0. \quad (\text{A.13})$$

Using B_1 and B_2 from (A.7) yields, for $n > 0$,

$$E(\pi^F) = -\frac{n(1-n)(1-\beta)^2 R^2 \theta_1^2 \theta_2}{(R-1)^2 (n\theta^F + (1-n)\theta^R)^2} \sigma_\varepsilon^2 \text{ and} \quad (\text{A.14})$$

$$E(\pi^R) = \frac{n^2 R^2 \theta^{F^2} \theta^R (1-\beta)^2}{(R-1)^2 (n\theta^F + (1-n)\theta^R)^2} \sigma_\varepsilon^2.$$

The solution for $E(\pi^F)$ is weakly negative for $0 < n \leq 1$ with $E(\pi^F) = 0$ for $n = 1$ and $E(\pi^F) < 0$ for $0 < n < 1$. At $n = 0$, $E(\pi^F) > 0$. The solution for $E(\pi^R)$ is smooth, continuous, and increasing in n with $E(\pi^R) = 0$ for $n = 0$ and $E(\pi^R) > 0$ for $n > 0$.

A.7 Proof of Proposition 4.

Solve for profits based on incorrect coefficients to the regression equation (8). Let \hat{c}_1 be the traders' estimate of c_1 . The coefficient $\hat{B}_2 = B_2(\hat{c}_1) = b_2(\hat{c}_1)$ as expressed in (A.4). Recall assumptions $B_3 = 0$, $B_1(c_{1t}) + B_2(c_{1t}) = 1/(R-1)$, and $c_{2t} = -B_2(c_{1t})c_{1t}$. From (7) and (10),

$$q^F = \theta^F R(\beta/(R-1) - \hat{B}_2) \varepsilon_{t+1} \text{ and}$$

$$q^R = \theta^R (\hat{c}_1 - (R-1)) \varepsilon_{t+1}.$$

The assumptions produce $p_{t+1} + d_{t+1} - Rp_t = R(1/(R-1) - B_2) \varepsilon_{t+1}$. Use (17) to obtain

$$E(\pi^F) = \theta^F R^2 (\beta / (R-1) - \hat{B}_2) (1 / (R-1) - \hat{B}_2) \sigma_\varepsilon^2 \text{ and}$$

$$E(\pi^R) = \theta^R R (\hat{c}_1 - (R-1)) (1 / (R-1) - \hat{B}_2) \sigma_\varepsilon^2.$$

Both q^F and q^R equal zero at $\hat{c}_1 = c_1^- = (R-1)$. Let $c_1^+(n)$ indicate the value of \hat{c}_1 at which $(1 / (R-1) - \hat{B}_2) = 0$. The closed form expression is needlessly long, but $c_1^+(n)$ is increasing in n as plotted in Figure 3 using parameters employed in the body of the paper.

Also needlessly complex to include are the first and second derivatives of $D(n, c_1)$ taken with respect to c_1 . $D_c(n, \hat{c}_1) < 0$ at $\hat{c}_1 < c_1$, $D_c(n, \hat{c}_1) = 0$ at $\hat{c}_1 = c_1$, and $D_c(n, \hat{c}_1) > 0$ at $c_1 < \hat{c}_1 < \bar{c}_1$. $D_{cc}(n, \hat{c}_1) > 0$ for all c_1 . $D(n, c_1) \rightarrow \infty$ as $c_1 \rightarrow \bar{c}_1$.

References

- Barberis, N., Shleifer, A., 2003. Style investing. *Journal of Financial Economics* 68, 191-199.
- Blume, L., Easley, D., O'Hara M., 1994. Market statistics and technical analysis: The role of volume. *The Journal of Finance* 49(1), 153-181.
- Börger, T., Sarin, R., 1997. Learning through reinforcement and replicator dynamics. *Journal of Economic Theory* 77, 1-14.
- 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.
- Brock, W.A., Hommes, C.H., 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control* 22, 1235-1274.
- Brock, W.A., LeBaron, B., 1996. A dynamic structural model for stock return volatility and trading volume. *The Review of Economics and Statistics* V#, 94-110.
- Brown, D.P., Jennings, R.J., 1989. On technical analysis. *The Review of Financial Studies* 2(4), 527-551.
- Cheung, Y.W., Friedman, D., 1998. A comparison of learning and replicator dynamics using experimental data. *Journal of Economic Behavior and Organization* 35, 263-280.
- de Fontnouvelle, P., 2000. Information dynamics in financial markets. *Macroeconomic Dynamics* 4, 139-169.
- Droste, E., Hommes C.H., Tuinstra, J., 2002. Endogenous fluctuations under evolutionary pressure in Cournot competition. *Games and Economic Behavior* 40, 232-269.
- Grossman, S.J., Stiglitz, J.E., 1980. On the impossibility of informationally efficient markets.

- The American Economic Review 70(3), 393-408.
- Grundy, B.D., McNichols, M., 1989. Trade and revelation of information through prices and direct disclosure. *The Review of Financial Studies* 2(4), 527-551.
- Hellwig, M.F., 1980. On the aggregation of information in competitive markets. *Journal of Economic Theory* 22, 477-498.
- Hong, H., Stein, J.C., 1999. A unified theory of underreaction, momentum trading, and overreaction in asset markets. *The Journal of Finance* 54(6), 2143-2184.
- Hopkins, E., 2002. Two competing models of how people learn in games. *Econometrica* 70(6), 2141-2166.
- Hussman, J.P., 1992. Market efficiency and inefficiency in rational expectations equilibria. *Journal of Economic Dynamics and Control* 16, 655-680.
- 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.
- Routledge, B.R., 1999. Adaptive learning in financial markets. *The Review of Financial Studies* 12(5), 1165-1202.
- Sethi, R., Franke, R., 1995. Behavioural heterogeneity under evolutionary pressure: Macroeconomic implications of costly optimisation. *The Economic Journal* 105, 583-600.
- Taylor, P.D., Jonker, L.B., 1978. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences* 40, 145-156.
- Timmermann, A., 1996. Excess volatility and predictability of stock prices in autoregressive

dividend models with learning. *Review of Economic Studies* 63, 523-557.

Table 1. Asymptotic behavior in simulation

Simulation duration is $T = 200,000$ with $\rho = 0.1$. The population means and standard deviations for the endogenous parameters are estimated from the last 10,000 observations of the simulation.

Coefficient	$\mu = 1/t$		$\mu = 1$		Full Information	
	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev
c_1	0.0581	0.0004	0.1722	0.0003	0.0098	0.0001
c_2	-0.5114	0.0077	-1.5973	0.0061	-0.0033	0.0073
c_3	---	---	---	---	0.5080	0.0049
B_1	8.8080	0.7745	10.9712	8.0724	0.0309	0.9518
B_2	41.1945	0.3860	39.0224	4.8968	49.8176	0.6455
B_3	0.0031	0.7749	0.6507	8.1448	0.2982	1.0869
b_3	-0.9874	0.0219	-0.6346	0.2216	-0.9615	0.0264
n	0.1236	0.0014	0.4150	0.0826	---	---




Figure 1: Comparison of the REE predicted values (solid black lines) to the terminal values of endogenous price parameters from simulation with 10,000 periods (data points). For b_3 , a line at -1 is also included. The simulations consist of 20 random draws of a dividend series. Each series is examined for $0 < n \leq 1$ using an incremental step size of 0.01.




Figure 2: Time series reflecting model behavior at $b_3 = -1$ boundary. $n = 0.05 < \underline{n}$, $b_3 = -2.0938$, $c_1 = 0.0372$. Top frame b_{3t} , middle frame c_{1t} , bottom frame $p_t - p_t^F$.




Figure 3: The phase space for the two processes, n_t (population) and c_{1t} (learning). $E(\Delta c_{1t})=0$ at $c_{1t} = c_1(n)$. $E(\Delta n_t)=0$ at $c_{1t} = c_1^-$ and $c_{1t} = c_1^+$. Above \bar{c}_1 is the infeasible region. For $n < \underline{n}$ the learning process is unstable at the REE fixed point.




Figure 4: Plot of REE fixed point $E(\pi^F)$ and $E(\pi^R)$ (solid black lines) and simulation average realizations of π^{Fe} (\circ) and π^{Re} ($+$) computed from the last 1,000 periods of 20 10,000 period simulations. Each series is examined for $0 < n \leq 1$ using an incremental step size of 0.01.




Figure 5: Phase plot in $(n_t, c_{1t} - c_1(n))$ space. Sample convergence of a simulation towards $(n^{**}, c_1^*(n^{**}) - c_1) = (0.1236, -0.006)$. Starting value $n_0 = 0.75$, and $\mu_t = 1/t$.

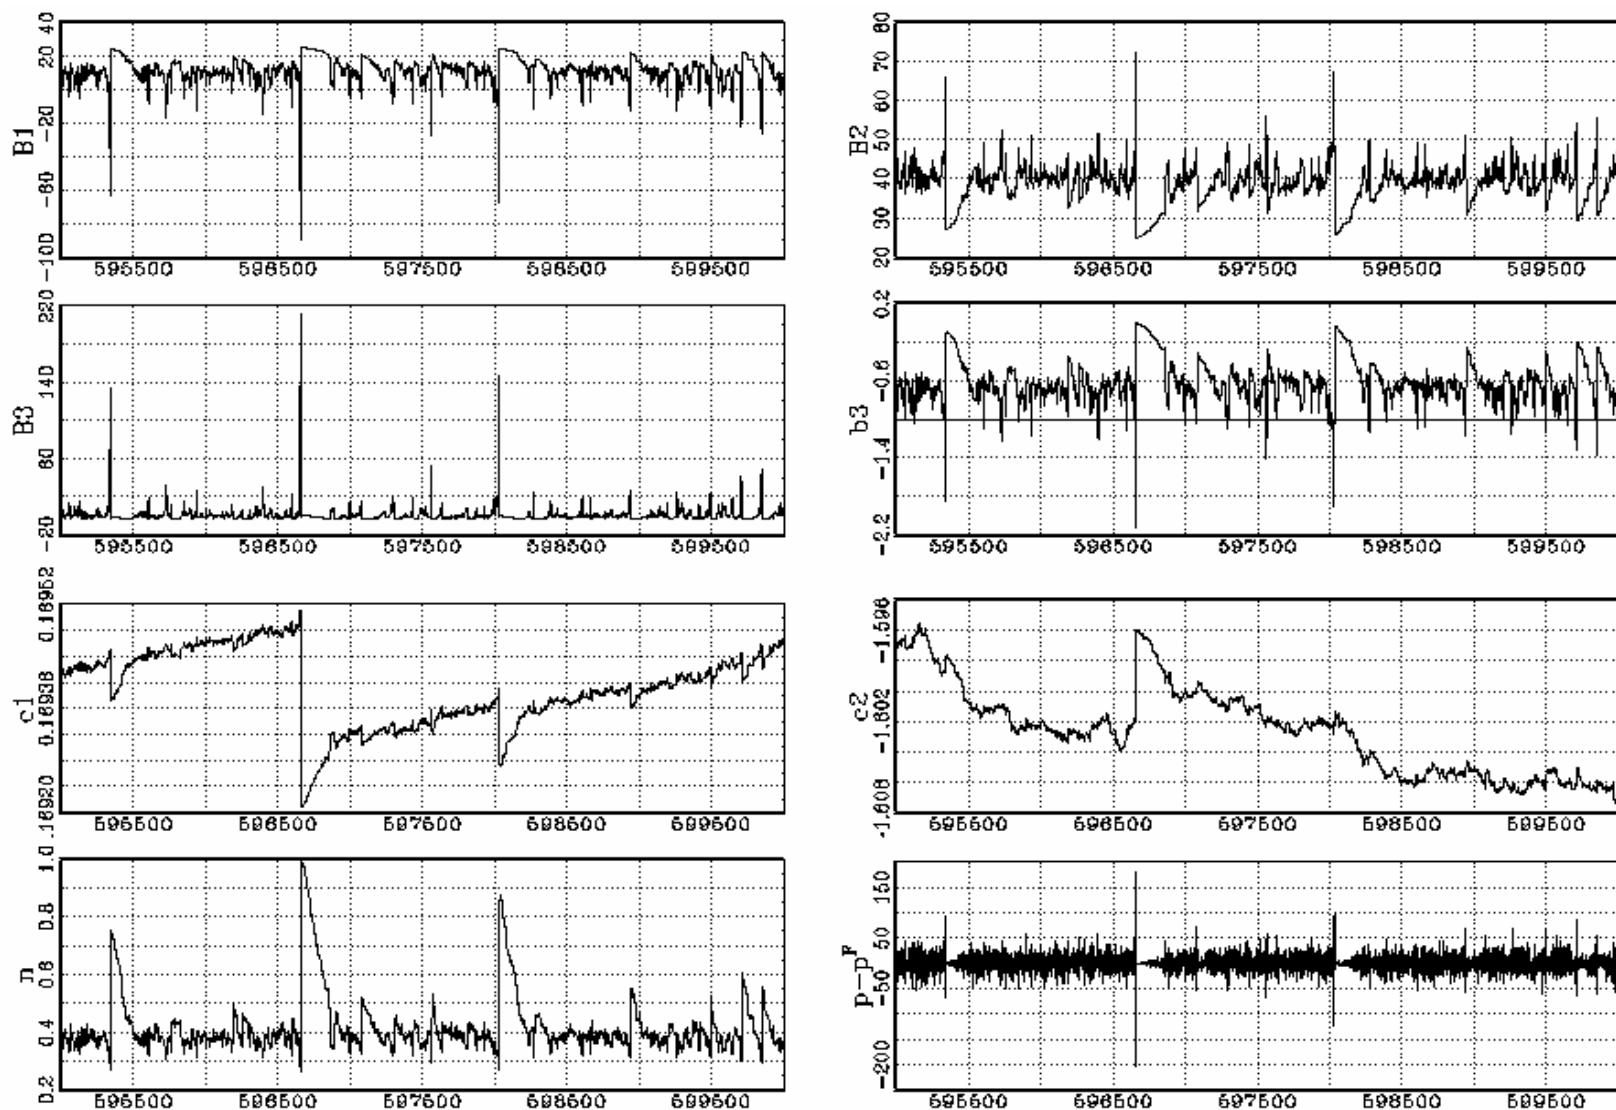


Figure 6: Final 5,000 observations of a $T = 600,000$ simulation. Starting value $n_0 = 0.75$, $\mu = 1$, $\rho = 0.1$.

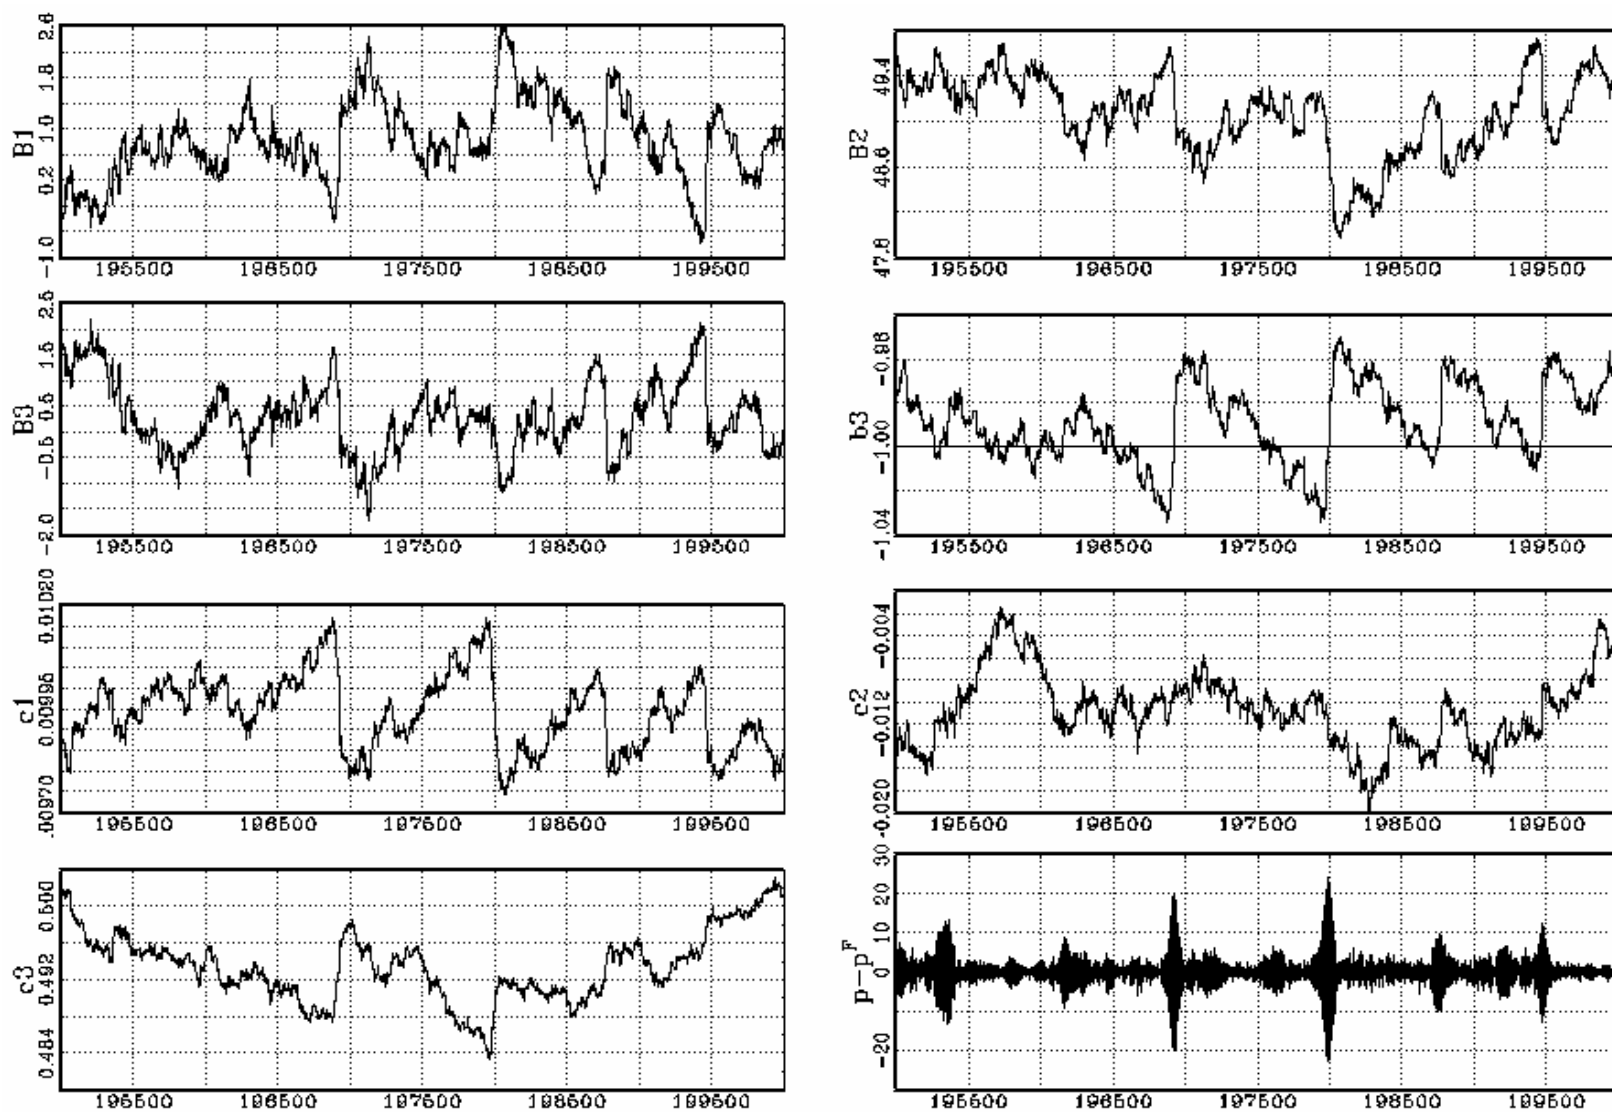


Figure 7: Traders access both fundamental and market information. Final 5,000 observations of $T = 200,000$.