Journal of Economic Behavior & Organization 76 (2010) 635–653

Contents lists available at ScienceDirect

Journal of Economic Behavior & Organization

journal homepage: www.elsevier.com/locate/jebo

Learning and adaptation's impact on market efficiency☆

David Goldbaum[a,*], Valentyn Panchenko[b]

[a] University of Technology Sydney, School of Finance and Economics, PO Box 123 Broadway, NSW 2007, Australia
[b] University of New South Wales, Australian School of Business, Australia

ARTICLE INFO

Article history:
Received 12 July 2006
Received in revised form 2 September 2010
Accepted 5 September 2010
Available online 21 September 2010

JEL classification:
G14
D82
D83
C61
C62

Keywords:
Market efficiency
Learning
Dynamic systems

ABSTRACT

A dynamic model with learning and adaptation captures the evolution in trader beliefs and trading strategies. Through a process of learning and observation, traders improve their understanding of the market. Traders also engage in a process of adaptation by switching between trading strategies based on past performance. The asymptotic properties are derived analytically, demonstrating that convergence to efficiency depends on the model of adaptation.

1. Introduction

Financial markets offer considerable evidence suggesting traders do not have a full and complete understanding of the process by which prices are determined. There also seems to be a lack of consensus among traders in interpreting what the current price signals about future payoffs. Evidence of disagreement includes, but is not limited to, excess volatility, high trading volume, the considerable variety in trading strategies, and the fervor with which traders seek to improve their models. This suggests that markets have yet to achieve the rational expectations equilibrium described by Grossman and Stiglitz (1980). Instead, traders appear to process information, develop strategies, and adapt to changing market conditions.

To gain insight into the impact of evolution in trading strategies on financial markets, this paper recasts Grossman and Stiglitz (GS) as a fully specified dynamic model. The absence of a rational expectations equilibrium in the presence of a revealing price ensures that there is always room for improvement in trading strategy. The two dynamics, learning and strategy adoption, capture two aspects of trader adaptation to the observed market environment. The single period terminal asset of GS offers tractability in a sufficiently rich environment to examine the key elements of adaptation. The asymptotic properties of the market governed by these two dynamic process are developed herein.

☆ We would like to thank the people at CeNDEF for providing the stimulating research environment in which this project developed. VP acknowledges the support under Australian Research Council's Discovery Projects funding scheme (project number DP0986718).

* Corresponding author. Tel.: +61 2 9514 7734; fax: +61 2 9514 7722.
E-mail address: david.goldbaum@uts.edu.au (D. Goldbaum).

0167-2681/$ – see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.jebo.2010.09.003

The GS model places informed traders in the market with rational uninformed traders. The former gain access to private information through costly research. The latter seek to extract the private information from the price. Each trader makes full use of the information to which he or she has access. This project retains much of the structure of the GS market as well as the notion that traders make full use of the available information.[1] In this behavior, they are rational, but here the uninformed traders are boundedly rational in that they are denied knowledge of the true relationship between the observed price and the expected payoff. The traders instead must, as in Bray (1982), attempt to learn the relationship through observation of market data. This is the learning process that GS presume to have already taken place prior to their analysis.[2] Distinct from Bray (1982), concurrent with the learning process, the traders choose between the two information options. Such a dynamic population is implicit in the GS model, but is not explicitly modeled in their static examination.[3] GS consider the fixed point equilibrium of the population process, having presumed prior convergence to rational expectations in the traders' beliefs. The importance of examining both as simultaneous processes is that the two dynamic processes, learning and the population process, interact, affecting market behavior both during evolution and in the asymptotic convergence.

Two types of population processes are examined and compared to explore the different implications for how they shape the market's asymptotic behavior. One family includes the Discrete Choice Dynamics introduced by Brock and Hommes (1997, 1998) and the other family includes Replicator Dynamics. Both processes have received extensive attention in the economics literature as tools for modeling evolving populations in a discrete choice setting.

The second key alteration to the GS model is to remove the random supply of the risky security, thereby removing the mechanism that ensures the existence of an equilibrium to the population process. As a result, the information advantage held by the informed traders only exists while the learning process is ongoing. Just how the dynamic model resolves the conflict between convergence in the learning process and the market's need for noise in the price offers insight into the evolution that results from traders' adaptation to endogenously changing market conditions. The resolution depends on the interaction between the two dynamic processes as it is shaped by the potential absence of the population fixed point. The nature of this interaction determines whether market efficiency is the limiting case.

The tractability of the present model is absent in the Goldbaum (2005) and Goldbaum (2006) simulations, also based on models of learning and strategy adoption. This examination's simpler market, partially the result of the terminal risky asset, enables analytical investigation of the relevant dynamic processes, eliminating the dependence on simulation based analysis.

The paper proceeds as follows. Section 2 develops the model and establishes the conditions for existence and stability of the rational expectations equilibrium as the fixed point to the dynamic processes. Section 3 considers the issue of market efficiency analytically and in simulation. Section 4 uses simulation to determine how the rate of learning is affected by the population process. Section 5 concludes.

2. Model and fixed point equilibrium

2.1. Market

Adopting a repeated GS framework, in each period a population of $N = N^I + N^U$ informed (I) and uninformed (U) traders trade a risky asset and a risk-free bond. The risk-free bond, with a price of one, pays R at the end of the period. The risky asset is purchased at the market determined price, p_t. At the end of the period it pays a randomly determined terminal value, u_t.

$$u_t = \bar{u} + \theta_t + \varepsilon_t \tag{1}$$

with, $\theta_t \sim N(0, \sigma_\theta^2)$, $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$, and $\text{cov}(\theta_t, \varepsilon_t) = 0$.

The determinates of u_t are IID. This process is known to the traders.

In each period, each trader maximizes a negative exponential utility function of end of period wealth. Under the assumption of normality in returns, the resulting demand for the risky asset is

$$q_{it}(p_t) = (E_{it}(u_t) - Rp_t)\gamma\phi_{it} \tag{2}$$

with $\phi_{it} = 1/\sigma_{it}^2$. Here, $1/\gamma$ is the coefficient of absolute risk aversion and $E_{it}(u_t) = E(u_t|F_{i,t})$ and $\sigma_{it}^2 = \text{var}(u_t|F_{i,t})$ are trader i's conditional expectations and variance, respectively, of his or her forecast error.

Define $\lambda_t = N_t^I/N$ to be the proportion of informed traders, leaving $(1 - \lambda_t)$ as the proportion of uninformed traders. Let q_t^k be the per capita demand for the risky security among group $k = I, U$ traders in period t. In a Walrasian equilibrium, the market price equates supply and demand for the risky asset. The supply is fixed, avoiding the introduction of exogenous

[1] This is in contrast to other models of multiple trader types with switching in which the market-based traders are limited in their effort to extract information from past prices rather than from contemporaneous information, as employed in Brock and Hommes (1998), Chiarella and He (2001), Föllmer et al. (2005), and Gaunersdorfer et al. (2008), among others.

[2] "They [traders] learn the relationship between the distribution of return and the price, and use this in deriving their demand for the risky assets" (p. 394).

[3] "We can calculate the expected utility of the informed and the expect utility of the uninformed. If the former is greater than the latter (taking account of the cost of information), some individuals switch from being uninformed to being informed (and conversely) (p. 394).

noise to the price. For convenience, set fixed net supply of the risky asset to zero. Market equilibrium requires a price that clears the market such that

$$0 = \lambda_t q_t^I(p_t) + (1-\lambda_t) q_t^U(p_t). \tag{3}$$

2.2. Information and forecasting

The informed traders receive a uniform signal revealing the value of θ_t. They rationally forecast

$$E(u_t|\theta_t) = \bar{u} + \theta_t. \tag{4}$$

The uninformed traders attempt to forecast end of period payoff based on the information contained in the price. They are as aware of the payoff process as are the informed traders. They thus know that unconditionally, $E(u_t) = \bar{u}$ and have an unconditional price expectation of $E(p_t) = \bar{u}/R$. Without knowledge of the true relationship between price and the underlying information, they engage in a process of least-squares learning by estimating[4]

$$u_t - \bar{u} = c_0 + c_1(p_t - E(p_t)) + e_t. \tag{5}$$

Eq. (5) is the uniformed traders' perceived model relating the observable p_t to the end of period u_t. It is a model of both information extraction and forecasting. The ability of the uniformed to extract information and predict the payoff is based on how close the perceived model is to the actual relationship. The uninformed traders update the coefficients according to the least-squares updating algorithm.

From Marcet and Sargent (1989a,b),

$$\mathbf{c}_t = \mathbf{c}_{t-1} + (Q_t^{-1} x_{t-1}(u_{t-1} - \bar{u} - \mathbf{c}_{t-1} x_{t-1}))'/t \tag{6}$$

$$Q_t = Q_{t-1} + (x_{t-1} x'_{t-1} - Q_{t-1})/t \tag{7}$$

with $x_t = \{1, p_t - \bar{u}/R\}$ and $\mathbf{c}_t = \{c_{0t}, c_{1t}\}$. Uninformed traders forecast

$$E(u_t|p_t) = \bar{u} + c_{0t} + c_{1t}(p_t - \bar{u}/R). \tag{8}$$

As reflected in (6), least-squares learning assigns equal weight to each past observation, implying infinite memory. This is consistent with the objective to explore the interaction between market efficiency and learning with traders taking full advantage of all available information. As will become evident with the model's development, error in the traders' model creates noise in the price process. Providing the traders full access to all of the historical data is a prerequisite to allowing the traders to learn the correct model. Finite memory or otherwise limiting the use of past information prevents convergence in learning, introducing unnecessary error into the system. A correctly specified model and least-squares learning applied to the unknown parameters make possible the traders' learning of the correct model. At this stage of analysis, it remains unclear how the paradox of the non-existence of a REE equilibrium is resolved by the dynamic system. Limiting trader memory or discounting more distant observations can resolve the paradox by ensuring asymptotic error in the model of the uninformed. This mechanism was employed by Branch and Evans (2006).

2.3. Learning fixed points

Let p_t be determined as the market clearing Walrasian price. Based on the demand of the mixed population of traders, price takes the form

$$p_t = b_0(\mathbf{c}_t, \lambda_t) + b_1(\mathbf{c}_t, \lambda_t)\theta_t. \tag{9}$$

The presence of $\mathbf{c}_t$ and λ_t in the price coefficients reflects the impact of trader beliefs on how information is incorporated into the price. Solve for the $\mathbf{c}_t$ and λ_t dependent market equilibrium to obtain the coefficients:

$$\begin{aligned} b_{0,t} &= \frac{(1-\lambda_t)c_{0,t}\phi_t^U}{\lambda_t R\phi^I + (1-\lambda_t)(R - c_{1,t})\phi_t^U} + \frac{\bar{u}}{R} \\ b_{1,t} &= \frac{\lambda_t \phi_t^I}{\lambda_t R\phi_t^I + (1-\lambda_t)(R - c_{1t})\phi_t^U} \end{aligned} \tag{10}$$

where[5]

$$\phi^I = 1/\sigma_I^2 = 1/\sigma_\varepsilon^2, \tag{11}$$

[4] An alternative would be to have the uninformed traders more simply consider the relationship between the price and payoff, $u_t = c_0 + c_1 p_t + e_t$. The alternate results in the replacement of σ_θ^2 that appear in subsequent analysis with $(\bar{u}^2 + \sigma_\theta^2)$. There is no substantive difference between the two approaches.

[5] Note that the presence of $b_{1,t}$ in ϕ_t^U means that (10) is an implicit solution. This will be formally handled in the discussion that follows regarding Rational Expectations and the Boundedly Rational Expectations. A reduced form solution is available, but pages long.

$$\phi_t^U = 1/\sigma_{U,t}^2 = ((\sigma_\varepsilon^2 + (1 - c_{1,t}b_{1,t})^2\sigma_\theta^2)^{-1}. \tag{12}$$

Were the parameters of the price equation, b_0 and b_1, set exogenously, the learning process would converge[6] towards

$$c_{0,t} = -c_{1,t}(b_0 - \bar{u}/R), c_{1,t} = 1/b_1. \tag{13}$$

The REE at which beliefs are consistent with the actual price determination process is,

$$c_0^* = 0, \quad c_1^* = R \tag{14}$$

and

$$b_0^* = \bar{u}/R, \quad b_1^* = 1/R. \tag{15}$$

This is consistent with the finding of GS. The solution is independent of the value of λ, reflecting that once the learning process has converged, the market clearing price does not depend on λ. Neither does the uninformed traders' model.

For computational manageability and ease of discussion, allow the uninformed traders to know $E(\varepsilon_t) = 0$, allowing recognition by the traders that the correct value for $c_{0,t} = 0$ so that $b_{0,t} = \bar{u}/R$. The uninformed traders are thus only required to estimate the single parameter c_1. The error in the uninformed traders' model is thus $c_{1,t} - c_1^*$. Deviations in $c_{1,t}$ from c_1^* are the source for deviation of $b_{1,t} = b_1(c_{1,t}, \lambda_t)$ from b_1^*.

As is apparent from (10), for $\lambda_t = 0$ and $c_{1,t} = c_1^*$ the price is undefined, reflecting the lack of fundamental information entering the market. For $\lambda_t = 0$ and $c_{1,t} \neq c_1^*$ the price is either undefined or not reflective of the underlying value.

2.4. Performance

The population process is driven by performance, here measured as profits. Define π_t^k as net profits to an investor in group k after deducting the cost of information acquisition, κ^k,

$$\pi_{i,t}^k = q_{i,t}^k(u_t - Rp_t) - \kappa^k, \quad k = I, U. \tag{16}$$

The presumption of the model is that $\kappa^I > \kappa^U$. Without loss in generality, set $\kappa^U = 0$ and $\kappa^I = \kappa$.

The $c_{1,t}$ and λ_t dependent expected profits are

$$E(\pi_t^I) = \gamma(1 - Rb_1(c_{1,t}, \lambda_t))^2\sigma_\theta^2\phi^I - \kappa \tag{17}$$

$$E(\pi_t^U) = -\gamma(R - c_{1,t})b_1(c_{1,t}, \lambda_t)(1 - Rb_1(c_{1,t}, \lambda_t))\phi_t^U(c_{1,t}, \lambda_t)\sigma_\theta^2. \tag{18}$$

Since the weighted demand of the two groups must sum to zero, uninformed trader profit can also be expressed as a function of the informed traders' trading profit (*pre* cost),

$$E(\pi_t^U) = \frac{-\lambda_t}{1 - \lambda_t}\gamma(1 - Rb_1(c_{1,t}, \lambda_t))^2\sigma_\theta^2\phi^I. \tag{19}$$

The informed traders' information advantage leads to nonnegative expected trading profits while the uninformed traders' profits become increasingly negative the greater the error in their model. The error in extraction is twofold. With $c_{1,t} \neq c_1^* = R$, the uninformed misinterpret the price information. Additionally, the error in $c_{1,t} \neq c_1^*$ feeds back into $b_{1,t} \neq b_1^* = R^{-1}$ so that the price fails to properly reflect the private information, θ_t.

Employment of the correct model by the uninformed traders produces zero trading profits so that, for $0 < \lambda_t$,

$$E(\pi_t^I - \pi_t^U | c_{1,t} = c_1^*) = -\kappa. \tag{20}$$

The analytical solution to profits is generally inaccessible to the traders where $c_{1,t} \neq c_1^*$ for two reasons, both stemming from the presence of $b_1(c_{1,t}, \lambda_t)$ in the profit equations.

The first challenge for the traders is the solution's dependence on λ_t. The value of b_1^* is independent of λ_t but $b_{1,t}$ does depend on λ_t when $c_{1,t} \neq c_t^*$. Knowing λ_t is equivalent to knowing the current trading strategy of the other market participants, which is reasonably placed beyond the reach of the traders. Without knowledge of λ_t, the traders are unable to correctly deduce the price function.

The second challenge to the trader's efforts to solve for expected returns is the complexity of the closed for solution of $b_{1,t}$ As observed in (10), the conditional variance terms of the two trading groups are factors in the pricing equation solutions. The presence originates with the dependence of the individual trader's demand function on the conditional variance, as seen in (2). The conditional variance of the informed traders' error is exogenous and easily computed, but the uninformed traders' conditional variance is endogenous, dependent on the price coefficient $b_{1,t}$, which in this case creates a extremely complex reduced form solution.

[6] Because of stochastic nature of the considered system, by convergence here and hereafter we mean almost sure convergence or convergence with probability one.

These two impediments lend credence to a boundedly rational solution. To address the latter, many papers follow Brock and Hommes (1998) in assuming that the uninformed traders presume a fixed standard error to their conditional forecast, setting $\sigma^2_{U,t} = \sigma^2_I$ so that $\phi^U_t = \phi^I$. While this presumption is correct only when $c_{1,t} = c^*_t$, there are two arguments in support of imposing this presumption generally. With conditional variance only present in the price through the traders' demand functions, if the traders are unable to derive closed form conditional variance and therefore use a substitute, the price then reflects the substitute. Under bounded rationality, if the true standard error is too complicated to be reasonably calculated by the traders, a simple rule such as $\sigma^2_{U,t} = \sigma^2_I$ can be presumed.

The bounded rationality assumption also allows the elimination of the forecast standard errors from the equilibrium price equation (and subsequently, from profits), making an analytical expression of the price and profits attainable and tractable. The resulting model equilibrium will be referred to as the boundedly rational expectations (BRE) solution. This solutions will be compared to that generated from rational expectations (RE) solution derived using the true expression of $b_1(c_{1,t}, \lambda_t)$ solved analytically with the aid of a computer. The complexity of the solution does not lend itself to properly comparative statics, but these can be inferred numerical from the analytical solution.

Let $\Psi(c_{1,t}, \lambda_t)$ be the denominator term of $b_{1,t}$ in (10),

$$\Psi(c_{1,t}, \lambda_t) = \lambda_t R\phi^I_t + (1-\lambda_t)(R - c_{1,t})\phi^U_t \tag{21}$$

$\Psi(c_{1,t}, \lambda_t)$ is also the coefficient on p_t in the aggregate demand equation. Aggregate demand is downward sloping if $\Psi(c_{1,t}, \lambda_t)$ remains positive.

From (9), a finite price requires $\Psi(c_{1,t}, \lambda_t) \neq 0$ and a reasonable market and price solution eliminates $\Psi(c_{1,t}, \lambda_t) < 0$. Let $S|c_{1,t}$ represent the set of all values of λ_t that result in $\Psi(c_{1,t}, \lambda_t) > 0$ given $c_{1,t}$. For the RE solution, all feasible values of λ_t produce $\Psi(c_{1,t}, \lambda_t) > 0$, and thus $S|c_{1,t} = [0, 1]\ \forall\, c_{1,t}$, excluding $\lambda_t = 0$ when $c_{1,t} = R$. Such is not the case under the BRE solution. Solve $\Psi(c_{1,t}, \lambda_t) = 0$ under BRE, using $\phi^I_t = \phi^U_t > 0$, to obtain a critical $\lambda_t = (c_{1,t} - R)/c_{1,t}$ for $c_{1,t} > R$. Let $\lambda^c(c_{1,t}) \in [0, 1]$ represent the lower bound on feasible values of λ_t producing a finite and reasonable price. Thus, for $c_{1,t} > c^*_1 = R$

$$\lambda^c(c_{1,t}) = \begin{cases} \dfrac{c_{1,t} - R}{c_{1,t}} & \text{for BRE} \\ 0 & \text{for RE.} \end{cases} \tag{22}$$

For $c_{1,t} > c^*_1$, as $\lambda_t \to \lambda^c_t$ from above, $b_{1,t} \to \infty$. The greater the upward bias in the uninformed traders' model, the fewer uninformed traders the market can absorb and still produce a finite price. Observe that $\lambda^c(c_{1,t})$ is a monotonically increasing function with $\lambda^c_t \to 0$ for $c_{1,t} \to c^*_1$ from above and $\lambda^c_t \to 1$ for $c_{1,t} \to \infty$.

That $S|c_{1,t}$ includes the full range of $\lambda\ \forall\, c_{1,t}$ under RE is explained by the endogenous $\sigma^2_{U,t}$. Increasing price errors cause this term to grow faster than the linear expected excess return in the numerator of the uninformed traders' demand. The result is that the large price deviation creates sufficient uncertainty among the uninformed traders to dampen their demand. Though $b_{1,t} \to \infty$ as $\lambda_t \to 0$ for $c_{1,t} > c^*_1$, unlike the BRE solution, $b_{1,t}$ remains bounded for $\lambda_t > 0$.

For $c_{1,t} < c^*_1$, the uninformed traders under-respond to price innovations. The under reaction by the uninformed traders does lead to pricing errors, but not the possibly unbounded price error produced by overreaction. The price remains bounded. For $c_{1,t} = c^*_1$, $b_{1,t} = b^*_1 = 1/R\ \forall \lambda \in (0, 1]$.

Define $d\pi_t$ as the expected performance difference,

$$d\pi_t = E(\pi^I_t - \pi^U_t) = g(c_{1,t}, \lambda_t). \tag{23}$$

Using (17) and (19) the expected profit differential can be solved as

$$d\pi_t = \frac{1}{(1-\lambda_t)}(\gamma(1 - Rb_1(c_{1,t}, \lambda_t))^2 \sigma^2_\theta \phi^I) - \kappa \tag{24}$$

where the function $b_1(c_{1,t}, \lambda_t)$ is specific to whether the expectations are BRE or RE. Use (10) in (24) to obtain

$$d\pi_t = \frac{(1-\lambda_t)(c_{1,t} - R)^2 \phi^{U2}_t}{(\lambda_t R\phi^I + (1-\lambda_t)(R - c_{1,t})\phi^U_t)^2} \frac{\gamma\sigma^2_\theta}{\sigma^2_\varepsilon} - \kappa \tag{25}$$

which, under the BRE simplifies to

$$d\pi_t = \frac{(1-\lambda_t)(c_{1,t} - R)^2}{(R + (1-\lambda_t)c_{1,t})^2} \frac{\gamma\sigma^2_\theta}{\sigma^2_\varepsilon} - \kappa.$$

Lemma 1 characterizes $d\pi_t$. The descriptions apply to both the BRE and RE.

Lemma 1. *The function $g(c_{1,t}, \lambda_t)$ displays the following characteristics*:

- L1.1 $d\pi_t = g(c_{1,t}, \lambda_t)$ *is continuous in both inputs for* $\lambda_t > \lambda_t^c$
- L1.2 $-\kappa \leq g(c_{1,t}, \lambda_t) < \infty$ *for* $\lambda_t > \lambda_t^c$
- L1.3 $g_c(c_{1,t}, \lambda_t) \begin{cases} > 0 \text{ for } c_{1,t} > c_1^* \text{ and } \lambda_t > \lambda_t^c \\ < 0 \text{ for } c_{1,t} < c_1^* \end{cases}$
- L1.4 $g_\lambda(c_{1,t}, \lambda_t) \begin{cases} = 0 \text{ for } c_{1,t} = c_1^* \\ < 0 \text{ for } \lambda_t > \lambda_t^c,\ c_{1,t} \neq c_1^* \text{ and } c_{1,t} > -c_1^*(2\phi^I/\phi^U - 1) \end{cases}$
- L1.5 $g(c_{1,t}, \lambda_t) \to \begin{cases} \infty \text{ as } \lambda_t \to \lambda^c(c_{1,t}) \text{ for } c_{1,t} > c_1^* \\ \gamma\sigma_\theta^2/\sigma_\varepsilon^2 \text{ as } \lambda_t \to 0 \text{ for } c_{1,t} < c_1^* \end{cases}$
- L1.6 $g(c_1^*, \lambda_t) = g(c_{1,t}, 1) = -\kappa$

Proof. In the BRE case, the characteristics of $g(c_{1,t}, \lambda_t)$ follow directly from (25). In the RE case, the characteristics of $g(c_{1,t}, \lambda_t)$ are derived numerically from (24) using the RE price solution of $b_{1,t}$. Observe from (11) and (12) that $\phi_t^U \leq \phi^I$.□

Proposition 1.

(a) *For* $c_{1,t} > c_1^*$ *and* $-\kappa \leq x < \infty\ \exists\, \lambda = \lambda^x(c_{1,t})$ *that solves* $x = g(c_{1,t}, \lambda)$. $\lambda^x(c_{1,t})$ *is monotonically increasing with* $\lambda^x(c_{1,t}) \to 0$ *for* $c_{1,t} \to c_1^*$ *from above and* $\lambda^x(c_{1,t}) \to 1$ *for* $c_{1,t} \to \infty$. *For* $a > b$, $\lambda^a(c_{1,t}) > \lambda^b(c_{1,t})$.

(b) *For* $c_{1,t} < c_1^*$, $g(c_{1,t}, 1) = -\kappa$ *and* $g(c_{1,t}, \lambda) \to (\gamma\sigma_\theta^2/\sigma_\varepsilon^2 - \kappa)$ *for* $\lambda \to 0$.
   (i) *For* $-c_1^*(2\phi^I/\phi^U - 1) \leq c_{1,t} < c_1^*$ *and* $-\kappa \leq x < (\gamma\sigma_\theta^2/\sigma_\varepsilon^2 - \kappa)\,\exists \lambda = \lambda^x(c_{1,t})$ *that solves* $x = g(c_{1,t}, \lambda)$. $\lambda^x(c_{1,t})$ *is monotonically decreasing with* $\lambda^x(c_{1,t}) \to 0$ *for* $c_{1,t} \to c_1^*$ *from below. For* $a > b$, $\lambda^a(c_{1,t}) > \lambda^b(c_{1,t})$.
   (ii) *For* $c_{1,t} < -c_1^*(2\phi^I/\phi^U - 1)$, $g_\lambda(c_{1,t}, \lambda_t)|_{\lambda=0} > 0$ *so that* $d\pi$ *is initially rising in* λ *and then falling.*

Proof of (a). Follows from L1.1, L1.3, L1.4, L1.5 and L1.6 of Lemma 1.□

Proof of (b). From L1.4 of Lemma 1, $d\pi$ is a monotonically decreasing function of λ_t. From L1.3 of Lemma 1, an increase in $c_{1,t}$ for $c_{1,t} > c_1^*$ or decrease in $c_{1,t}$ for $c_{1,t} > c_1^*$ increases the $d\pi_t$ function, thus increasing the value of $\lambda_t < \lambda^0(c_{1,t})$ producing $d\pi_t = x$.□

When $c_{1,t} \neq c_1^*$, error is introduced into the price by the uninformed traders' erroneous beliefs. For a large λ_t the knowledge of the informed traders dominates the market to keep the error in the price small. In this case, the informed traders do not, in expectation, recoup the cost of their information advantage and under-perform the uninformed traders. Lowering λ_t magnifies the impact of the error in the uninformed traders' model on the price. In this case, the informed traders are able to out-perform the uninformed traders and may recoup the cost of the private information.

Introduce notation $\lambda^0(c_{1,t})$ as a special case of $\lambda^x(c_{1,t})$ with $x = 0$. From Proposition 1, the nature of $\lambda^0(c_{1,t})$ depends on whether $c_{1,t}$ is less than, equal to, or greater than c_1^*. For $c_{1,t} = c_1^*$, $\lambda^0(c_{1,t})$ does not exist since $g(c_1^*, \lambda) = -\kappa$ for $\lambda > 0$ and $g(c_1^*, 0) > 0$. Let $\lambda^+(c_{1,t})$ represent the portion of $\lambda^0(c_{1,t})$ for which $c_{1,t} > c_1^*$. From Proposition 1, $\lambda^+(c_{1,t})$ is a continuous function within (0,1) with $\lambda^+(c_{1,t}) \to 0$ as $c_{1,t} \to c_1^*$, $\lambda^+(c_{1,t}) \to 1$ as $c_{1,t} \to \infty$, the first and second derivatives of $\lambda^+(c_{1,t})$ are $\lambda_c^+ > 0$, $\lambda_{cc}^+ < 0$ and $\lambda^+(c_{1,t}) > \lambda^c(c_{1,t})$.

For $c_{1,t} < c_1^*$, trading profits are defined by the limits at the two extremes for λ_t, with zero trading profits when $\lambda_t = 1$ and trading profits approaching $\gamma\sigma_\theta^2/\sigma_\varepsilon^2$ as $\lambda_t \to 0$. Let $\lambda^-(c_{1,t})$ represent the portion of $\lambda^0(c_{1,t})$ for which $c_{1,t} < c_1^*$. For $\kappa < \gamma\sigma_\theta^2/\sigma_\varepsilon^2$, $\lambda^-(c_{1,t})$ is a continuous function within (0,1) with $\lambda^-(c_{1,t}) \to 0$ as $c_{1,t} \to c_1^*$, $\lambda^-(c_{1,t}) \to 1$ as $c_{1,t} \to -\infty$, and first derivative, $\lambda_c^- < 0$. For small κ, $\lambda_{cc}^- < 0$. As $\kappa \to \gamma\sigma_\theta^2/\sigma_\varepsilon^2$ from below, $\lambda^-(c_{1,t})$ flattens and becomes convex for $c_{1,t}$ near c_1^* while remaining concave for extremely low values of $c_{1,t}$. Informed trading profits above $\gamma\sigma_\theta^2/\sigma_\varepsilon^2$ can only be achieved if $c_{1,t} < -c_1^*(2\phi^I/\phi_t^U - 1)$ and in this case, the maximum trading profit is achieved by an interior value of λ_t. If $\kappa > \gamma\sigma_\theta^2/\sigma_\varepsilon^2$, then $\lambda^-(c_{1,t})$ is a cone-shaped function for which the highest value of $c_{1,t}$ able to produce $g(c_{1,t}, \lambda_t) = 0$ is less than $-c_1^*(2\phi^I/\phi_t^U - 1)$. To simplify notation, let $\Gamma_t = 2\phi^I/\phi_t^U - 1$ and observe that $1 \leq \Gamma_t$. A value of $c_{1,t}$ less than $-c_1^*\Gamma_t$ implies an extremely large error in the market-based traders' beliefs.

For future reference, define

Condition A: $\kappa < \gamma\sigma_\theta^2/\sigma_\varepsilon^2$.

Fig. 1 plots the RE $\lambda^0(c_{1,t})$ for a range of κ given $\gamma\sigma_\theta^2/\sigma_\varepsilon^2$, capturing the features developed in Lemma 1 and Proposition 1.

While formal analysis requires examination of all possible market conditions, a violation of Condition A should be viewed as an extreme market condition. A low $\sigma_\theta^2/\sigma_\varepsilon^2$ reduces the ability of the informed traders to exploit their information advantage in the market as knowledge of θ_t contributes little in predicting the value of u_t. When Condition A is violated, the relatively high information costs for relatively low predictive ability from the private information means that no matter the number of uninformed traders and the magnitude of the downward bias in the uninformed market traders within the range $c_{1,t} \in [-c_1^*\Gamma_t, c_1^*)$, the informed traders cannot recoup the cost of their information.

The effects of $c_{1,t}$ and λ_t on $d\pi_t$ can be observed in Fig. 2(a) where $d\pi_t$ is plotted as a function of λ_t for a selection of $c_{1,t}$ values. The intersection of each curve with the horizontal axis indicates the value of λ_t^0 for the particular $c_{1,t}$.

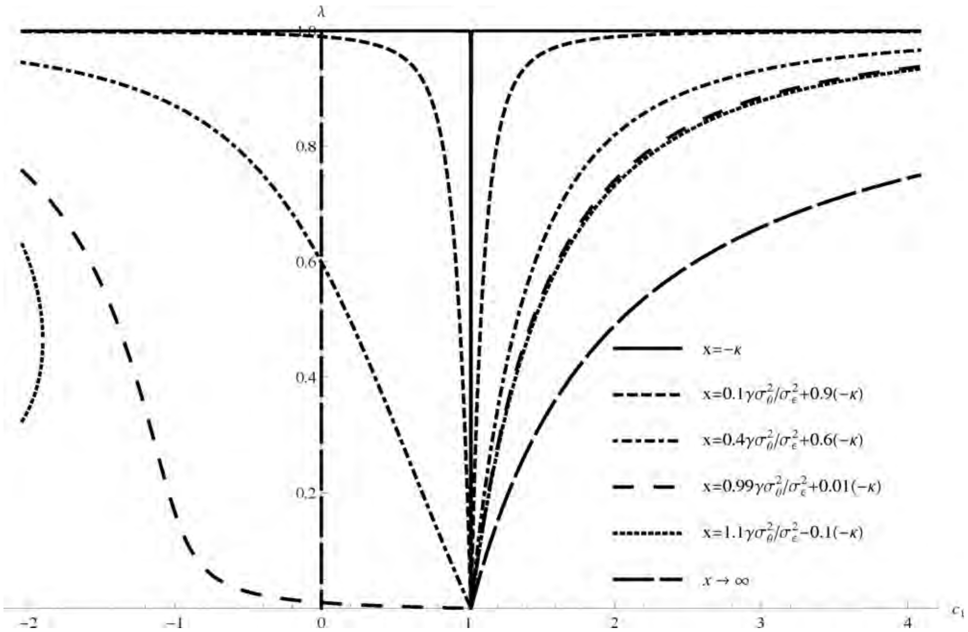

Fig. 1. Contour plot of $d\pi$. Each contour is $\lambda^x(c_{1,t})$ for a given $d\pi = x$. (The same figure can also be used to depict trading profits as a fraction of $\gamma\sigma_\theta^2/\sigma_\varepsilon^2$ by considering $\kappa = 0$.)

Fig. 2(b) plots λ^0 and λ^c as functions of c_1.

Each period, the traders must make a decision about whether to be informed or uninformed, a decision that would normally depend on $d\pi_t$ but without knowledge of λ_t, traders are not able to derive it analytically. The traders instead forecast performance using an average of the past realizations of profits,

$$\bar{\pi}_t^k = \bar{\pi}_{t-1}^k + (\pi_{t-1}^k - \bar{\pi}_{t-1}^k)/t, \bar{\pi}_0^k \text{ given}, \; k = I, U. \tag{26}$$

Let $d\bar{\pi}_t = \bar{\pi}_t^I - \bar{\pi}_t^U$. The evolution of $d\bar{\pi}_t$ according to (26) is the third and final dynamic process of the model. Its realization is an input into the population process. The value of $d\bar{\pi}_t$ has no direct impact on the state of the market.

In developing a performance measure, it is common in financial market models to limit the time horizon or place greater weight on more recent observations. As with the learning process, such limits would be counter to this paper's objectives of exploring the behavior of the market in the absence of arbitrary limits. In the GS model, it is unambiguous to the traders that the uninformed earn greater average profits in the presence of a revealing price. A cumulative performance measure allows the traders of the present model to converge to the same conclusion. As with the learning process, a limited memory artificially imposes a mechanism that can create a REE, this time by hindering the population's ability to learn of the superiority of the uninformed strategy. Limited use of past performance information is a source of instability in Brock and Hommes (1998) and many of the related examinations.

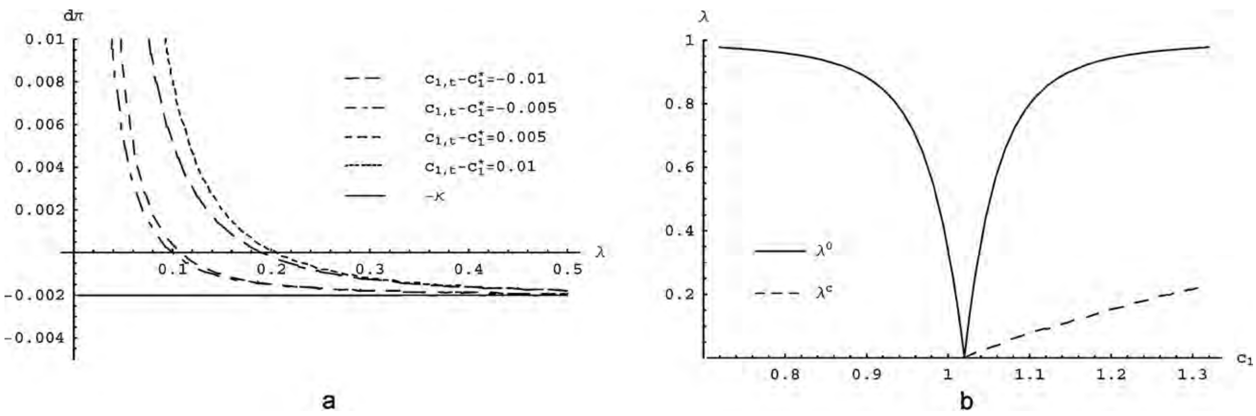

Fig. 2. (a) Plot of $d\pi$ as a function of λ given four different values for the model error, $c_1 - c_1^*$. $d\pi = g(c_1, \lambda^0) = 0$ at the intersection with the x axis. (b) Plot of λ^0 and λ^c as functions of c_1.

2.5. Population processes

Consider two families of population processes often employed in evolutionary discrete choice settings. Consider a population in which relative performance between strategy options produces a shift away from the inferior strategy or strategies towards the superior strategy or strategies. In such a setting, relative performance determines the innovation of λ_t. Such a process is distinct from a population that directly reflects relative performance. In the latter case, performance maps directly to the level of λ_t.

2.5.1. Innovation Population Dynamics

The *Innovation Population Dynamics* (IPD) process is characterized by the following:

Features of the IPD

(a) $\Delta\lambda_t = \lambda_t - \lambda_{t-1} = f(d\pi_t, \lambda_{t-1})$
(b) $f(0, \lambda_{t-1}) = 0$
(c) $f(x, \lambda_{t-1})$ is monotonically increasing in x and continuous
(d) $\lambda_t \in [0, 1]$, $\lambda_t \in (0, 1)$ for finite $d\pi_t$

Any strictly dominated strategy is driven towards negligible use as adherents defect. An interior fixed point to λ exists if there exists $\lambda \in (0, 1)$ that generates $d\pi = 0$. Such a population process is consistent with that considered by GS.

One convenient approach consistent with the IPD is to define

$$f(d\pi_t, \lambda_{t-1}) = \begin{cases} r(d\pi_t)(1-\lambda_{t-1}) \text{ for } d\pi_t \geq 0 \\ r(d\pi_t)\lambda_{t-1} \text{ for } d\pi_t < 0 \end{cases} \tag{27}$$

where $-1 < r(x) < 1$ for finite x, and $r(x)$ is monotonically increasing and continuous with $r(0) = 0$. This is the Branch and McGough (2008) form of the Replicator Dynamics (RD) process, an example of an IPD process. The Replicator Dynamic has its roots in evolutionary biology where a more fit population reproduces at a greater rate, changing the population proportion over time. In the adaptation to economics, including the Sethi and Franke (1995) application to a population facing a discrete choice, the movements in the population represent shifts in the population from an inferior to a superior strategy.

2.5.2. Population Level Dynamics

The *Level Population Dynamics* (LPD) process is characterized by the following:

Features of the LPD

(a) $\lambda_t = f(d\pi_t)$
(b) $f(0) = 1/2$
(c) $f(x)$ is monotonically increasing, continuous and has finite second derivatives (class C^2)
(d) $\lambda_t \in [0, 1]$, $\lambda_t \in (0, 1)$ for finite $d\pi_t$

In the LPD case, the function f maps directly from $d\pi_t$ to λ_t. A fixed point may exist at any value of $\lambda \in (0, 1)$, but the only fixed point value of λ that produces $d\pi = 0$ is at $\lambda = 1/2$. At any fixed point for which $\lambda \neq 1/2$, λ reflects the superior performance of the majority group. The Discrete Choice Dynamic (DCD) process is well established in the heterogeneous agent literature, having been introduced by Brock and Hommes (1997) as an adaptation of the randomized discrete choice model of Manski and McFadden (1981).

The DCD model assigns heterogeneity to the population of traders by randomizing the individual agent's choice. The process can be thought to capture either unmodeled idiosyncratic aspects of the individual's utility function or as actual randomness in the individual agent's decision process.[7]

The heterogeneous population explanation is based on the notion that unmodeled differences between traders leads to a distribution of perceived relative fitness for each strategy. The measure of relative fitness employed by the modeler, $d\bar{\pi}_t$ in this case, is the measure of relative performance at the center of the distribution of the population's beliefs, but is not unique among the population and, importantly, not necessarily uniquely rational.

The unmodeled aspects include different or additional information that the individual finds important for strategy evaluation. The strategy indicated to be superior according to the $d\bar{\pi}_t$ measure attracts a majority of the traders, but a tail of the distribution ranks the indicated inferior strategy to be the more attractive option. Heterogeneity leads traders to make a choice that is correct according to that individual trader's metric, but leaves a persistent population of underperformers according to the modeler's fitness measure.

Alternatively, with random preferences, a random idiosyncratic component is added to a shared non-stochastic utility measure, creating a distribution in preferences. The appeal of the random preference explanation for heterogeneity in choice

[7] Both aspects are discussed in McFadden (1981) in the original context of the randomized discrete choice model, and in De Fontnouvelle (2000) in a dynamic heterogeneous population setting.

is that no trader suffers from persistent underperformance. In each period some traders select the inferior strategy as a result of the random component in preference. The likelihood of this event is a decreasing function of $|d\bar{\pi}_t|$. For a stable $d\bar{\pi}_t$, the population proportions are stable as well, but in each period it is a different group of traders who randomly and temporarily employ the inferior strategy. No trader remains adherent to the inferior strategy.

One can think of the random use of the inferior strategy as individual experimentation. This is the approach explicitly modeled by Diks and Dindo (2008) who employ a parameter that captures the willingness of agents to experiment by employing the inferior strategy. The result is a persistent population of informed trades despite the strategy's inferior performance in an IPD type population process.

Branch and McGough (2008) offer a comparison of the DCD and RD processes. They demonstrate that the RD is able to produce chaotic population dynamics similar to those previously examined under DCD in an unstable cobweb model. The instability is rooted in the short memory of the agents when considering past performance. With short memory, the agents are unable to reach a consistent opinion of which of the two options offers long-run average superior performance, producing instability in the population proportion.

2.6. Asymptotic behavior and stability

To summarize the system which has been developed, there is a system of beliefs captured by the endogenously determined parameters and a system of state variables. Let Φ_t be a vector of the parameters, $\Phi_t = (c_{1,t} \quad Q_t \quad d\bar{\pi}_t)'$. Each element of Φ_t is updated by the recursive Eqs. (6), (7) and (26), respectively. Let Z_t be the vector of state variables, $Z_t = (u_t\ p_t\ \lambda_t)'$.

Of the state variables, only λ_t is useful for predicting the next state and it is hidden to the traders. In the LPD case, $\lambda_t = f(d\bar{\pi}_t)$, while for IPD $\lambda_t = f(d\bar{\pi}_t, \lambda_{t-1})$. The traders seek to forecast the value of u_t with the uninformed traders basing their forecast on p_t. Let W_t be a vector of the random determinates of u_t, $W_t = (\theta_t\ \varepsilon_t)'$, while vector A defines known constant determinates of u_t and p_t, $A = (\bar{u}\ \bar{u}/R)'$. The realization of $X_t = (u_t\ p_t)'$ can be expressed linearly in W_t as

$$X_t = A + B(\Phi_t, \lambda_t)W_t$$

with

$$B(\Phi_t, \lambda_t) = \begin{pmatrix} 1 & 1 \\ b_1(\Phi_t, \lambda_t) & 0 \end{pmatrix}.$$

The fixed point to a least-squares learning process is a point at which the perceived law of motion is consistent with the actual law of motion. The recursively determined parameters attain a fixed point only if the dynamic state equations produce evolution in the state variables that are consistent with the learning agents' beliefs.

To aid discussion, it is useful to introduce terminology specific to each of the dynamic processes. A "learning fixed point" refers to a fixed point to the learning process captured in (6) and (7), holding fixed the population parameter, λ_t. At the learning fixed point, the uninformed traders employ the correct model for extracting information on u_t from the observed p_t.

Similarly, a "population fixed point" refers to the a fixed point to the population process itself, $\lambda_t = \lambda^*$, to be developed below, while the parameters of the learning process are fixed.

A "system fixed point" is the traditional system wide fixed point in which, in this case, both a learning fixed point and population fixed point have been achieved.

2.6.1. System fixed point

Proposition 2 below establishes the learning fixed point. The nature of the population fixed point dependents on the population process as developed in Proposition 3.

Proposition 2. *$c_{1,t} = c_1^* = R$ is the learning fixed point.*

Proof. Given $\lambda_t = \lambda > 0$, the fixed point expressed in (12) and (15) follows from (10) and (13). $c_{1,t} = c_1^* = R$ is the fixed point to the learning process. Since c_1^* is independent of λ for $\lambda > 0$, $c_{1,t} = c_1^*$ is the fixed point for the learning process. From (10), for $\lambda = 0$, the price equation is undefined if $c_{1,t} = c_1^*$.[8] □

Proposition 3. *Presume that Condition A holds.*

A. *Given an IPD process*
  i. *For $c_{1,t} = c_1 \neq c_1^*$, $\exists \lambda_{IPD}^{fp}$, $\lambda_{IPD}^{fp} > \lambda^c(c_1)$, that is a fixed point to the population process.*
  ii. *$\lambda_{IPD}^{fp} \to 0$ as $c_1 \to c_1^*$.*
  iii. *For $c_{1,t} = c_1^*$, no fixed point exists to the population process.*

[8] The fixed point can also be derived from the associated ordinary differential equation of the learning process, which requires $b_{1,t}(1 - b_{1,t}c_{1,t}) = 0$. $b_{1,t} = 0$ is the solution for when p_t contains no information, as when $\lambda_t = 0$. For $\lambda_t \neq 0$, $b_{1,t} \neq 0$ and $b_{1,t}c_{1,t} = 1$ implies $c_{1,t} = R$.

B. *Given an LPD process*

i. *For* $c_{1,t} = c_1$, $\exists \lambda^{fp}_{LPD}(c_1)\ \forall c_1$, *that is a fixed point to the population process.*

ii. *For* $c_{1,t} = c_1^*$, $\lambda^{fp}_{LPD} = \lambda^{fp}_{LPD}(c_1^*) \in (0, 1/2]$ *based on* $d\pi(c_1^*, \lambda) = -\kappa$.

Proof of 3A. By Features of the IPD(a), a fixed point in the IPD process requires λ such that $f(g(c_1, \lambda), \lambda) = 0$. By IPD(b) this is obtained from $g(c_1, \lambda) = 0$. By Proposition 1, for a fixed $c_1 \neq c_1^*$ there exists $\lambda^0 > \lambda^c(c_1)$ producing $d\pi = 0$. $\lambda^{fp}_{IPD} = \lambda^0$ is thus a fixed point to the population process. (ii) follows directly from Proposition 2. Using (20) and L1.6 of Lemma 1 for (iii), when $c_{1,t} = c_1^*$, $\lambda_t - \lambda_{t-1} = f(-\kappa, \lambda_{t-1})$, and from Features IPD, $-\lambda_{t-1} < f(-\kappa, \lambda_{t-1}) < 0$, $\forall \lambda_{t-1} \in (0, 1]$.□

Proof of 3B. By the Features of the LPD(a) and Lemma 1, λ^{fp}_{LPD} is a fixed point in the population process if λ^{fp}_{LPD} solves $\lambda = f(g(c_1, \lambda))$. By LPD(c) and LPD(d) λ^{fp}_{LPD} is unique. $d\pi(c_1^*, \lambda) = -\kappa$ so Features LPD(b), LPD(c), and $\kappa > 0$ ensure $0 < \lambda^{fp}_{LPD} \leq 1/2$ for $c_{1,t} = c_1^*$.□

According to Proposition 3, with IPD, the existence of a fixed point to the population process depends on the presence of error in the uninformed traders' model. Proposition 3A parallels the GS findings. There is an equilibrium interior population proportion if the price is not perfectly revealing, as is the case with $c_{1,t} \neq c_1^*$. When $c_{1,t} = c_1^*$, the price is informationally efficient, in which case there is no equilibrium to the population process. The fixed points to the two processes do not coexist and thus no system fixed point exists.

A system fixed point does exist under LPD and it is unique. Proposition 3B(i) follows from the fact that f is monotonically increasing in $d\pi$ and g is monotonically decreasing in λ for $c_{1,t} \neq c_1^*$ and $c_{1,t} > -c_1^* \Gamma_t$. At the fixed point, the superior performance of the uninformed traders supports their majority position. The uninformed strategy outperforms paying to be informed, but in each period, in accordance with the discrete choice model, a group of traders chooses to be informed.

The LPD fixed point is consistent with Branch and Evans (2006) for whom the fixed mapping of the DCD between relative performance and the population proportion means that a stable finite difference in performance produces a stable population of traders in which both strategies are in use.

2.6.2. *Stability*

As developed below, in the LPD case, the asymptotic behavior of the system is characterized by the stability of the established fixed point. In the absence of a fixed point in the IPD case, an attractor for the system is demonstrated to exist and is employed to describe the system's behavior.

The termination of the risky asset at the end of each period eliminates the possibility of price bubbles and the associated drift in the parameters that can result from a trending price unhinged from its fundamental value. Consequently, the threat of non-convergence arises from the possibility that the interaction of the two processes results in cycles or chaotic behavior within the feasible range of the endogenous parameters, rather than from unbounded escape from fundamentals.

Lemma 2. *Given* $c_{1,t} = c_1^*$ *and* $\lambda_t > 0$, $d\bar{\pi}_t \to -\kappa$.

Proof. By (25) $d\pi_t = -\kappa$ for $c_{1,t} = c_1^*$ and $\lambda_t > 0$. By the LLN, $d\bar{\pi}_t \to d\pi$. (More trivially, for $\mathbf{c}_t = \mathbf{c}^*$ all traders are equally informed. There are no trades and no trading profits net of cost.)□

Proposition 4. *The unique fixed point of the LPD process is asymptotically locally stable.*

Proof. By substituting Eq. (9) for the state variable p_t and $\lambda_t = f(d\bar{\pi}_t)$ into recursive Eqs. (6), (7) and (26), the random dynamical system for the LPD process can be represented in the form of three equations for vector of parameters $\Phi_t = (c_{1,t} \quad Q_t \quad d\bar{\pi}_t)'$:

$$c_{1,t} = c_{1,t-1} + (Q_t^{-1}\theta_{t-1} b_1(c_{1,t-1}, d\bar{\pi}_{t-1})(u_{t-1} - \bar{u} - c_{1,t-1} b_1(c_{1,t-1}, d\bar{\pi}_{t-1})\theta_{t-1}))/t, \tag{28}$$

$$Q_t = Q_{t-1} + (\theta_{t-1} b_1(c_{1,t-1}, d\bar{\pi}_{t-1}))^2 - Q_{t-1})/t, \tag{29}$$

$$d\bar{\pi}_t = d\bar{\pi}_{t-1} + \left(\frac{1}{1 - f(d\bar{\pi}_{t-1})} (\bar{u} + \theta_t - R b_1(c_{1,t-1}, d\bar{\pi}_{t-1})\theta_t) \times (u_t - R b_1(c_{1,t-1}, d\bar{\pi}_{t-1})\theta_t)\phi^I - \kappa - d\bar{\pi}_{t-1} \right) / t. \tag{30}$$

Note that under BRE, $b_1(c_{1,t-1}, d\bar{\pi}_{t-1}) = f(d\bar{\pi}_{t-1})/(f(d\bar{\pi}_{t-1})R + (1 - f(d\bar{\pi}_{t-1})(R - c_{1,t-1}))$, under RE we use an explicit solution for $b_1(c_{1,t-1}, d\bar{\pi}_{t-1})$, but do not report it here, for brevity.

Following Evans and Honkapohja (2001), under regularity conditions verified in Appendix A, the stability of the above system of stochastic difference equations can be investigated by analyzing the stability of the associated system of ordinary differential equations (ODEs) given by:

$$\frac{d}{d\tau} \begin{bmatrix} c_1 \\ Q \\ d\bar{\pi} \end{bmatrix} = \begin{bmatrix} Q^{-1} b_1(c_1, d\bar{\pi})(1 - c_1 b_1(c_1, d\bar{\pi})\sigma_\theta^2 \\ b_1^2(c_1, d\bar{\pi})\sigma_\theta^2 - Q \\ \dfrac{1}{1 - f(d\bar{\pi})}(1 - R b_1(c_1, d\bar{\pi}))^2 \sigma_\theta^2 \phi^I - \kappa - d\bar{\pi} \end{bmatrix}. \tag{31}$$

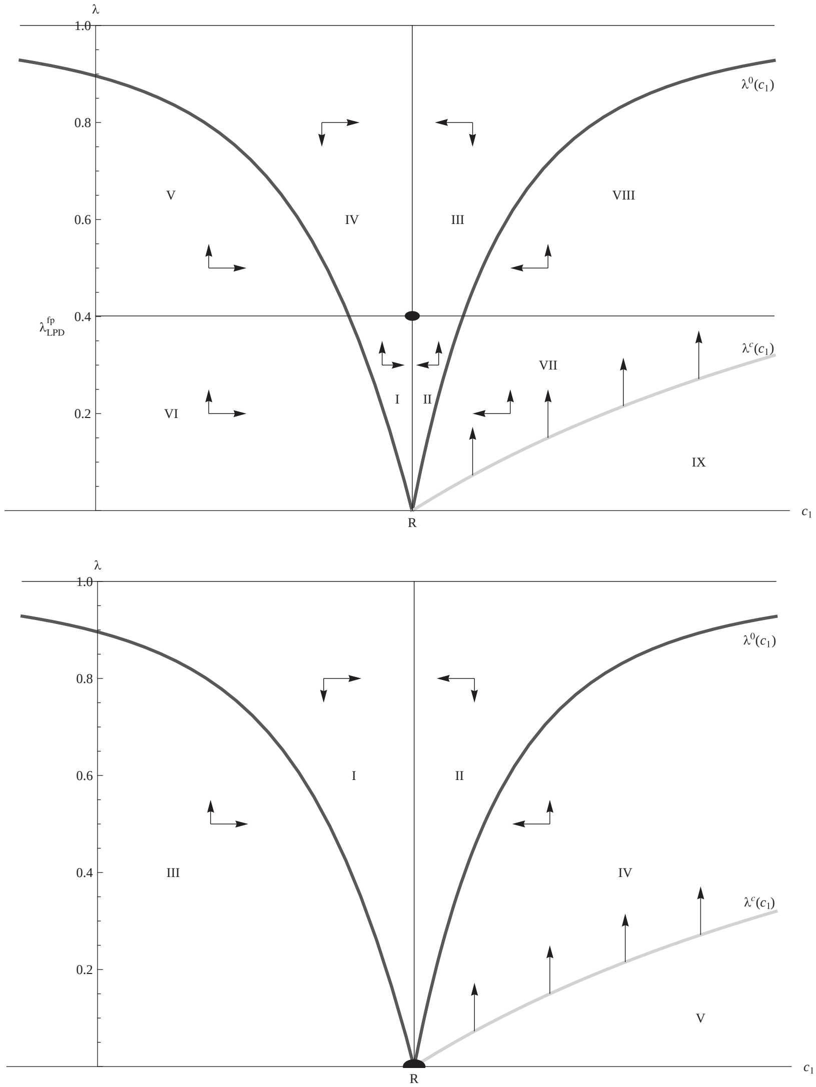

Fig. 3. (a) Phase space around the LPD fixed point. (b) Phase space around the IPD point of attraction.

Local stability of the above system of ODEs is analyzed by computing the eigenvalues of the corresponding Jacobian matrix J evaluated at fixed point $c_1^* = R$ and $d\bar{\pi}^* = -\kappa$. Provided that $f(-\kappa)$ and $f'(-\kappa)$ are finite, the Jacobian matrix J evaluated at the fixed point is equal to

$$J = \begin{pmatrix} -1/f(-\kappa) & 0 & 0 \\ \dfrac{2\sigma_\theta^2(1-f(-\kappa))}{R^3 f(-\kappa)} & -1 & 0 \\ 0 & 0 & -1 \end{pmatrix}.$$

Note that the solution is identical under BRE and RE. The corresponding eigenvalues are $(-1/f(-\kappa); -1; -1)$, which are all negative since $0 < f(-\kappa) < 1$. This condition is required for the local stability of the fixed point for the associated ODEs.□

The eigenvalues of the Jacobian establish the local stability of the fixed point. The stability is the natural outcome of a fixed point that is interior to the region above λ^0.

Global convergence can be obtained by restricting parameters Φ_t to the basin of attraction (cf. D set in Marcet and Sargent (1989a,b)). Fig. 3 is useful for discussing stability λ_t.[9] The basin of attraction for the LPD system is the full parameter space

[9] Because only two variables are present of the larger dynamic system, the two frames are not a complete phase space, but they do capture the dynamics at play. The missing dimension is the path through which λ_t influences $\partial\lambda_t/\partial t$. Directly, λ_t determines $d\pi_t$, influencing $d\bar{\pi}_t$ which determines $\partial\lambda_t/\partial t$.

excluding the region below λ^c in the BRE case (labeled region IX). Though the trajectories in region V and VIII point away from the fixed point, they remains within the basin of attraction because the increase in λ_t returns the system to a region above $\lambda^0(c_{1,t})$ which is again convergent. The learning that improves the accuracy of $c_{1,t}$ decreases instances in which $\lambda_t < \lambda^0(c_{1,t})$ asymptotically towards a zero probability event.

As a result of the unbounded support for ε_t, the possibility that a large negative realization of $d\pi_t$ could be sufficient to induce a realization of $\lambda_t < \lambda_t^c$ cannot be eliminated under BRE. In the context of the least-squares learning process, the constraint $\lambda_t > \lambda_t^c(c_{1,t})$ is, through inversion, a λ_t dependent upper bound on c_1. Without knowledge of λ_t, the traders are not in a position to impose limits on $c_{1,t}$. Still, it is worth observing that from (L1.5) the boundary $\lambda^c(c_{1,t})$ is naturally reflecting since $g(c_{1,t}, \lambda_t) \to \infty$ as $\lambda_t \to \lambda_t^c$ leading to a near certain increase in λ_t away from the boundary at $\lambda^c(c_{1,t})$. In this market, the population process accommodates for errors in the beliefs of the traders, generally returning the market to a more stable region of the parameter space.

The point of interest in the IPD system is $c_1 = R$ and $\lambda = 0$. It exists at the locus of the four regions defined by λ^-, λ^+, and λ^c as can be seen in frame (b) of Fig. 3. The inability to enclose the point within a stable region means that the point is not locally stable. As will be developed, system convergence derives from the interaction of the governing dynamic processes that maintains asymptotic convergence.

The dynamic system in the case of the IPD consists of four recursive equations for $\Phi_t = (c_{1,t} \quad Q_t \quad d\bar{\pi}_t)'$ and λ_t. The speed of recursive updating of $\Phi_t = (c_{1,t} \quad Q_t \quad d\bar{\pi}_t)'$ reduces over time (because of $1/t$ factor) while the speed of λ_t-updating does not change over time. That is the reason the analysis of this system conceptually differs from the traditional analysis of recursive algorithms used to prove the stability in the LPD case. A recursive learning system with replicator dynamics was recently considered by Guse (2010). The system is also related to the literature on dynamical systems with slow-fast dynamics or systems with multiple time scales (see, e.g., Lordon, 1997).

The IPD system is demonstrated to have a point of attraction ($\lambda^* = 0$, $c^* = R$) despite the absence of a fixed point. As will be developed, the evolution of λ_t is driven by three processes occurring at three different time scales. The slow improvement in the accuracy of $c_{1,t}$, as captured by a decrease in $(c_{1,t} - c_1^*)^2$, allows a slow convergence of $\lambda_t \to 0$. The interaction between the slow convergence in $d\bar{\pi}_t$ and the fast evolution in λ_t causes oscillating deviations in λ_t from its trending convergence at a medium time-scale. Finally, the random realizations of profits as determined by θ_t and ε_t introduce small random noise in λ_t in a fast time-scale.

At the slowest time-scale, $\lambda^0(c_{1,t})$ captures that accuracy in $c_{1,t}$ is a prerequisite for $\lambda_t \to 0$. For λ_t to decline requires $d\bar{\pi}_t < 0$. To have $d\bar{\pi}_t < 0$ requires $d\pi_t < 0$ which is obtained when λ_t is above $\lambda^0(c_{1,t})$. Since $\lambda^0(c_{1,t}) \to 0$ only as $c_{1,t} \to c_1^*$, λ_t is bounded away from zero for $c_{1,t} \neq c_1^*$.

To understand the oscillatory process in λ_t, consider the system at a starting point in which $d\pi_t = 0$ and $\lambda_t > \lambda^0(c_{1,t})$. The former means that $\Delta\lambda_t = 0$ while later means that $d\pi_t < 0$. Though the random realizations of profit differentials can differ from their expectations, the accumulation of realized profit differentials based on $d\pi_t < 0$ leads to $d\bar{\pi}_t < 0$, starting a process of general decline in λ_t. A convergence of $\lambda_t \to \lambda^0(c_{1,t})$ leads to an increase in $d\pi_t$ that eventually becomes positive when λ_t passes below $\lambda^0(c_{1,t})$. At $\lambda_t = \lambda^0(c_{1,t})$, $d\pi_t = 0$ but as a cumulative average with $1/t$ updating, $d\bar{\pi}_t$ remains positive so that λ_t continues to decline. Once $\lambda_t < \lambda^0(c_{1,t})$, the accumulation of $d\pi_t > 0$ realizations cause increase in $d\bar{\pi}_t$, which slows the downward evolution in λ_t. The evolution in λ_t reverses and increases once $d\bar{\pi}_t$ becomes positive. Again, because of the slow-fast process, the evolution in λ_t is faster than the evolution in $d\bar{\pi}_t$. As λ_t increases, it eventually becomes greater than $\lambda^0(c_{1,t})$, generating $d\pi_t < 0$ but $d\bar{\pi}_t$ remains positive until the accumulation of negative $d\pi_t$ realizations is sufficient to lower $d\bar{\pi}_t$ to zero. While $d\bar{\pi}_t > 0$, λ_t continues to rise, though it is in the $\lambda_t > \lambda^0(c_{1,t})$ region. Eventually the accumulation causes $d\bar{\pi}_t$ to turn negative, reversing the direction of λ_t.

Features of the oscillations:

1. The distribution of $d\pi_t$ is asymmetric, with a lower bound of $-\kappa$ as determined by the cost of information. This lower bound is realized when $c_{1,t} = c_1^*$ or $\lambda_t = 1$. The expected profit differential is unbounded from above as $\lambda_t \to \lambda^c(c_{1,t})$ for $c_{1,t} > c_1^*$. As a result, the rate of change in λ_t tends to be slow during the downward portion of the oscillation. Once $\lambda_t < \lambda^0(c_{1,t})$, relatively large magnitude positive realizations quickly reverse the trajectory of λ_t so that only a small number of periods are spent with $d\bar{\pi}_t > 0$. Overall, the decline in λ_t is slow and drawn out while it increases relatively quick.
2. The oscillation in λ_t is inevitable and unavoidable. An uninterrupted decent in λ_t towards zero requires improvement in the accuracy of $c_{1,t}$ sufficient to maintain a path of λ_t convergence that remains between $\lambda^-(c_{1,t})$ and $\lambda^+(c_{1,t})$ as $\lambda_t \to 0$. This path is eliminated by the asymptotic slow convergence in $c_{1,t}$ and the fast evolution in λ_t.
3. Each cycle of oscillation includes an instance of λ_t crossing $\lambda^0(c_{1,t})$ both from above and from below. When the fast-moving λ_t crosses $\lambda^0(c_{1,t})$, $d\pi_t$ changes sign, but the change in sign in the slow moving $d\bar{\pi}_t$ occurs only with sufficient accumulation of $d\pi_t$.

The random realizations of θ_t and ε_t mean that the realized profit differential need not match the expectation according to $d\pi_t$. The affect on λ_t is filtered through the $1/t$ updating of $d\bar{\pi}_t$, but it does mean that each time step innovation in λ_t is randomly determined. The same random process also drives the value of $c_{1,t}$ so that with each time step, the value of $\lambda^0(c_{1,t})$ reflects the value of the randomly determined $c_{1,t}$. The general time and location at which λ_t crosses $\lambda^0(c_{1,t})$ is largely driven by underlying dynamics of the state variables, but is also affected by the random realizations of θ_t and ε_t.

Proposition 7 below is established using the horizontal distances in Fig. 2(b) between $\lambda^-(c_{1,t})$, c_1^*, $\lambda^+(c_{1,t})$, and $\lambda^c(c_{1,t})$. Since both $\lambda^+(c_{1,t})$ and $\lambda^c(c_{1,t})$ are monotonically increasing functions and under Condition A $\lambda^-(c_{1,t})$ is monotonically decreasing, all can be inverted, expressing c_1 as a function of λ. Let φ represent the horizontal distance between $\lambda^c(c_{1,t})$ and $\lambda^+(c_{1,t})$ in proportion to the horizontal distance $\lambda^+(c_{1,t})$ and $\lambda^-(c_{1,t})$. Inverting $\lambda^-(c_{1,t})$, $\lambda^+(c_{1,t})$, and $\lambda^c(c_{1,t})$ to express c_1^-, c_1^+, and c_1^c as respective functions of λ yields

$$c_1^+ = R\frac{(1-\lambda)(\gamma\sigma_\theta^2\phi^U - \kappa\sigma_\varepsilon^2(\lambda\phi^I + (1-\lambda)\phi^U) + \lambda\sigma_\varepsilon\sigma_\theta\phi^I\sqrt{\gamma\kappa(1-\lambda)}}{(1-\lambda)(\gamma\sigma_\theta^2 - \kappa(1-\lambda)\sigma_\varepsilon^2)\phi^{U2}} \tag{32}$$

$$c_1^- = R\frac{(1-\lambda)(\gamma\sigma_\theta^2\phi^U - \kappa\sigma_\varepsilon^2(\lambda\phi^I + (1-\lambda)\phi^U) - \lambda\sigma_\varepsilon\sigma_\theta\phi^I\sqrt{\gamma\kappa(1-\lambda)}}{(1-\lambda)(\gamma\sigma_\theta^2 - \kappa(1-\lambda)\sigma_\varepsilon^2)\phi^{U2}} \tag{33}$$

$$c_1^c = R/(1-\lambda). \tag{34}$$

In the BRE case, (32) and (33) simplify to

$$c_1^+ = R\frac{(1-\lambda)(\gamma\sigma_\theta^2 - \kappa\sigma_\varepsilon^2) + \lambda\sigma_\varepsilon\sigma_\theta\sqrt{\gamma\kappa(1-\lambda)}}{(1-\lambda)(\gamma\sigma_\theta^2 - \kappa(1-\lambda)\sigma_\varepsilon^2)}$$

$$c_1^- = R\frac{(1-\lambda)(\gamma\sigma_\theta^2 - \kappa\sigma_\varepsilon^2) + \lambda\sigma_\varepsilon\sigma_\theta\sqrt{\gamma\kappa(1-\lambda)}}{(1-\lambda)(\gamma\sigma_\theta^2 - \kappa(1-\lambda)\sigma_\varepsilon^2)}.$$

From these inverted expressions can be derived,

$$\varphi(\lambda_t) = \frac{c_1^c(\lambda_t) - c_1^+(\lambda_t)}{c_1^+(\lambda_t) - c_1^-(\lambda_t)} = \frac{1}{2}\left(\Theta\frac{\phi^U}{\phi^I} + \Theta^{-1}\frac{\phi^I - \phi^U}{\phi^I} - 1\right) \tag{35}$$

with $\Theta = \sigma_\theta\sqrt{\gamma}/(\sigma_\varepsilon\sqrt{(\kappa(1-\lambda))})$. In the BRE case, (35) simplifies to

$$\varphi(\lambda_t) = \frac{1}{2}\left(\Theta - 1\right).$$

Under Condition A, $\Theta - 1 > 0$. Under RE, Condition A ensures $\varphi > 0$ if $\phi_t^U \leq 2\phi^I$ which is equivalent to $(1 - c_{1,t}b_{1,t})^2\sigma_\theta^2 \leq \sigma_\varepsilon^2$.

Proposition 5. *Given Condition A, in finite time, there exist $\xi > 0$ such that for $\lambda_t < \xi$ $d\pi_t > 0$.*

Proof. Follows from $(c_{1,t} - c_1^*)^2 > 0$ in finite time and $\lim_{\lambda\to 0}(c_1^+(\lambda) - c_1^-(\lambda)) = 0$.□

Proposition 6. *Given Condition A and $\lambda_t > 0$ there exists a not inconsequential distance between $c_1^c(\lambda_t)$ and $c_1^+(\lambda_t)$ and between $c_1^+(\lambda_t)$ and $c_1^-(\lambda_t)$.*

Proof. For $\kappa < \gamma\sigma_\theta^2/\sigma_\varepsilon^2$, $\lim_{\lambda\to 0}(\varphi(\lambda)) = k$ with $0 < k < \infty$.□

Proposition 7. *Given Condition A, all three time-scale convergence properties continue as $t \to \infty$.*

Proof. $\partial\varphi/\partial\lambda > 0$, indicating that φ decreases with a decline in λ but is bounded away from zero in the limit. By Proposition 6, there exists a buffer region generating $d\pi > 0$ that is of measureable size relative to the region generating $d\pi < 0$. This region ensure that during conversion, as $(c_{1,t} - c_1^*)^2 \to 0$, each incursion of $\lambda_t < \lambda^0(c_{1,t})$ occurs closer to zero, $d\pi_t < \infty$ remains possible. Let c_1^x be the value of $c_1 > c_1^*$ solving $d\pi = x$ given λ (so that c_1^+ is just the special case for $x = 0$). A byproduct of $\partial\varphi/\partial\lambda > 0$ is that as $\lambda \to 0$, for $x > 0$, $c_1^x - c_1^+/c_1^c - c_1^+$ increases to a bounded limit. Thus, as incursions of $\lambda_t < \lambda^0(c_{1,t})$ occur closer to zero, the proportion of the horizontal distance between $\lambda^+(c_1)$ and $\lambda^c(c_1)$ contributing to $0 < d\pi \leq x$ increases. This contributes to a slower process of reversal in the direction of λ_t and a slower rate of exit as it returns to the region above $\lambda^0(c_{1,t})$. Thus, each oscillation as $(c_{1,t} - c_1^*)^2$ decreases is of diminished magnitude and extremes closer to zero. Consider as proof the alternative in which $\varphi(\lambda) \to 0$ as $\lambda \to 0$, then as the accuracy in $c_{1,t}$ improves and λ_t crossed below $\lambda^0(c_{1,t})$ close to zero, there would be a decreasing buffer offering $0 < d\pi_t < x$ for all x, $0 < x < \infty$. In this case, as $(c_{1,t} - c_1^*)^2 \to 0$ the realization of $\lambda_t < \lambda^0(c_{1,t})$ generates $d\pi_t \to \infty$ which results in a sustained increase in λ_t. $\lim_{\lambda\to 0}(\varphi(\lambda)) = 0$ generates a process by which improvement in the accuracy of $c_{1,t}$ generates oscillations with increasing extremes in λ_t between zero and one. A $\lim_{\lambda\to 0}(\varphi(\lambda)) = 0$ also invites realization in the invalid region of $\lambda_t < \lambda^c(c_{1,t})$, resulting in a breakdown in the market.□

It is most natural to consider the evolution of the system with $\lambda^0(c_{1,t})$ in the shape of a cone, as depicted in Fig. 3(b). Early in the learning process when the error in $c_{1,t}$ is large, λ_t will cross $\lambda^0(c_{1,t})$ while still relatively far from zero. As the accuracy in $c_{1,t}$ improves, λ_t will achieve a value closer to zero before crossing $\lambda^0(c_{1,t})$. It is the funnel shape of $\lambda^0(c_{1,t})$ that creates the environment by which Proposition 7 is true. The probability of crossing $\lambda^c(c_{1,t})$ declines with the passage of time because the system is not allowed to enter the region where the horizontal distance between $\lambda^+(c_{1,t})$ and $\lambda^c(c_{1,t})$ is vanishingly small until $E(c_{1,t} - c_1^*)^2$ is itself vanishingly small.

This is no longer true when Condition A does not hold. When Condition A does not hold, values of $c_{1,t} \in [-c_1^* \Gamma_t, c_1^*)$ produce $d\pi_t < 0 \; \forall \lambda > 0$. Thus, for $c_{1,t} < c_1^*$, the system can become arbitrarily close to $\lambda_t = 0$ even when $E(c_{1,t} - c_1^*)^2$ remains large. The eventual and inevitable switch in $c_{1,t}$ to $c_{1,t} > c_1^*$ can generate an unbounded positive $d\pi_t$. The emergence of diminishing oscillation in λ_t is not assured when Condition A does not hold.

As with the LPD case, the stability can only be considered locally. The unbounded support for ε_t keeps the probability of landing the system under BRE below $\lambda^c(c_{1,t})$, and thus outside of the supported region D, non-zero. The decrease in φ with a decrease in λ_t, the same feature that ensures convergence under Condition A, also contributes to decreasing likelihood of exiting the supported parameter space with the passage of time. As noted, the when Condition A is not satisfied, the point at $\lambda = 0$ and $c_1 = R$ cannot be considered an attractor. By arguments similar to those employed in the proof of Proposition 7(b), as $\kappa \to \gamma\sigma_\theta^2/\sigma_\varepsilon^2$ from below, Condition A remains satisfied, but the probability of escape from D is increased.

3. Market efficiency

3.1. Analysis

This section examines the implications for market asymptotic efficiency under the IPD and LPD regimes. The market is efficient when the price fully reflects all the available information, in this case the private signal θ_t. Let p_t^{EM}, indicate the efficient market price. The efficient market price is the rational expectations equilibrium price,

$$p_t^{EM} = R^{-1}\theta_t.$$

Note that with $b_1^* = 1/R$ producing $p_t^* = p_t^{EM}$ market efficiency is achieved if $c_{1,t} = c_1^*$ or if $\lambda_t = 1$. Efficiency is thus embedded in the model. It is natural to think of the pricing error, $\left|p_t - p_t^{EM}\right|$, as a measure of market efficiency, but this distance is dependent on not just the parameters of the price equation, but also the realization of θ_t. For this reason, a convenient and more relevant measure is the variance of this distance, $\sigma_{p,t}^2 = \text{var}_t(p_t - p_t^{EM})$,

$$\sigma_{p,t}^2 = (b_{1,t} - b_1^*)^2\sigma_\theta^2, \tag{36}$$

which can also be expressed as $\sigma_{p,t}^2 = R^{-2}(1 - Rb_{1,t})^2\sigma_\theta^2$, revealing $d\pi_t$ in (24) to be a λ_t dependent multiple of $\text{var}_t(p_t - p_t^{EM})$.

Recall from Proposition 2 that λ_t^0 reflects the accuracy of the uniformed traders' model, converging towards zero as the model becomes increasingly accurate. The closer λ_t^0 is to zero the greater the market's ability to absorb uninformed traders without causing substantial mispricing. The accuracy of the market price thus depends on λ_t relative to λ_t^0. The innovations to λ_t, in turn, are driven by profits as determined by price accuracy.

Proposition 8. *Under IPD and Condition A, the error in the market price converges asymptotically to a fixed positive value.*

Proof. Using (37) in (24) and solving for $\sigma_{p,t}^2$ yields

$$\sigma_{p,t}^2 = (d\pi_t + \kappa)(1 - \lambda_t)\sigma_\varepsilon^2/R^2\gamma. \tag{37}$$

From Proposition 7(a), $d\pi_t$ oscillates around zero so that $\sigma_{p,t}^2$ oscillates around $\kappa(1 - \lambda_t)\sigma_\varepsilon^2/R^2\gamma$ which, with the convergence of $\lambda_t \to 0$, becomes simply $\kappa\sigma_\varepsilon^2/R^2\gamma$. Since the oscillations in $d\pi_t$ diminish with the system's convergence, so too does the oscillation in $\sigma_{p,t}^2$.□

Proposition 9. *The LPD process ensures* $\sigma_{p,t}^2 \to 0$.

Proof. Because of existence of asymptotically stable fixed point under LPD (Proposition 4), λ_t converges to $\lambda_{LPD}^{fp} > 0$, and $b_{1,t}$ converges to b_1^* so that $\sigma_{p,t}^2 \to 0$.□

As the model improves under LPD, the environment for convergence in learning approaches the fixed λ environment of Bray (1982).

3.2. Simulations

Simulations illustrate the asymptotic behavior of the model and give character to market behavior under LPD and IPD. Under LPD, the asymptotic behavior of the market is unaffected by whether traders are presumed to behave according to BRE or RE. Under RE, there is nothing to prevent convergence to the fixed point. The BRE solution with $\phi_t^U = \phi^I$ means that there exist combinations of $c_{1,t}$ and λ_t producing $\Psi(c_{1,t}, \lambda_t) \le 0$, and thus the absence of a reasonable market clearing price. For BRE under the LPD, the probability of realizing $\Psi(c_{1,t}, \lambda_t) \le 0$ converges to zero as the system stabilized around the fixed point. If the simulation survives the early stages of learning and population evolution, then convergence is very likely.

The simulations under the IPD reveal the three components of the noisy oscillatory convergence of λ_t towards zero. The pricing error cycles with the oscillations in λ_t, but does not improve with the improvements in the model accuracy.

Fig. 4 through Fig. 6 demonstrate the different behavior produced the LPD and IPD. Each figure contains six frames plotting sample times-series of select endogenous parameters produced by simulations of the model. The left column displays, from

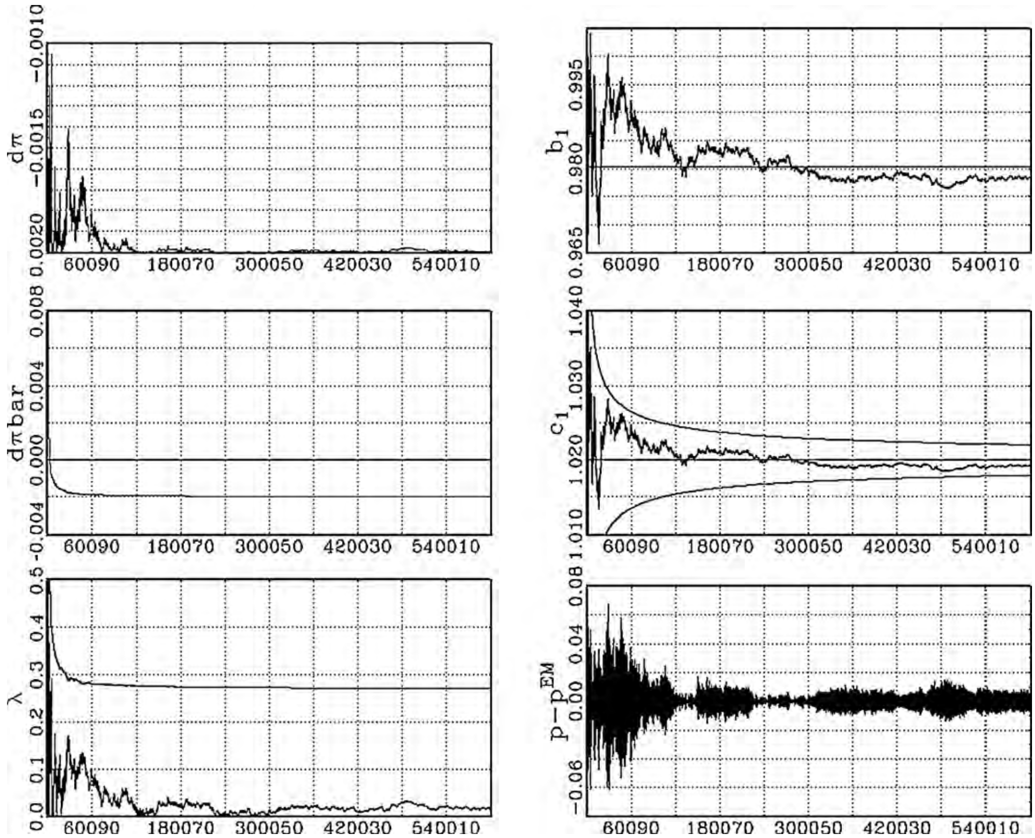

Fig. 4. Population governed by Discrete Choice Dynamics (of the LPD family). T = 600,000 periods. Column 1: Row 1: Time t expected profits, $d\pi_t$; Row 2: cumulative performance measure, $d\bar{\pi}_t$; Row 3: λ_t (black) and λ_t^0 (grey). Column 2: Row 1: $b_{1,t}$; Row 2: $c_{1,t}$ with 90 percent confidence bands; Row 3: $p_t - p_t^{EM}$. R = 1.02, κ = 0.002.

top to bottom, the state-dependent expected profit differential, $d\pi_t$, the cumulative performance measure, $d\bar{\pi}_t$, and a plot of λ_t (black) and λ_t^0 (grey). The right column displays, from top to bottom, $b_{1,t}$, $c_{1,t}$, and the deviations from market efficiency, $p_t - p_t^{EM}$. Included in the plot of $c_{1,t}$ are the classic regression model 90 percent confidence bands computed from the parameter's asymptotic variance, $\sigma_{c_1}^2 = R^{-2}\sigma_\theta^2/t$.

Fig. 4 plots a typical DCD based simulation. Each trader elects to be informed with probability $\Pr_t(x_{i,t} = I)$. For a large population of traders, the Law of Large Numbers applies. From the original discrete choice model of Manski and McFadden,

$$\lambda_t \to \Pr_t(x_{it} = I) = \frac{\exp(\rho\bar{\pi}_t^I)}{\exp(\rho\bar{\pi}_t^I) + \exp(\rho\bar{\pi}_t^U)} = (\tanh(\rho(d\bar{\pi}_t)/2) + 1)/2. \tag{38}$$

Consider a population sufficiently large such that the relationship can be treated as an identity, $\lambda_t = \Pr_t(x_{i,t} = I)$. For $d\bar{\pi}_t$ near zero, there is little perceived difference and the populations are of nearly equal in proportion. The greater the difference, the smaller the proportion of traders who use the inferior strategy in that period. The parameter $\rho \geq 0$ sets the population's "intensity of choice", defining how sensitive the population is to the difference in performance between the options.

The figure shows the convergence in learning where $c_1^* = R = 1.02$ and convergence in the population to the fixed point of $\lambda_{LCD}^{fp} = 0.2689$. The learning process progresses towards an increasingly accurate model, producing a convergence towards market efficiency in the price. In the latter part of the sample, price efficiency tracks the accuracy of $c_{1,t}$, which remains well within the 90 percent confidence bands. The error introduced by the uninformed traders diminishes as the traders' model improves.

The examination of the Replicator Dynamic employs the function

$$r(d\bar{\pi}_t) = \tanh(\delta d\bar{\pi}_t/2) \tag{39}$$

for use in (27). The parameter $\delta \geq 0$ sets the strength of the populations' response to disparity in the perceived performance of the two options.

The RD simulation plotted in Figs. 5 and 6 employs the same payoff stream for the random security as employed in the DCD simulation plotted in Fig. 4. As seen in Fig. 6, a pattern of oscillations in λ_t overlays its convergence towards zero. The values of $d\pi_t$, $d\bar{\pi}_t$, and $p_t - p^{EM}$ all cycle in accordance with the oscillations in λ_t. Notice the asymmetry in $d\pi_t$ with

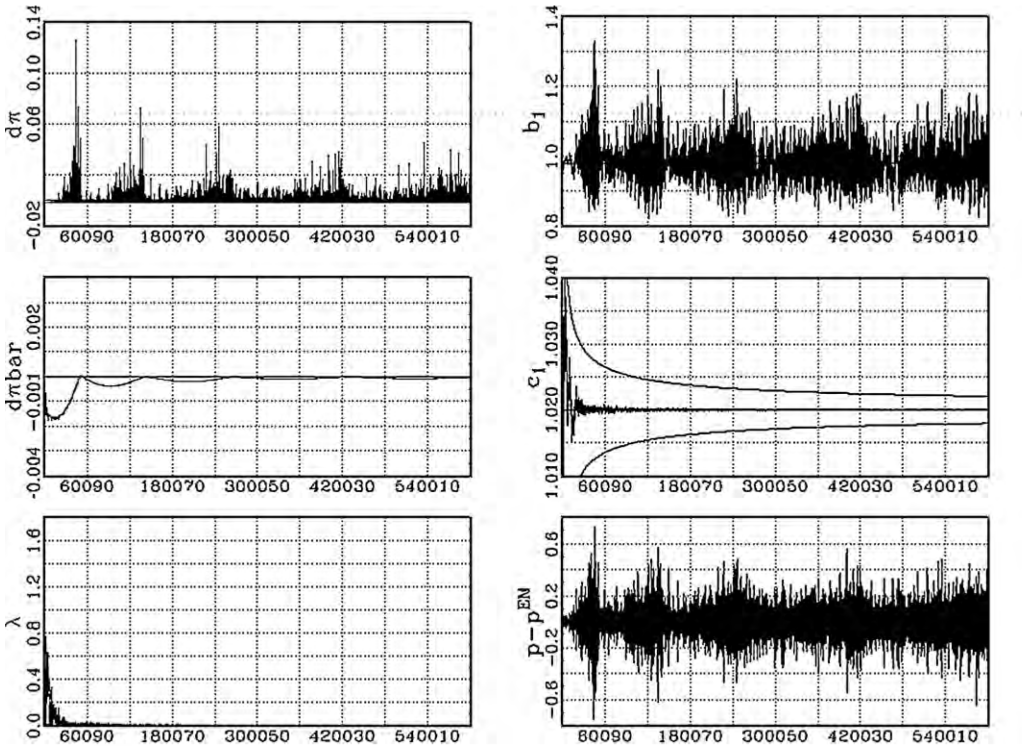

Fig. 5. Population governed by Replicator Dynamics (of the IPD family). $T = 600{,}000$ periods. Column 1: Row 1: Time t expected profits, $d\pi_t$; Row 2: cumulative performance measure, $d\bar{\pi}_t$; Row 3: λ_t (black) and λ_t^0 (grey). Column 2: Row 1: $b_{1,t}$; Row 2: $c_{1,t}$ with 90 percent confidence bands; Row 3: $p_t - p_t^{EM}$. $R = 1.02$, $\kappa = 0.002$, $\rho = 0.1$.

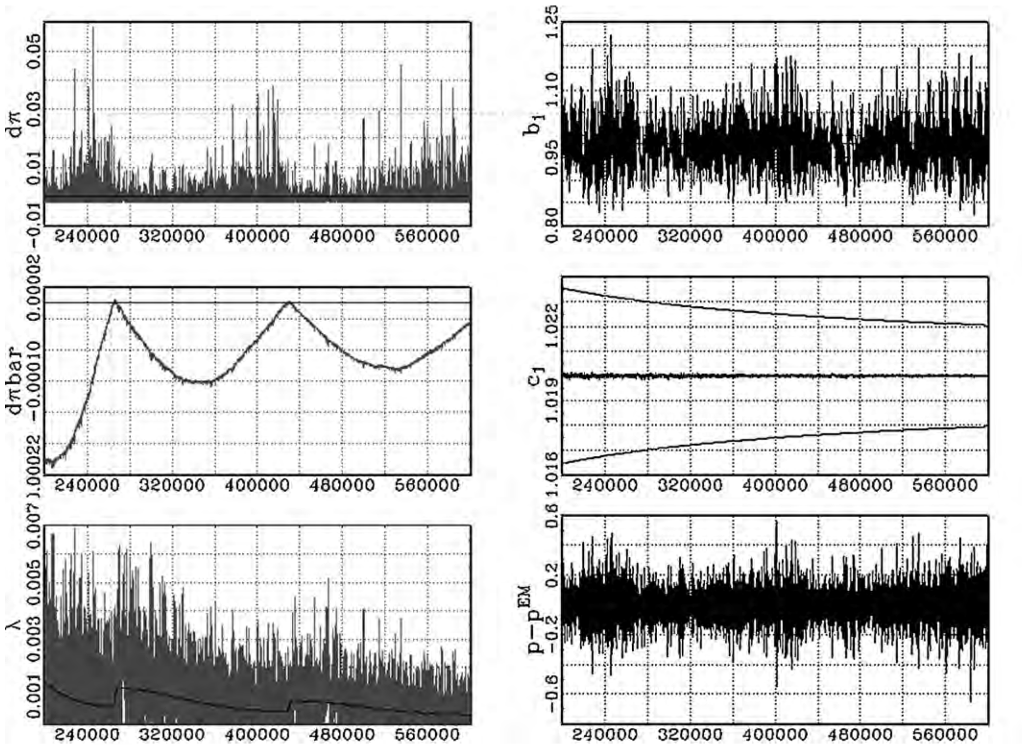

Fig. 6. Population governed by Replicator Dynamics (of the IPD family). $t = 200{,}000$–$600{,}000$. Column 1: Row 1: Time t expected profits, $d\pi_t$; Row 2: cumulative performance measure, $d\bar{\pi}_t$; Row 3: λ_t (black) and λ_t^0 (grey). Column 2: Row 1: $b_{1,t}$; Row 2: $c_{1,t}$ with 90 percent confidence bands; Row 3: $p_t - p_t^{EM}$. $R = 1.02$, $\kappa = 0.002$, $\rho = 0.1$.

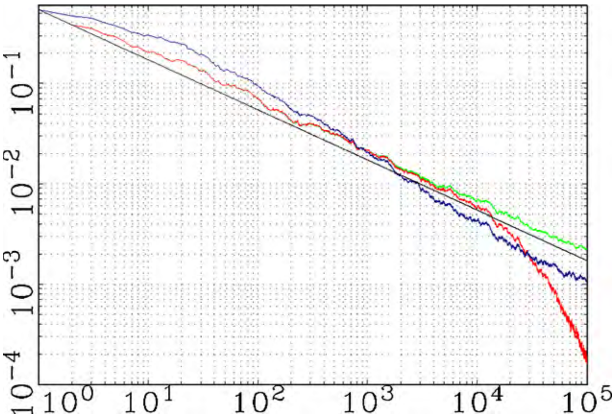

Fig. 7. The average $|c_{1,t} - c_1^*|$ across 100 simulations plotted against time. Light grey = constant λ_t, dark grey = DCD (a LPD example), grey = RD (a IPD example). (In the color web version of the article, these are green, blue, and red, respectively.)

$-\kappa \leq d\pi_t < \infty$ which influences the oscillations in λ_t. The λ_t^0 series is noisy as a result of the random element of $c_{1,t}$, but as the frequency and depth with which $\lambda_t < \lambda_t^0$ is realized increases, so rises the frequency and magnitude of the realizations of $(\pi_t^I - \pi_t^U) > 0$. Over time, the positive realized profits contribute to the slow rise in $d\bar{\pi}_t$ until $d\bar{\pi}_t > 0$ is realized and reverses the direction of λ_t. Notice that each subsequent oscillation in λ_t starts at a lower peak and ends at a lower trough, while the peaks and troughs in the deviation of $p_t - p^{EM}$ remain constant across cycles.

4. Rate of learning

Comparing the evolution in $c_{1,t}$ in Figs. 4 and 5 suggests that the rate of learning is different under the RD population process than under the DCD population process. The relative values of λ_t and λ_t^0 impact the rate of learning, producing the different simulation outcomes.

Underlying Fig. 7 are 100 iterations of 100,000 periods each. Three versions of the model are applied to each of these series. All three models start with $\lambda_0 = 0.8$. The base model has a fixed $\lambda_t = \lambda = 0.8$ so that the learning process is the only dynamic process in the simulation. The second set of simulations is based on the DCD model. The final set of simulations is based on the RD model. The average of the distance $|c_{1,t} - c_1^*|$ taken over the 100 simulations for each model is plotted against time. In the log–log scale the solid straight line is a plot of $1/\sqrt{t}$. The three jagged lines are the output from the three models.

The convergence of the base model runs parallel to the plot of $1/\sqrt{t}$, indicating the same rate of convergence. The convergence of the learning process of both the DCD and RD models diverge from the base model. The most pronounced example of this occurs at the end of the RD simulations. Learning is accelerated by periods of large pricing errors, which occur when λ_t approaches or drops below λ_t^0. In comparison, the increased rate of learning in the DCD model simulations are short lived. Learning is accelerated early in the simulations when the price exhibits increase volatility relative to the fixed λ simulations, a time when λ_t is near $\lambda^0(c_{1,t})$. Once λ_t stabilizes at the fixed point, λ_{LPD}^{fp} while $\lambda^0(c_{1,t}) \to 0$ with the improvements in $c_{1,t}$, learning slows to a rate of $1/t$. Very early in the simulation, learning in the DCD setting is hampered by the large innovations in λ_t from period to period that substantially alter the relationship between price and payoff. The rate of learning recovers once λ_t settles with the accumulation of data.

Proposition 7 was established based on the rate of the learning presumed to be unaffected by the market environment. Here, it is discovered that the rate at which the traders learn $c_{1,t}$ increases when the pricing errors are large. This has the potential to decrease the oscillations in λ_t, allowing the system to emerge from below λ_t^0 through improvement in $c_{1,t}$ rather than an increase in λ_t.

5. Conclusion

The asymptotic convergence properties of a financial market with learning and adaptation have been derived analytically. The analysis reveals that, having explicitly modeled the dynamic processes and accounted for their interaction; market efficiency is possible, the lack of a fixed point need not be seen as a paradox, and a non-revealing price can be modeled as endemic to the market rather than the product of an *ad hoc* imposition of noise. The process of population dynamics matters to the solution. The Innovation Population Dynamic (IPD) generates convergence in both learning and the population towards an attractor, but the dual convergence is demonstrated to generate persistent pricing errors. The Level Population Dynamic (LPD) produces convergence towards a fixed point that allows for the emergence of an efficient price.

Market efficiency arises from the Level Population Dynamic process because it produces a persistent population of informed traders, despite their underperformance relative to the uninformed traders. The population process is different from that envisioned by Grossman and Stiglitz but, in the case of the Discrete Choice Dynamics, the process has been employed extensively in the dynamic choice literature based on a strong empirical and theoretical foundation. The Discrete Choice model offers a couple of intuitive explanations of how such a group of trader can remain present, but other models capable of generating a persistent population of informed traders may exist as well, offering their own explanations.

The original paradox observed by Grossman and Stiglitz is a product of the discontinuity in expected profits that arises in the absence of informed traders. This results in the absence of a Rational Expectations Equilibrium. Given that the uninformed traders have the correct model for extracting the private signal from the price, no barrier exists to prevent the entire population of traders from attempting to adopt the uninformed strategy, achieving the point of discontinuity. The Grossman and Stiglitz model resolve the paradox by creating an equilibrium away from the point of discontinuity by injecting noise into the price.

Like Grossman and Stiglitz, the Innovation Population Dynamics generates an environment in which, asymptotically, the entire population converges towards full adoption of the uninformed strategy. Contrary to Grossman and Stiglitz, in the presence of the learning process the discontinuity at the point of attraction ceases to be a problem in need of a solution. Importantly, the rate at which the traders adopt the uninformed strategy is tied to the rate at which the traders improve their understanding of the market. In the presence of error, excessive reliance on the uninformed strategy produces a pricing error that maintains a population of informed traders. As the uninformed model improves, the proportion of informed traders declines, but the informed strategy cannot be completely abandoned while error exists in the uninformed traders' model, even as the error declines to zero.

The system exists and operates out of equilibrium. It is not necessary to create an equilibrium for the population process through assumption. Error is introduced into the price by the uninformed traders' use of an imperfect model. The population process tunes the market to maintain a persistent error in the price despite the diminishing error in the model. The process is a financial market version of the Malthusian Trap. Increasing adoption of the increasingly accurate model for extracting information from the price produces a persistent pricing error. The pricing error creates the space necessary for a population of informed traders to maintain a profitable presence in a competitive market.

These two versions of the model are, of course, simple abstractions of more complex behavior. The learning and adoption process can be seen as capturing the efforts by traders to learn from and adapt to the evolving market of which they have an imperfect understanding.

Appendix A.

Evans and Honkapohja (2001, pp. 124–125) state assumptions (A.1)–(A.3) and (B.1)–(B.2) under which the stability analysis for the system of the random difference equations can be investigated via the associated system of ODEs. Here we verify that assumptions are satisfied for our setting.

Rewrite Eqs. (28)–(30) more compactly is terms of vector of parameters $\Phi_t = (c_{1,t} \quad Q_t \quad d\bar{\pi}_t)'$, vector of state variables $X_t = (u_t\ p_t)'$, vector of constants $A = (\bar{u} \quad \bar{u}/R)'$, vector of the random determinates of u_t, $W_t = (\theta_t\ \varepsilon_t)'$, functions $H(\Phi, X)$ and $\rho(\Phi, X)$, and a sequence of gains γ_t, describing how the vector of parameters Φ_t is updated, that is,

$$\Phi_t = \Phi_{t-1} + \gamma_t H(\Phi_{t-1}, X_{t-1}) + \gamma_t^2 \rho(\Phi_{t-1}, X_{t-1}),$$

$$X_t = A + B(\Phi_t)W_t, \quad \text{where}$$

$$B(\Phi_t) = \begin{pmatrix} 1 & 1 \\ b_1(\Phi_t) & 0 \end{pmatrix}.$$

Function $H(\Phi, X)$ deduced from the system Eqs. (28)–(30) is

$$H(\Phi, X) = \begin{pmatrix} Q^{-1}\theta b_1(c_1, d\bar{\pi})(u - \bar{u} - c_1 b_1(c_1, d\bar{\pi})\theta) \\ \left(\theta b_1(c_1, d\bar{\pi})\right)^2 - Q \\ \dfrac{1}{1 - f(d\bar{\pi})}(\theta - Rb_1(c_1, d\bar{\pi})\theta)(u - Rb_1(c_1, d\bar{\pi})\theta)\phi^I - \kappa - d\bar{\pi} \end{pmatrix}$$

and $\rho(\Phi, X)$ is given by

$$\rho(\Phi, X) = \begin{pmatrix} \left(\left(\theta b_1(c_1, d\bar{\pi})\right)^2 - Q\right)^{-1} \theta b_1(c_1, d\bar{\pi})(u - \bar{u} - c_1 b_1(c_1, d\bar{\pi})\theta) \\ 0 \\ 0 \end{pmatrix}.$$

(A.1) Positive, nonstochastic, nonincreasing gain sequence γ_t satisfies

$$\sum_{t=1}^{\infty} \gamma_t = \infty \quad \text{and} \quad \sum_{t=1}^{\infty} \gamma_t^2 < \infty$$

For OLS learning used here decreasing sequence of gains is $\gamma_t = 1/t$, which satisfies above conditions.

Define open set $D \subset \mathbb{R}^3$ around equilibrium point Φ^*.

(A.2) For any compact $S \subset D$ there exist C_1, C_2, q_1, q_2 such that for any $\Phi \in S$

$$H(\Phi, X) \le C_1(1 + |X|^{q_1})$$
$$\rho(\Phi, X) \le C_2(1 + |X|^{q_2}).$$

This assumption imposes polynomial bounds on functions $H(\Phi, X)$ and $\rho(\Phi, X)$. Near the equilibrium Φ^*, c_1 is close to R, $b_1(c_1, d\bar{\pi})$ is close to $1/R$ and, hence, $c_1 b_1(c_1, d\bar{\pi})$ and $Rb_1(c_1, d\bar{\pi})$ are close to 1. Finite $b_1(c_1, d\bar{\pi})$ is ensured by the feature of the LPD, where it is assumed that function $f(d\bar{\pi})$ is bounded between (0,1) for finite $d\bar{\pi}$. The required $Q \neq 0$ (nonsingular variance–covariance matrix in terms of OLS) is satisfied given that θ is a random variable (with nonzero variance). Moreover, random variables θ and u are assumed to be Normally distributed and have finite second moments. Hence assumption A.2 is satisfied.

(A.3) For any compact $S \subset D$ and any $X \in S$, function $H(\Phi, X)$ is twice continuously differentiable with bounded partial second derivatives on S.

In the feature of the LPD it was assumed that function $f(d\bar{\pi})$ is of class C^2. Under this assumption partial second derivatives of $H(\Phi, X)$ exist and by similar arguments used in (A.2) they are finite.

(B.1) Vector of the random determinates of u_t, $W_t = (\theta_t\ \varepsilon_t)'$ is iid with finite absolute moments. W_t is assumed to be Normally distributed, hence B.1 is satisfied.

(B.2) For any compact set $S \subset D$:

$\sup_{\Phi \in S} \left|B(\Phi)\right| \le M$ and $B(\Phi)$ satisfy Lipschitz conditions on S.

The condition holds since $b_1(c_1, d\bar{\pi})$ remains finite and has finite partial derivatives in the neighborhood of equilibrium point Φ^* (by the same arguments as in A.2).

Define $h(\Phi) = \lim_{t\to\infty} EH(\Phi, X_t(\Phi))$. The associated ODE is then defined as $d\Phi/d\tau = h(\Phi)$ and under (A.1)–(A.3) and (B.1), (B.2) $h(\Phi)$ is locally Lipschitz.

References

Branch, W., Evans, G.W., 2006. Intrinsic heterogeneity in expectations formation. Journal of Economic Theory 127, 264–295.

Branch, W., McGough, B., 2008. Replicator dynamics in a cobweb model with rationally heterogeneous expectations. Journal of Economic Behavior and Organization 65, 224–244.

Bray, M., 1982. Learning, estimation, and the stability of rational expectations. Journal of Economic Theory 26, 318–339.

Brock, W.A, Hommes, C.H., 1997. A rational route to randomness. Econometrica 65, 1059–1095.

Brock, W.A., Hommes, C.H., 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. Journal of Economic Dynamics & Control 22, 1235–1274.

Chiarella, C., He, X., 2001. Asset price and wealth dynamics under heterogeneous expectations. Quantitative Finance 1, 509–526.

De Fontnouvelle, P., 2000. Information dynamics in financial markets. Macroeconomic Dynamics 4, 139–169.

Diks, C., Dindo, P., 2008. Informational differences and learning in an asset market with boundedly rational agents. Journal of Economic Dynamics and Control 32, 1432–1465.

Evans, G., Honkapohja, S., 2001. Learning and Expectations in Macroeconomics. Princeton University Press, Princeton, NJ.

Föllmer, H., Horst, U., Kirman, A., 2005. Equilibria in financial markets with heterogeneous agents: a probabilistic perspective. Journal of Mathematical Economics 41, 123–155.

Gaunersdorfer, A., Hommes, C.H., Wagener, F.O.O., 2008. Bifurcation routes to volatility clustering under evolutionary learning. Journal of Economic Behavior and Organization 67, 27–47.

Goldbaum, D., 2005. Market efficiency and learning in an endogenously unstable environment. Journal of Economic Dynamics and Control 29, 953–978.

Goldbaum, D., 2006. Self-organization and the persistence of noise in financial markets. Journal of Economic Dynamics and Control 30, 1837–1855.

Grossman, S.J., Stiglitz, J.E., 1980. On the impossibility of informationally efficient markets. The American Economic Review 70, 393–408.

Guse, E., 2010. Heterogeneous expectations, adaptive learning, and evolutionary dynamics. Journal of Economic Behavior and Organization 74, 42–57.

Lordon, F., 1997. Endogenous structural change and crisis in a multiple time-scales growth model. Journal of Evolutionary Economics 7, 1–21.

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, 1306–1322.

McFadden, D., 1981. Econometric Models of Probabilistic Choice in Structural Analysis of Discrete Data with Econometric Applications. MIT Press, Cambridge, MA.

Sethi, R., Franke, R., 1995. Behavioural heterogeneity under evolutionary pressure: macroeconomic implications of costly optimization. The Economic Journal 105, 583–600.