Divergent behavior in markets with idiosyncratic private information*

David Goldbaum†

February 22, 2016

Abstract

Perpetually evolving divergent trading strategies is the natural consequence of a market with idiosyncratic private information. In the face of intrinsic uncertainty about other traders' strategies, participants resort to learning and adaptation to identify and exploit profitable trading opportunities. Model-consistent use of market-based information generally improves price performance but can inadvertently produce episodes of sudden mispricing. The paper examines the impact of trader's use of information and bounded rationality on price efficiency.

Keywords: Heterogeneous Agents, Efficient Markets, Learning, Dynamics, Computational Economics
(JEL Codes: G14, C62, D82)

*I thank Carl Chiarella and members of SydneyAgents for feedback throughout the development of this paper. I gratefully acknowledge the financial support of the Paul Woolley Centre for Capital Market Dysfunctionality.

†Economics Discipline Group, University of Technology Sydney, PO Box 123 Broadway, NSW 2007 Australia, david.goldbaum@uts.edu.au

1 Introduction

Financial markets exhibit extraordinary diversity in investor trading strategies. Widespread among traders are attempts to extract rent through market participation. Vigorous trading and extensive market commentary suggests a lack of uniformity among market participants and possible disagreement as to the true price determination process.

This paper explores a process by which reasonable data-driven adaptation and learning by market participants shape market evolution. The developed model places traders into an imperfect information environment in which the rational expectations equilibrium is analytically inaccessible to the traders for its dependence on a hidden endogenous state variable. An optimizing approach has traders update trading strategies through learning and adaptation. The process can continue without end due to the model's absence of a fixed point. In the developed setting market-based strategies have a role, potentially improving market efficiency, in extracting information from market observables. To the market's potential detriment, while they are able to trade profitably, the traders lack all of the information necessary to employ the market information without error and without potentially distorting the market price.

The financial market setting draws on models of divergent beliefs, learning, and adaptive behavior in financial markets. Foundational investigations such as Hellwig (1980) and Grossman and Stiglitz (1980) considered the role of markets in aggregating and filtering information and the equilibrium implications of the market participants trading on others' private information extracted from the market. Investigations such as Frankel and Froot (1990), De Long et al. (1990a), and De Long et al. (1990b), consider the possible sustainability of multiple beliefs in static settings. Subsequent analysis considered heterogeneous traders in dynamic settings that endogenize current market impact. One approach has traders choose between discrete information options based on past performance. For the models developed in Brock and LeBaron (1996), Brock and Hommes (1998), De Grauwe et al. (1993), and Giardina and Bouchaud (2003), among others, the popularity of a particular information source depend directly on relative performance. Relative performance determines the innovation in popularity in Sethi and Franke (1995), Branch and McGough (2008), and Goldbaum (2005).[1] Another source of

[1] Wealth accumulation to those using the particular information or strategy is another mechanism generating evolution in market impact, as in Chiarella and He (2001), Farmer and Joshi (2002), Chiarella et al. (2006), and Sciubba (2005). Other mechanisms have been considered as well. Lux (1995), for example, relies on investor sentiment, Routledge (1999) incorporates dispersion through random encounters, while dispersion occurs over a

evolution in markets comes from traders updating how they use information in developing a trading strategy. Statistical learning tools, such as Marcet and Sargent (1989a), Marcet and Sargent (1989b), and Evans and Honkapohja (2001). Non-statistical approaches such as the genetic algorithms in LeBaron et al. (1999) and Bullard and Duffy (1999) offer mechanisms by which traders can improve available trading tools, generating evolution in market behavior as strategies improve, following the lessons suggested by recent past events.

Failure by a fundamental trader dominated market to achieve perfect efficiency creates an opportunity for market-based traders to extract information from the price. To make the non-fundamental information viable, models such as developed in Grossman and Stiglitz (1980), Evans and Ramey (1992), Brock and Hommes (1998), Droste et al. (2002), and Chiarella and He (2003) offer market-based trading as a low-cost alternative to acquiring the same information known to an informed group of traders. The alternative approach to information adopted in this paper handicaps fundamental information with private idiosyncratic noise, as in Brock and LeBaron (1996), while making contemporaneous the market extraction of the information, as in Grossman and Stiglitz (1980). In this environment, market-based trading offer the potential to take advantage of the market's filtering properties to gain profitable information not possessed by any individual fundamentally informed trader. The resulting competitive or even superior market-based information, achieved without imposition of cost on the private fundamental information, is, to my knowledge, unique to the developed model.

Market-based trading strategies, particularly low cost trend-following rules, introduce instability in the dynamic financial market system in Brock and LeBaron (1996), Brock and Hommes (1998), De Grauwe et al. (1993), Giardina and Bouchaud (2003), Goldbaum (2003), Lux (1998), and Panchenko et al. (2013). In contrast, for the developed model in which market-based information is capable of generating profits while improving market efficiency, the model captures an environment inherently supportive of traders' use of market-based information. Model mispricing is thus not hardwired into the market-based trading strategy. Mispricing arises only circumstantially when traders, for historical or path dependent reasons, use market-based information inappropriately for the particular realized state.

What the modeled traders do not know for certain is how to interpret the market information. Addressing this uncertainty is confounded by the self-referential aspect of beliefs and by

network in Panchenko et al. (2013).

the model's state-dependent mapping between market observables and investment fundamentals. The former can be overcome with a convergent learning process but only in the absence of continued evolution in the latter.[2] Following Goldbaum and Panchenko (2010), the analysis points to the role of the process governing the population dynamics in shaping market behavior. The presence of the fixed point anchors the asymptotic behavior of the market. The interaction between the self-referential learning and the adaptive behavior can be either a source of asymptotic stability or instability in the absence of a fixed point.

Boundedly rational behavior is a common feature of models exploring trader heterogeneity, divergent beliefs, and learning. Analysis of the model includes an exploration of bounded rationality imposed through memory length. Brock and LeBaron (1996) and LeBaron et al. (1999) highlight the stabilizing influence of long memory on the dynamic system. Long memory, for example, helps to stabilize the inherently destabilizing cobweb model in Branch and McGough (2008) to produce asymptotically similar behavior between the replicator dynamic (RD) and the discrete choice dynamic (DCD) processes.

Incorporating individual rationality of an agent unaware of the correct model, the misspecified equilibrium of Branch and Evans (2006) and the mixed expectations equilibrium of Guse (2010) both describe a fixed point supporting heterogeneous beliefs. The bounded rationality appears in the form of an under-parameterized model that nonetheless appears consistent with the actual law of motion.[3] Goldbaum (2006) imposes constraints consistent with a rational expectations equilibrium solution on the traders' behavior, though the solution itself is hidden. Relaxing these rationality restrictions and incorporating other boundedly rational behavior introduces a variety of mechanisms through which mispricing arises.

The analysis of this paper offers new insight into the market consequences of imperfect idiosyncratic information by analytically extending the model in Goldbaum (2006) and developing a number of new application treatments. The model is attractive for supporting divergent beliefs without arbitrary costs or limitations on choice. Explored are the role of dynamic processes, bounded rationality, and memory in shaping near term evolution and asymptotic behavior.

The remainder of the paper proceeds as follows. Section 2 introduces the structure of the

[2] Bullard (1994), Bullard and Duffy (2001), and Chiarella and He (2003) offer examples with non-convergence in learning.

[3] The mis-specification of the under-parameterized minimum squared variance model is absorbed into the error term. In Guse (2010), the mis-specified model persists supported by a cost advantage.

financial market, trader behavior, and the available information. Also developed in the section is a rational expectations equilibrium as it depends on a hidden endogenous state variable for reference against which to compare market price and beliefs under boundedly rational settings. Both the replicator and discrete choice dynamic models for driving the state variable are evaluated. A model of trader learning is also developed. Different notions of bounded rationality as they may apply to traders in the financial market are discussed in this section as well. Section 3 offers computation analysis of the market, highlighting the interactions between the state variable, learning, and rationality. Conclusions are drawn in Section 4.

2 The Model

Analysis of the model reveals that an equilibrium cannot be achieved with traders employing fundamental information alone. There is a role for market-based information in support of portfolio decisions. With divergent beliefs, the relationship between the price and payoff is found to depend on the unobserved extent to which traders rely on fundamental versus market-based information. A dynamic model is developed based on traders estimating the relationship.

2.1 Information and model development

A large population of N traders trade a risky dividend-paying asset and a risk-free bond paying R. The risky asset can be purchased at price p_t in period t and is subsequently sold at price p_{t+1} after paying the holder dividend d_{t+1} in period $t+1$. The dividend process follows an exogenous AR(1) process

$$d_{t+1} \quad = \quad \phi d_t + \epsilon_{t+1},\ \phi \in (0,1) \tag{1}$$

normalized to mean zero with innovations distributed $\epsilon_t \sim \text{IIDN}(0, \sigma_\epsilon^2)$. Available to the traders for time t trading is a combination of public and private fundamental information as well as market-based information,

$$Z_{it} = \{s_{it}, p_t, d_t, p_{t-1}, d_{t-1}, \dots\}.$$

The dividend d_t is paid at the start of the period and its value is public knowledge at the time of trade. The Walrasian price p_t is not yet realized but can be conditioned on when the trader submits a demand function. Each trader has access to a private idiosyncratically noisy signal, s_{it}, centered on next period's dividend,

$$\begin{aligned} s_{it} &= d_{t+1} + e_{it} \\ e_{it} &\sim \text{IIDN}(0, \sigma_e^2). \end{aligned} \tag{2}$$

In forming demand for the risky asset, traders use available information to forecast the future payoff, $p_{t+1}+d_{t+1}$. The population is heterogeneous in how much weight to place on fundamental versus market-based information. The proper balance turns out to be state dependent and hidden so that the selective use of information will be consistent with the developed model. At one extreme for trader type, the "fundamental" trader completely discounts market variables as a source of useful information, relying entirely on pubic and private fundamental information to form expectations. At the other extreme, the "market-based" trader uses public information to the exclusion of the noisy private signal. In each period, the traders select

$$\begin{aligned} Z_{it} &= Z_{it}^F \cup Z_{it}^M \\ Z_{it}^F &= \{s_{it}, d_t, d_{t-1}, \dots\} \\ Z_{it}^M = Z_t^M &= \{p_t, d_t, p_{t-1} d_{t-1}, \dots\}. \end{aligned}$$

Traders choose between these two extreme positions.[4]

The equity demand component of the agents' optimal control problem collapses to a spot market decision with negative exponential utility. Given prices $\{p_t\}_{t=0}^{\infty}$, optimal equity demand is given by

$$\hat{q}_{it}^k(p_t) = (E(p_{t+1} + d_{t+1} | Z_{it}^k) - Rp_t)/\gamma\sigma_{kt}^2, k = F, M \text{ and for all } t \tag{3}$$

with $\sigma_{kt}^2 = \text{Var}_{it}(p_{t+1} + d_{t+1} | Z_{it}^k)$. The competitive equilibrium consists of a population of $N = N_t^F + N_t^M$ optimizing traders and a price series $\{p_t\}_{t=0}^{\infty}$ that clears the spot market in equities in each period.

[4]Alternatively, allow every trader use of Z_{it}. There is no REE solution as the trader should optimally ignore the private signal and rely on p_t to extract d_{t+1}, a strategy that is inconsistent with an informative price if universally employed. See Goldbaum (2006) for the derivation and for simulations based on this alternate model.

An equilibrium-consistent fundamental trader belief has p_t that is linear in d_t and d_{t+1} of the form

$$p_t = \frac{1}{R-\phi}\left((1-\alpha)\phi d_t + \alpha d_{t+1}\right) \tag{4}$$

with $\alpha \in [0,1]$.[5] Iterated expectations applied to (1) and (4) produce the fundamental information based forecast of the excess payoff to the risky asset,

$$E(p_{t+1} + d_{t+1}|Z^F_{it}) = \left(\frac{R}{R-\phi}\right) E(d_{t+1}|Z^F_{it}). \tag{5}$$

Independence of the expected payoff from α is convenient from a modeling perspective as the market clearing price is robust to the trader's belief, right or wrong, rational or irrational, regarding the value α. A deficiency in the fundamental understanding of the price process captured by ignorance of the particular value of α does not impact the market. A modeling choice to impose or relax rationality through α does not affect demand for the risky asset. Beliefs that deviate from (4) do affect the market.

Fundamental traders project d_{t+1} on their available Z^F_{it} information, obtaining the mean squared error minimizing forecast with

$$E(d_{t+1}|Z^F_{it}) = (1-\beta)\phi d_t + \beta s_{it} \tag{6}$$

and $\beta = \sigma^2_\epsilon/(\sigma^2_\epsilon + \sigma^2_e)$. Thus, fundamental traders' uncertainty, the consequence of awareness of the idiosyncratic component of their signal, leads to fundamental trader reliance on d_t when forecasting d_{t+1}.

The market-based traders employ a forecasting model that is linear in all relevant variables consistent to forecasting the following period's payoff,

$$E(p_{t+1} + d_{t+1}|Z^M_t) = c_{0t} + c_{1t}p_t + c_{2t}d_t. \tag{7}$$

Let q^k_t be the average demand of the population of type k traders, $k = F, M$. Based on

[5] The fundamental trader belief is supported by an initial guess of an unconstrained linear price, $p_t = b_0 + b_1 d_t + b_2 d_{t+1}$, subsequently verified by the market clearing solution expressed in (11) - (13). The equilibrium-consistent market-based trader belief, including (23), produce the zero intercept and constrained coefficients incorporated into (4).

individual expectations (5) and (7),[6]

$$q_t^F = \left(\frac{R}{R-\phi}((1-\beta)\phi d_t + \beta d_{t+1}) - Rp_t\right)/\gamma\sigma_{Ft}^2, \quad (8)$$

$$q_t^M = (c_{0t} + c_{2t}d_t - (R - c_{1t})p_t)/\gamma\sigma_{Mt}^2. \quad (9)$$

Though no individual fundamental trader knows the value of d_{t+1}, it is reflected in q_t^F without noise as aggregation filters the idiosyncratic component of s_{it}. With portion n_t of traders using the fundamental approach and $1 - n_t$ employing the market-based approach, a consistent Walrasian price function is

$$p_t = p_t(n_t, \mathbf{c}_t) = b_{0t} + b_{1t}(n_t, \mathbf{c}_t)d_t + b_{2t}(n_t, \mathbf{c}_t)d_{t+1} \quad (10)$$

in which $\mathbf{c}_t$ represents a vector of the coefficients in (7). The coefficients of (10) solve the market clearing condition, $n_t q_t^F + (1 - n_t)q_t^M = 0$ at

$$b_0(n_t, \mathbf{c}_t) = \frac{c_{0t}(1-n_t)\kappa_t^2}{n_t R + (1-n_t)(R - c_{1t})\kappa_t^2}, \quad (11)$$

$$b_1(n_t, \mathbf{c}_t) = \frac{n_t \frac{R}{R-\phi}(1-\beta)\phi + (1-n_t)c_{2t}\kappa_t^2}{n_t R + (1-n_t)(R - c_{1t})\kappa_t^2}, \quad (12)$$

$$b_2(n_t, \mathbf{c}_t) = \frac{n_t \frac{R}{R-\phi}\beta}{n_t R + (1-n_t)(R - c_{1t})\kappa_t^2} \quad (13)$$

where $\kappa_t = \sigma_{Ft}/\sigma_{Mt}$.

The extent to which the market clearing price reflects the public d_t or the private d_{t+1} depends on the confidence of the fundamental traders in their signal (β), the beliefs of the market-based traders about the relationship between market observables and future payoffs (c_{0t}, c_{1t}, c_{2t}), the traders' relative uncertainties in predicting future payoffs (κ_t^2), and the proportion of the market employing the fundamental strategy (n_t). Naturally, also present in the price coefficients are the opportunity cost of investing in the risky asset (R) and the AR(1) coefficient of the dividend process (ϕ).

Let $\hat{\pi}_t^k$ represent the performance measure associated with type k strategy using information up to the realizations of time t dividends. With each fundamental trader trading based on

[6]Formally, $q_t^F = R(((1-\beta)\phi d_t + \beta(d_{t+1} + \frac{1}{n_t N}\sum e_{it}))/(R-\phi) - p_t)/\gamma\sigma_{Ft}^2$ but with a large $n_t N$ population of fundamental traders, the last term is approximately zero.

idiosyncratic information, $\hat{\pi}_t^F$ is the average of the fundamental population. Two processes for governing how popularity evolves in response to performance differentials are considered for how they alter the long run and evolutionary processes of the system. Let innovation population dynamic (IPD) identify the set of processes in which the performance differential determines the innovation in popularity. The example adopted for analysis is the 2-choice version of the more general K choice replicator dynamic (RD) model found in Branch and McGough (2008). The model generates the transition equation

$$n_{t+1} = g(\hat{\pi}_t^F - \hat{\pi}_t^M, n_t) = \begin{cases} n_t + r(\hat{\pi}_t^F - \hat{\pi}_t^M)(1 - n_t) & \text{for } \hat{\pi}_t^F \geq \hat{\pi}_t^M \\ n_t + r(\hat{\pi}_t^F - \hat{\pi}_t^M)n_t & \text{for } \hat{\pi}_t^F < \hat{\pi}_t^M \end{cases} \tag{14}$$

with

$$r(x) = \tanh(\delta x/2) \tag{15}$$

driving the n_t process. Unlike its biological origins, the RD as employed need not be absorbing at the boundaries. According to (14), estimates of superior performance by the counterfactual strategy move the population away from the boundary.[7]

Let level population dynamic (LPD) identify the set of processes in which the performance differential determines the level of popularity. The example adopted for analysis is the discrete choice dynamics (DCD) process, employed in Brock and Hommes (1998), which identifies popularity as a direct function of the performance differential,

$$n_{t+1} = f(\hat{\pi}_t^F - \hat{\pi}_t^M) = \frac{1}{2}(1 + \tanh(\rho(\hat{\pi}_t^F - \hat{\pi}_t^M)/2)). \tag{16}$$

The parameters δ and ρ play similar roles in setting the sensitivity of the trader population to the magnitude of $\hat{\pi}_t^F - \hat{\pi}_t^M$. Under the RD process, the more successful strategy attracts adherents from the less successful strategy, consistent with the process described in Grossman and Stiglitz (1980). Under the DCD, $\hat{\pi}_t^F - \hat{\pi}_t^M$ maps directly into n_t with the superior strategy always employed by the majority of the population.

Conditional variance, $\sigma_{kt}^2 = \text{Var}_{it}(p_{t+1} + d_{t+1}|Z_{it}^k)$, is derived using (10) and the appropriate

[7] Parke and Waters (2007) and Guse (2010) have to contend with absorbing boundary conditions. The boundaries of the present model are shown to be reflective.

(5) or (7) with,

$$\begin{aligned}
p_{t+1} + d_{t+1} - E(p_{t+1} + d_{t+1}|Z_{it}^F) &= \left(1 + b_{1t+1} + \phi b_{2t+1} - \frac{R}{R-\phi}\right)\phi d_t \\
&+ \left(1 + b_{1t+1} + \phi b_{2t+1} - \beta\frac{R}{R-\phi}\right)\epsilon_{t+1} \\
&+ b_{2t+1}\epsilon_{t+2} - \beta\frac{R}{R-\phi}e_{it} \qquad (17) \\
p_{t+1} + d_{t+1} - E(p_{t+1} + d_{t+1}|Z_t^M) &= (\phi(1 + b_{1t+1} + \phi b_{2t+1}) - c_{1t}(b_{1t+1} + \phi b_{2t+1}) - c_{2t})d_t \\
&+ (1 + b_{1t+1} + \phi b_{2t+1} - c_{1t}b_{2t+1}))\epsilon_{t+1} + b_{2t+1}\epsilon_{t+2}. \qquad (18)
\end{aligned}$$

In the order in which they appear in (17), fundamental trader error arrises (i) when the market misprices dividends, (ii) as a result of down-weighting private information about the future dividend due to the noise in the signal, (iii) the unobservable e_{t+2} component of p_{t+1}, and (iv) the realized idiosyncratic noise. Market based trader error arrises as a consequence of (i) inconsistence between c_{1t} and c_{2t}, (ii) misinterpretation of the market information through error in c_{1t}, and (iii) the unobservable e_{t+2} component of p_{t+1}. The developed analytical solution will identify conditions under which certain sources of error can be eliminated.

2.2 Solution

The market-based traders are capable of possessing correct beliefs consistent with the actual pricing function (10).

Definition 1. An n_t-dependent Rational Expectations Equilibrium describes a market in which the coefficients of the market-based strategy in (7) correctly reflect the projection of $p_{t+1} + d_{t+1}$ on d_t and p_t. Further, the fundamental strategy employs beliefs about the price function consistent with (4) and forecast dividends according to (6).

Recall $\kappa^2 = \sigma_F^2/\sigma_M^2$. The n_t-dependent Rational Expectations Equilibrium (REE(n_t)) solution is the b_2 and κ^2 that solve (21) and (24) of the following so that, for $n_t \in (0,1]$,

$$
\begin{aligned}
p_t^* = p_t^*(n_t) &= b_1^*(n_t)d_t + b_2^*(n_t)d_{t+1} && (19)\\
b_1^*(n_t) &= \frac{n_t(1-\beta)\phi}{(R-\phi)\left(n_t+(1-n_t)\kappa^*(n_t)^2\right)} && (20)\\
b_2^*(n_t) &= \frac{n_t\beta+(1-n_t)\kappa^*(n_t)^2}{(R-\phi)\left(n_t+(1-n_t)\kappa^*(n_t)^2\right)} && (21)\\
c_1^*(n_t) &= \frac{R}{(R-\phi)b_2^*(n_t)} = \frac{R\left(n_t+(1-n_t)\kappa^*(n_t)^2\right)}{n_t\beta+(1-n_t)\kappa^*(n_t)^2} && (22)\\
c_2^*(n_t) &= \frac{\phi}{R-\phi}(R-c_1^*(n_t)) = -\frac{n_tR(1-\beta)\phi}{(R-\phi)\left(n_t\beta+(1-n_t)\kappa^*(n_t)^2\right)} && (23)\\
\kappa^*(n_t)^2 &= \frac{\sigma_F^*(n_t)^2}{\sigma_M^*(n_t)^2} = 1+\frac{(1-\beta)R^2}{(R-\phi)^2b_2^*(n_t)^2} && (24)
\end{aligned}
$$

where k^{*2} follows from

$$
\begin{aligned}
\sigma_F^*(n_t)^2 &= \left((1-\beta)\left(\frac{R}{R-\phi}\right)^2 + b_2^*(n_t)^2\right)\sigma_\epsilon^2 && (25)\\
\sigma_M^*(n_t)^2 &= b_2^*(n_t)^2\sigma_\epsilon^2. && (26)
\end{aligned}
$$

As $n_t \to 0$, then $b_1^*(n_t) \to 0$, $b_2^*(n_t) \to 1/(R+\phi)$, $c_1^*(n_t) \to R$, and $c_2^*(n_t) \to 0$. For $n_t = 0$, then $b_1^*(0) = \phi/(R-\phi)$, $b_2^*(0) = 0$ as derived from the consistent solution $c_1^*(0) = 0$ and $c_2^*(0) = R$.

Accepting the partition of information into fundamental and market-based, the REE(n_t) deviates from a true rational expectations equilibrium in that the derived market-based traders' beliefs at the REE(n_t) relies on treating expectations of all future n_τ, $\tau \geq t+1$, as a point estimate unchanging from to the current n_t. Deviations would have consequence on the price stream. A constant n_t may or may not be consistent with the process driving n_t at the REE(n_t) solution.

Let p_t^0 and p_t^1 represent the price at the two information extremes based on the accuracy of the private signal. With zero content in the signal, with $\sigma_\epsilon \to \infty$, then $\beta = 0$. Zero error, with $\sigma_\epsilon = 0$, results in $\beta = 1$. For $n_t \neq 0$,[8]

$$
\begin{aligned}
p_t^0 &\equiv p_t^*(1)|_{\beta=0} = \frac{\phi}{R-\phi}d_t\\
p_t^1 &\equiv p_t^*(n_t)|_{\beta=1} = \frac{1}{R-\phi}d_{t+1}.
\end{aligned}
$$

[8]Prices p_t^0 and p_t^1 also correspond to the Fama (1970) semi-strong and strong form efficient prices.

Let p_t^F represent the price at the extreme of a market populated by only fundamental traders. With $n_t = 1$,

$$p_t^F \equiv p_t^*(1) = \frac{(1-\beta)\phi}{R-\phi} d_t + \frac{\beta}{R-\phi} d_{t+1}. \tag{27}$$

The opening for profitable employment of the market-based information follows from $p_t^F \in [p_t^0, p_t^1]$, introducing predictability in the price as a consequence of d_{t+1} contributing to the value of both p_t and p_{t+1}. The presence of the market-based traders moves the market towards the efficient market price, as reflected in $p_t^*(n_t) \in [p_t^F, p_t^1]$ with $\lim_{n_t \to 0} p_t^*(n_t) = p_t^1$. Since $p_t^*(0) = p_t^0$ there is a Grossman and Stiglitz (1980) type discontinuity at $n_t = 0$.

Observe that $b_1^*(n_t) + \phi b_2^*(n_t) = \phi/(R-\phi) \forall n_t$. Let

$$\alpha_t^* = \alpha^*(n_t) = \frac{n_t \beta + (1-n_t)\kappa^*(n_t)^2}{n_t + (1-n_t)\kappa^*(n_t)^2}, \tag{28}$$

allowing the REE(n_t) price to be expressed as $p_t^* = \frac{1}{R-\phi}\left((1-\alpha_t^*)\phi d_t + \alpha_t^* d_{t+1}\right)$. With $\alpha_t^* \in [\beta, 1]$, the REE($n_t$) price can be interpreted as the present discounted value reflecting the aggregation of the market's forecast of future dividends. The extent to which the REE(n_t) price reflects the public d_t or the private d_{t+1} depends on the quality of the signal, as reflected in β, and the traders' choice of information, as reflected in n_t.

Contributing to the robustness of the model, fundamental trader ignorance of n_t does not undermine the existence of the REE(n_t) solution. Equivalent to previously identified freedom from the fundamental traders' beliefs regarding, the REE(n_t) solution requires only that each fundamental trader holds beliefs regarding price formation consistent with $b_{1t} + \phi b_{2t} = \phi/(R-\phi)$, free from particular knowledge of b_{1t} and b_{2t} individually. Without consequence on the REE(n_t) solution or the market clearing price, each fundamental trader can independently employ the correct $p_t^*(n_t)$, mistakenly employ $p_t^*(m_t)$ for $m_t \neq n_t$, or naively employ the p_t^F, p_t^0, p_t^1, or any other price structure consistent with (4). Given a price realization consistent with (4), accuracy of the fundamental trader forecast is a reflection of the accuracy of their forecast of d_{t+1} and is independent of the believed α.

Useful for consistency of trader behavior when encountering non-equilibrium pricing, the condition imposed on market-based trader beliefs to ensure $b_{1t} + \phi b_{2t} = \phi/(R-\phi)$ is weaker than the conditions necessary to generate the REE(n_t). It only requires $c_{2t} = (R - c_{1t})\phi/(R-\phi)$

as evidenced by substituting for c_{2t} in (12) to produce $b_1(n_t, c_{1t}) + \phi b_2(n_t, c_{1t}) = \phi/(R - \phi)$ regardless of the value of c_{1t}. Thus, the condition $c_{2t} = c_2^*(c_{1t})$ implied by (23) is a sufficient condition to support the price structure underpinning the fundamental strategy. That is, in order for the fundamental traders' forecast to conform to the requirements of a REE(n_t), the market-based traders need only employ a c_{2t} value that is REE(n_t) consistent with c_{1t} without necessarily employing the correct REE(n_t) implied $c_{1t} = c_1^*(n_t)$.

For $c_{2t} = c_2^*(c_{1t})$, all trader beliefs are consistent with the price determination process. For the market-based traders, error arises from ignorance of n_t that leads to errors in estimating d_{t+1}, not erroneous beliefs about how prices are formed according to n_t. In the error terms expressed in (17) and (18), the market-based traders' employment of $c_{2t} = c_2^*(c_{1t})$ eliminates the first term of (18) and the resulting $b_{1t} + \phi b_{2t} = \phi/(R - \phi)$ eliminates the first term of (17) as well.

The payoff to the risky asset is never without some level of uncertainty. The closer p_t is to p_t^1, the more p_{t+1} reflects the yet unobserved d_{t+2}, increasing the market's uncertainty about investing in the risky asset. The same level of uncertainty also arises for $n_t = 1$ and $\beta = 0$ or for $n_t = 0$. In both of these latter cases, p_{t+1} depends on d_{t+1} only but no time t trader has information concerning the value of d_{t+1}.

Realizing the REE(n_t) requires that the market-based traders employ $\mathbf{c}_t = \mathbf{c}^*(n_t)$. It is appropriate to ascertain whether the traders can deduce $\mathbf{c}^*$ analytically from their knowledge of the market. From (23), c_2^* can be expressed in terms of c_1^*. For a known zero net supply of the risky asset, the traders can derive analytically that $c_0^* = 0$. For market-based traders incorporating these two conditions into their understanding of the market, only c_1^* remains to be derived. From (22), solving for c_1^* requires knowledge of n_t. Reasonably, n_t is not directly observable. For an n_t that is the endogenous product of a dynamic system, the question becomes whether some fixed point n^{fp} value can be identified that is consistent with REE(n_t).

2.3 Performance

Define performance in terms of individual profit,

$$\pi_{it}^k = q_{it}^k (p_{t+1} + d_{t+1} - Rp_t). \tag{29}$$

Using the REE(n_t) consistent $c_0 = 0$ and $c_{2t} = c_2^*(c_{1t})$, and the market clearing condition for b_{2t} from (13), (29) generates, for $n_t \in (0, 1]$,

$$E(\pi_t^F) = (1 - n_t)\Delta_t \tag{30}$$

$$E(\pi_t^M) = -n_t\Delta_t \tag{31}$$

so that $E(\pi_t^F - \pi_t^M) = \Delta_t$. Here,

$$\Delta_t = \Delta(c_{1t}, n_t) = \left(\frac{n_t(1-\beta)R + (1-n_t)(R-c_{1t})\kappa^2}{(n_tR + (1-n_t)(R-c_{1t})\kappa^2)^2}\right)\left(\frac{R}{R-\phi}\right)^2 \frac{(R-c_{1t})\beta\sigma_\epsilon^2}{\sigma_M^2}. \tag{32}$$

The REE(n_t) expected profit differential, $E(\pi_F^* - \pi_M^*)$, based on $c_{1t} = c_1^*(n_t)$, reduces to

$$\Delta^*(n_t) = -\left(\frac{1-\beta}{n_t + (1-n_t)\kappa^*(n_t)^2}\right)^2 \left(\frac{R}{R-\phi}\right)^2 \frac{n_t\sigma_\epsilon^2}{\sigma_M^{*2}}. \tag{33}$$

That $\Delta^*(n_t) < 0$ for all $n_t \neq 0$ reveals the benefit to extracting filtered information from the REE market over direct access to noisy information. The fundamental traders only profit in the presence of error in the market-based traders' model, as c_{1t} deviates sufficiently from $c_1^*(n_t)$, allowing Δ_t to be positive.

A fixed point to the entire dynamic system requires the REE(n_t) solution combined with a fixed point to the population process. The fixed point condition depends on the population regime.

Proposition 1. *Given a level population dynamic (LPD) for n_t, the REE(n_t) competitive equilibrium has a unique fixed point n^{fp} at which $n^{fp} = f(\Delta^*(n^{fp}))$.*

Proof. Under the LPD population process, $n_{t+1} = f(\hat{\pi}_t^F - \hat{\pi}_t^M)$ according to (16) and at the REE(n_t), $\hat{\pi}_t^F - \hat{\pi}_t^M = \Delta^*(n_t)$. For $\rho < \infty$, $f(x)$ is continuous and monotonically increasing in x. A fixed point solution is n^{fp} such that $n^{fp} = f(\Delta^*(n^{fp}))$. Since $\lim_{n_t \to 0} \Delta^*(n_t) = 0$ and $\Delta^*(n_t)$ is monotonically decreasing as n_t increases to one, a unique n^{fp}, $0 < n^{fp} \leq 1/2$, such that $n^{fp} = f(\Delta^*(n^{fp}))$ exists. $\square$

Figure 1 captures the existence of the fixed point under the DCD population process. Since the slope of $f(x)\mid_{x=0}$ increases with ρ in Figure 1, the value of $n^{fp} \in (0, 1/2]$ decreases with increasing ρ. At the extremes, $\rho = 0$ results in a horizontal $f(\pi^F - \pi^M)$ and $n^{fp} = 1/2$ while

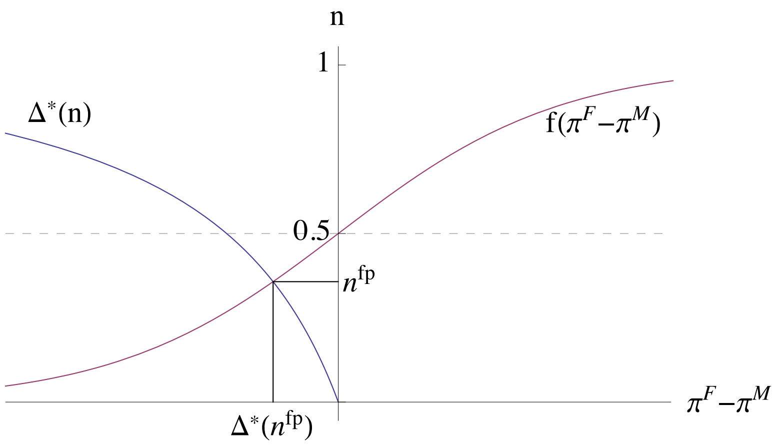

Figure 1: n^{fp} for DCD population dynamics.

$\rho \to \infty$ approaches a step function in $f(\pi^F - \pi^M)$ so that $n^{fp} \to 0$. With $E(\pi^F_{fp} - \pi^M_{fp}) < 0$, the inferior profits of the fundamental strategy at n^{fp} support the realization of $n^{fp} < 1/2$. The DCD fixed point is inconsistent with the Grossman and Stiglitz (1980) notion of an equilibrium in which the expected performance differential is zero.

Proposition 2. *Given an innovation population dynamic (IPD) for n_t, the REE(n_t) competitive equilibrium excludes a fixed point in n_t.*

Proof. Under the IPD population process, the fixed point condition requires the existence of an n^{fp} such that $n^{fp} = g(\Delta^*(n^{fp}), n^{fp})$. With $n = g(\Delta, n)$ if and only if $\Delta = 0$, the fixed point requires $\Delta^*(n^{fp}) = 0$. Since no such n^{fp} exists, there can be no fixed point to the RD population process. □

The existence of an REE depends on the existence of a (n^{fp}, c_1) combination for which $n^{fp} = n^*(c_1^*(n^{fp}))$. Such a point does not exist since for $c_{1t} = c_1^*(n_t)$, $\Delta^*(n_t) < 0$ for all $n_t \in (0, 1]$ and $E(\pi_t^F - \pi_t^M) > 0$ for $n_t = 0$. The Grossman and Stiglitz (1980) discontinuity means that the $n_t = 0$ boundary is reflecting rather than absorbing.

2.4 Learning

Consider a fixed $n_t = n \in (0, 1]$ for all t. Allow traders to update the market-based model based on empirical observations. Least-squares learning involves a process of updating the coefficients of (7) according to

$$\begin{aligned}\hat{\mathbf{c}}_t &= \hat{\mathbf{c}}_{t-1} + \lambda_t (Q_t^{-1} x_{t-2}(p_{t-1} + d_{t-1} - \hat{\mathbf{c}}_{t-1} x_{t-2}))' && (34)\\ \hat{Q}_t &= \hat{Q}_{t-1} + \lambda_t (x_{t-1} x_{t-1}^{'} - \hat{Q}_{t-1}) && (35)\end{aligned}$$

with $\lambda_t = 1/t$ and $x_t = \{1, p_t, d_t\}$. Here, $\hat{c}_t$ reflects the time t estimate of the corresponding coefficient of (7) based on the learning algorithm while $\hat{Q}_t$ is the estimate of the variance-covariance matrix for x_t used in the estimation of $\hat{c}_t$.

Proposition 3. *Given a fixed n, $\sigma_{kt}^2 = \sigma_k^2(n, \mathbf{c}_t)$, a sequence of market clearing prices $\{p_t\}_{t=0}^{\infty}$, and least squares updating of beliefs and performance, the REE(n) competitive equilibrium is locally stable.*

Proof. See Appendix □

By Proposition 3, the self-referencing system of prices and beliefs is locally stable at the REE(n) under least-squares learning. As the fixed point to the learning process, $\mathbf{c}^*(n)$ are the rational expectations coefficients for the market-based traders and REE(n) is a rational expectations equilibrium for market-based traders.

The parameters λ_t regulates the learning process. With $\lambda_t = 1/t$, the traders update the market-based model consistent with the standard least-squares learning algorithm of Marcet and Sargent (1989b), giving equal weight to each observation. Least-squares learning is a natural choice for the traders given a fixed hidden state variable. The perpetually evolving state of the RD processes can make other parameter updating processes seem reasonable. For $\lambda_t = \lambda$, $0 < \lambda < 1$, the traders update with a constant gain by which the contribution of past observations to the current parameter estimate decays exponentially.

The variables included in the vector x_t can be adapted to the presumed rationality of the market-based traders. Imposing $c_{0t} = 0$ and $c_{2t} = c_2^*(c_{1t})$, x_t becomes a scalar, $x_t = p_t - \phi d_t/(R - \phi)$.

Allowing n_t to evolve over time imposes on the traders a need to estimate other n_t-dependent variables as accommodation to the inability of the traders to derive the values analytically. Traders need to estimate both π_t^F and π_t^M as inputs to the model adoption decision. Each trader must also estimate the σ_{kt}^2, $k \in \{F, M\}$, appropriate for the model adopted as an input to demand. Consider the updating algorithms

$$\hat{\sigma}^2_{kt} = \hat{\sigma}^2_{kt-1} + \theta_t((p_t + d_t - E(p_t + d_t|Z^k_{t-1}))^2 - \hat{\sigma}^2_{kt-1}) \tag{36}$$

$$\hat{\pi}^k_t = \hat{\pi}^k_{t-1} + \mu_t(\pi^k_{t-1} - \hat{\pi}^k_{t-1}),\ k = F, M \tag{37}$$

Like λ_t, the parameters θ_t and μ_t regulate the learning process. Under least-squares learning, with $\theta_t = \mu_t = 1/t$, $\hat{\sigma}^2_{kt}$ and $\hat{\pi}^k_t$ are simple sample averages of all past observations. A constant gain biases weight towards the more recent observations.

2.5 Evolution without a Fixed Point

Imposing the constraints $c_{0t} = c^*_0 = 0$ and $\hat{c}_{2t} = c^*_2(\hat{c}_{1t})$ of the REE(n_t) solution, the system, represented by equations (1), (10), (14), (34), (35) and (37), can be evaluated using phase space dynamics in the (c_1, n) plane. The phase space analysis is facilitated by imposing $\mu_t = \mu = 1$ in (37) so that the time t realization of $\pi^F(c_{1t}, n_t) - \pi^M(c_{1t}, n_t)$ alone identifies n_{t+1}. Though not accessible to the traders, let $\hat{\sigma}^2_{tk} = \sigma^*_k(n_t)^2$ of (25) and (26) for $k \in \{F, M\}$ so that the employed conditional variances are correct expressions reflecting the current n_t rather than the history-dependent estimate (36).

Market-based trader beliefs are unchanging if $\hat{c}_{1t} = c^*_1(n_t)$. At this REE$(n_t)$, the market-based model correctly reflects the relationship between the observables p_t and d_t and the expected payoff of the following period, $E(p_{t+1} + d_{t+1})$. The function $c^*_1(n_t)$ is monotonically increasing for $0 < n_t \leq 1$ with $c^*_1(n) \to R$ for $n \to 0$ and $c^*_1(1) = R/\beta$.

The population process is at a steady state if $\Delta(\hat{c}_{1t}, n_t) = 0$. The coefficient c_{1t} appears twice in the numerator of $\Delta(\hat{c}_{1t}, n_t)$. Let $c^+_1(n_t)$ and c^-_1 represent the two functions capturing combinations of $\hat{c}_{1t}$ and n_t consistent with $\Delta(\hat{c}_{1t}, n_t) = 0$ in (32). For $0 < n_t \leq 1$, the former is monotonically increasing and everywhere above $c^*_1(n_t)$,

$$c^+_1(n_t) = R\left(1 + (1-\beta)\frac{n_t}{(1-n_t)\kappa^2}\right). \tag{38}$$

The latter is a constant, $c^-_1 = R$, located below $c^*_1(n_t)$. Expected profits are zero at $\hat{c}_{1t} = c^+_1(n_t)$ because the resulting market clearing price is the efficient market price, p^1_t, at which expected profits are zero regardless of the individual trader's position taken in the market. Expected

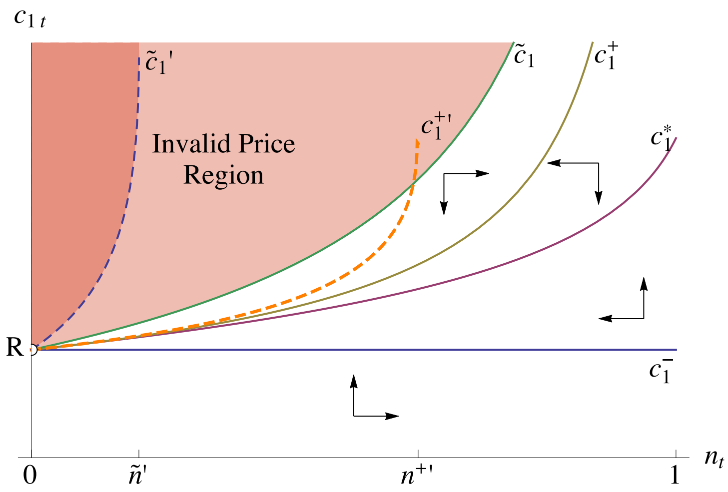

Figure 2: Phase space in n_t and $\hat{c}_{1t}$ for the RD population process. $c_1^*(n_t)$ is the REE(n_t) value of $\hat{c}_{1t}$ and the attractor to the learning process for a given n_t. For $c_1^- < \hat{c}_{1t} < c_1^+(n_t)$ the market-based model is sufficiently accurate to earn profits at the expense of the fundamental strategy, leading to a decline in n_t. For $\hat{c}_{1t} < c_1^-$ and for $c_1^+(n_t) < \hat{c}_{1t} < \tilde{c}_1(n_t)$ the fundamental strategy dominates the market-based strategy so that from these regions n_t is increasing. Above $\tilde{c}_1(n_t)$, the aggregate demand curve for the risky security is upward sloping and no positive price exists to clear the market. The dashed lines reflect an alternate specification for which the current-period implication of $\hat{c}_{1t} \neq c_1^*(n_t)$ is recognized when calculating the market-based model error. The market is more tolerant of error, as reflected in $c_1^{+'}$ and $\tilde{c}_1'$, when market-based trader are increasingly uncertain in the face of large price deviations (developed in Section 2.6.2).

profits are zero at $\hat{c}_{1t} = c_1^-(n_t)$ because the market traders expect the risky asset to offer the same return as the risk-free bond and thus there is no trading at the market clearing price.

A third relevant function included in the phase space is $\tilde{c}_1(n_t)$. The expression $n_t R + (1 - n_t)(R - \hat{c}_{1t})\kappa_t^2$ appears in the denominator of the two pricing coefficients, $b_1(\hat{c}_{1t}, n_t)$ and $b_2(\hat{c}_{1t}, n_t)$ as well as the denominator of $\Delta(\hat{c}_{1t}, n_t)$. The negative of the expression is the slope of the risky asset's aggregate demand function so that when it is zero the market demand function is horizontal and different from zero, producing an infinite market clearing price (based on a zero net supply). Let $\tilde{c}_1(n_t)$ be the function

$$\tilde{c}_1(n_t) = R\left(1 + \frac{n_t}{(1 - n_t)\kappa^2}\right), \tag{39}$$

producing $n_t R + (1 - n_t)(R - \tilde{c}_1(n_t))\kappa^2 = 0$. For $0 < n_t \leq 1$, $\tilde{c}_1(n_t)$ is monotonically increasing and everywhere above $c_1^+(n_t)$. Combinations of c_{1t} and n_t approaching the function from below or from the right generate $p_t(c_{1t}, n_t) \to \infty$ and $\Delta(\hat{c}_{1t}, n_t) \to \infty$.

Above $\tilde{c}_1(n_t)$, the combination of n_t and $\hat{c}_{1t}$ do not allow for a reasonable market clearing price. The precarious nature of the market in the vicinity of $\tilde{c}_1(n_t)$ is the consequence of the excessive influence of the market-based traders. As a group, they have an upward sloping demand function in price. From the perspective of the market-based traders, an increase in the price is interpreted as an indication of good news about the underlying d_{t+1}, increasing demand. At $\hat{c}_{1t} = c_1^*(n_t)$, the market-based model correctly accounts for the influence of the market-based trader population on the price. As a consequence, the aggregate demand for the risky asset remains downward sloping in p_t. For $\hat{c}_{1t} > c_1^*(n_t)$, the market-based model projects too large a deviation in d_{t+1} based on the observed p_t. The market-based traders thus take too large a position relative to the underlying reality. For $\hat{c}_{1t} > \tilde{c}_1(n_t)$, the position produces an upward-sloping aggregate market demand function.[9]

The traders themselves cannot be relied upon to recognize dangerous market conditions introduced by their own belief. Implicit in the trader's use of $\hat{c}_{1t}$ is that it is a reasonable approximation of $c_1(n_t)$ for the current unobserved n_t. For any $\hat{c}_{1t} \in (R, R/\beta]$ there exists n_1 and n_2, $0 < n_1 < n_2 \leq 1$ for which $\hat{c}_{1t} = \tilde{c}_1(n_1)$ and $\hat{c}_{1t} = c_1^*(n_2)$. The market-based traders' belief that $c_{1t} = \hat{c}_{1t}$ is reasonable if the unobserved n_t is near n_2 but disastrously wrong, generating substantial mispricing, if n_t is near n_1. The greater distances between $c_1^*(n_t)$ and $c_1^+(n_t)$ and between $c_1^+(n_t)$ and $\tilde{c}_1(n_t)$ as n_t increases reflect a market more tolerant of trader error.

Given n_t, $c_1^*(n_t)$ is an attractor for $\hat{c}_{1t}$. For $\hat{c}_{1t}$ between $c_1^-(n_t)$ and $c_1^+(n_t)$, $E(\Delta(\hat{c}_{1t}, n_t)) < 0$ so that n_t tends to decline. In this range, the market-based model, while not necessarily perfectly correct for extracting information from the price, is more accurate than the average fundamental trader relying on a noisy signal. Outside this range, with $\hat{c}_{1t} < c_1^-$ or $c_1^+(n_t) < \hat{c}_{1t} < \tilde{c}_1(n_t)$, the inaccuracy in the market-based model is large enough that the user of the fundamental information expects to earn profits at the expense of the market-based traders and therefore n_t tends to increase in this region.

All four functions of the phase space radiate out from the point $n_t = 0$ and $\hat{c}_{1t} = R$ but because of the discontinuity at $n_t = 0$, none take a value of R at $n_t = 0$. Therefore, though the four functions come arbitrarily close, they never intersect. The failure of $c_1^*(n_t)$ to intersect with either $c_1^-(n_t)$ or $c_1^+(n_t)$ graphically captures the absence of a fixed point to the RD dynamic

[9] As an alternate interpretation, for $\hat{c}_{1t} > c_1^*(n_t)$, the market-based model can be seen as underestimating the influence of the market-based traders on the price since, for $\hat{c}_{1t} < R/\beta$, there exists $n > n_t$ such that $\hat{c}_{1t} = c_1^*(n)$.

system.

2.6 Bounded Rationality

How closely the traders adhere to rationality strongly influences the near-term evolution and asymptotic characteristics of the market. Constraints on the parameters reflect traders making full use of the information and knowledge of price determination to impose rationality on beliefs and behavior. Relax these conditions and the market is no longer constrained to exhibit features of the rational expectations equilibrium. Features of the boundedly rational market are thus self-fulfilling.

Considered analytically in this section and computationally in the following section are model treatment variations based on (i) the consequence of allowing $\hat{c}_{2t} \neq c_2^*(\hat{c}_{1t})$, both with and without appropriate accommodation by the fundamental traders, (ii) the presence or absence of a fixed point as determined by the population process, and (iii) memory length as captured by λ_t, and μ_t.

2.6.1 Knowing fundamental traders

The fundamental traders hold beliefs consistent with the market so long as the market produces $b_{1t} + \phi b_{2t} = \phi/(R - \phi)$. The condition is violated when $\hat{c}_{2t} \neq c_2^*(\hat{c}_{1t})$. The disruption to pricing is profound when the fundamental traders try to incorporate the deviation from proper pricing into their beliefs.

As previously asserted, traders aware of the market and its structure can deduce that $c_0 = 0$ regardless of the unobservable n_t. They may also incorporate a feature of the REE(n_t) by imposing $\hat{c}_{2t} = c_2^*(\hat{c}_{1t})$ according to (23). Let $b_2(n_t, \hat{c}_{1t})$ represent the market clearing value of b_2 with $\hat{c}_{1t}$ set freely but $\hat{c}_{2t} = c_2^*(\hat{c}_{1t})$. This REE($n_t$) constrained $b_2(n_t, \hat{c}_{1t})$ has up to three roots of which only one is real at any particular value of $n_t \in (0, 1]$. Naturally, $b_{1t} = \phi/(R - \phi) - \phi b_{2t}$ completes the pricing equation.

For $R < \hat{c}_{1t} < R/\beta$, Figure 3 includes an example of the $b_2(n_t, \hat{c}_{1t})$ solution. Let $\tilde{n}$ represent the value of n_t such that $\hat{c}_{1t} = \tilde{c}_1(n_t)$ so that $\tilde{n}(\hat{c}_{1t}) = \tilde{c}_1^{-1}(\hat{c}_{1t})$. For $n_t \in (\tilde{n}, 1]$, $b_2(n_t, \hat{c}_{1t})$ is increasing as n_t declines with $b_2(n_t, \hat{c}_{1t}) \to +\infty$ as $n_t \to \tilde{n}(c_1)$. For the invalid price region $n \in (0, \tilde{n})$, $b_2(n_t, \hat{c}_{1t})$ is negative.

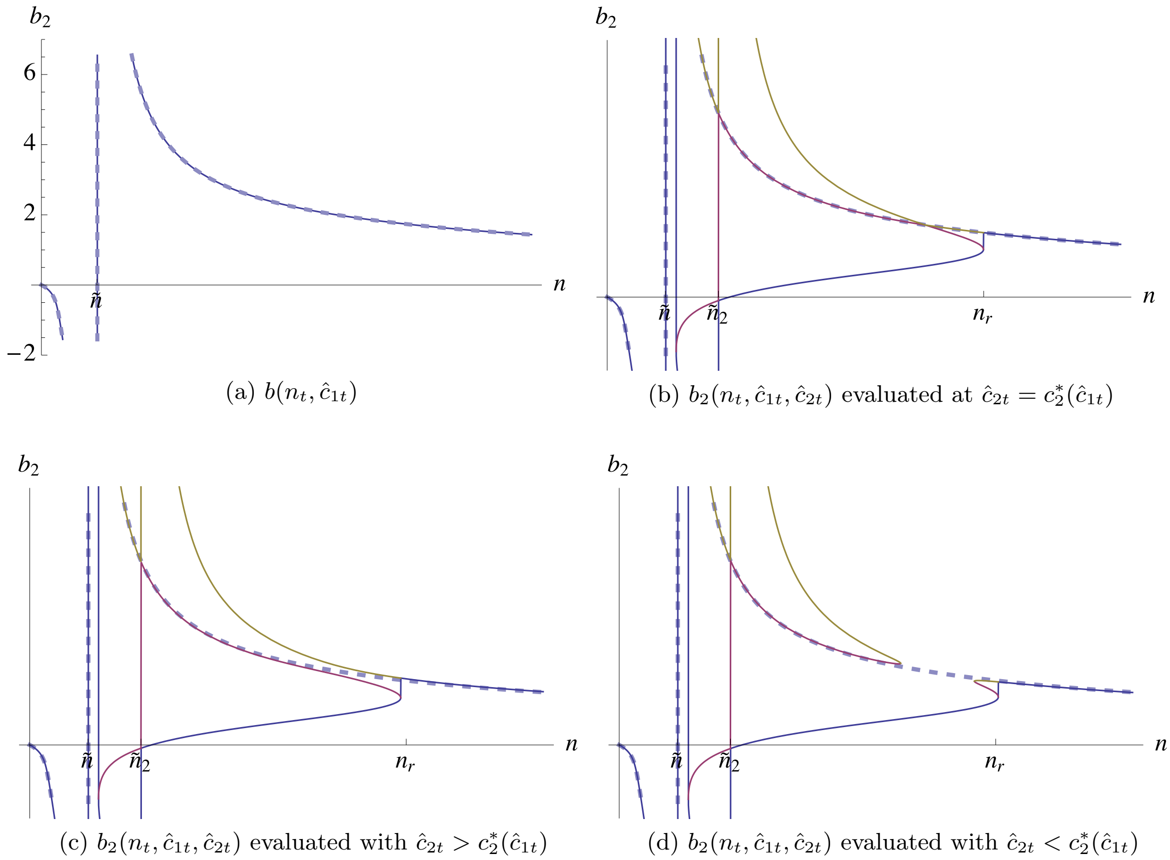

Figure 3: Real roots of $b_2(n_t, \hat{c}_{1t})$ and $b_2(n_t, \hat{c}_{1t}, \hat{c}_{2t})$ in the presence of fundamental traders able to account for the deviation from $b_{1t} + \phi b_{2t} = \phi/(R - \phi)$. $R = 1.02$, $\phi = 0.5$, and $\sigma_\epsilon = \sigma_e = 1$. With $c_1 = 1.1$ then $\tilde{n}(1.1) = 0.034$. The $b_2(n_t, \hat{c}_{1t})$ solution is included in each frame (dashed line). For $\hat{c}_{2t} = c_2^*(\hat{c}_{1t})$, for all values of n_t, one of the roots of $b_2(n_t, \hat{c}_{1t}, \hat{c}_{2t})$ is equal to the root of $b_2(n_t, \hat{c}_{1t})$. For $\hat{c}_{2t} > c_2^*(\hat{c}_{1t})$, the point of tangency that ensures continuity in the roots of $b_2(n_t, \hat{c}_{1t}, \hat{c}_{2t})$ that match the root of $b_2(n_t, \hat{c}_{1t})$ is lost. For $\hat{c}_{2t} < c_2^*(\hat{c}_{1t})$, a gap opens between the two points of intersection in two of the roots tracing $b_2(n_t, \hat{c}_{1t})$.

Market-based traders unaware or uninterested in imposing $\hat{c}_{2t} = c_2^*(\hat{c}_{1t})$ introduce deviations from $b_{1t} + \phi b_{2t} = \phi/(R - \phi)$. Consider a sophisticated fundamental trader able to account for these deviations.[10] Let $b_2(n_t, \hat{c}_{1t}, \hat{c}_{2t})$ represent the market clearing value of b_2 with both $\hat{c}_{1t}$ and $\hat{c}_{2t}$ set freely. The unconstrained $b_2(n_t, \hat{c}_{1t}, \hat{c}_{2t})$ has five roots. For $n_t > n_r$, where $n_r \in (\tilde{n}, 1)$, only one of the roots is real. For $n_t \leq n_r$, up to three of the roots are real, producing three possible market clearing prices.

Figure 3 highlights the challenges in identifying market clearing price produced by this scenario. Rather than imposing market discipline, the sophisticated fundamental traders exploiting market-based trader error increase intractability with possible multiple market clearing prices depending on the state.

As seen in frame 3b, evaluated at $\hat{c}_{2t} = c_2^*(\hat{c}_{1t})$, one of the three real roots of $b_2(n_t, \hat{c}_{1t}, \hat{c}_{2t})$ coincides with the real root of $b_2(n_t, \hat{c}_{1t})$. The other two other roots are real for values of $n_t \in (0.034, 0.22)$ starting at a single point at $n_t = 0.22$ and bifurcating with one branch increasing as n_t declines and the other decreasing. The denominator of $b_2(n_t, \hat{c}_{1t})$ is linear in n_t, crossing zero once at $\tilde{n}_t$. The denominator $b_2(n_t, \hat{c}_{1t}, \hat{c}_{2t})$ is a binomial in n_t and crosses zero twice. This explains the additional root converging to $+\infty$ as $n_t \to 0.065$ from above.

For $\hat{c}_{2t} \neq c_2^*(\hat{c}_{1t})$, the three real roots do not always coincide with the constrained solution for a range of n_t. As depicted in frame 3c, for $\hat{c}_{2t} > c_2^*(\hat{c}_{1t})$, since the roots do not intersect, to maintain $b_2(n_t, \hat{c}_{1t})$ over the range of n_t requires jumping from one root to another. For $\hat{c}_{2t} < c_2^*(\hat{c}_{1t})$ seen in frame 3d, there is a range over which the two roots most closely approximating $b_2(n_t, \hat{c}_{1t})$ become imaginary. The only remaining real root produces a b_2 that declines as n_t declines.

2.6.2 Concurrent conditional variance

Another challenge for the traders is how to evaluate the error associated with forecasting the payoff, a component of submitted demand. The error variance can be derived analytically for each forecast strategy only with knowledge of the true n_t-dependent pricing relationship. Equations (25) and (26) capture the REE(n_t) values while (36) is driven by the data. The ratio $\sigma_F^*(n)^2/\sigma_M^*(n)^2$ is monotonically increasing in n. At the default simulation parameter

[10] Here, the fundamental traders exploit deviations from $b_{1t} + \phi b_{2t} = \phi/(R - \phi)$. To do so, they need to know the correct $\alpha^*(n_t)$, a level of knowledge not granted to them outside of this subsection.

values considered in Section 3, the ratio ranges from 1.52 to 3.08 for increasing $n \in (0,1]$. In simulation, whether the traders employ (25) and (26) or (36) to compute conditional variance has negligible effect since both estimates track closely to the REE(n_t) values.

The inconsequence follows from the fact that in both computations, the conditional variance is determined independent of the current market realization. By (36), $\hat{\kappa}_t$ is determined by events up to $t-1$ while (25) and (26) presume a price consistent with REE(n_t). Both formulas prevent the traders from using the current market realization as an input into the uncertainty. With κ_t in the denominator of both c_1^+ and $\tilde{c}_1$, an increase in relative uncertainty among the employers of the market-based strategy decreases the market-based traders' price impact by decreasing the size of their position.

To incorporate this uncertainty, allow the traders to incorporate the uncertainty produced from their own error. For $\hat{c}_{2t} = c_2^*(\hat{c}_{1t})$, but $\hat{c}_{1t} \neq c_1^*(n_t)$, the conditional variances of (25) and (26) become

$$\sigma_F^2(n_t, \hat{c}_{1t}) = \left((1-\beta)\left(\frac{R}{R-\phi}\right)^2 + b_2^2(n_t, \hat{c}_{1t})\right)\sigma_\epsilon^2 \tag{40}$$

$$\sigma_M^2(n_t, \hat{c}_{1t}) = \left(\left(\frac{R}{R-\phi} - \hat{c}_{1t} b_2(n_t, \hat{c}_{1t})\right)^2 + b_2^2(n_t, \hat{c}_{1t})\right)\sigma_\epsilon^2, \tag{41}$$

reflecting the dependence of the market-based strategy on the accuracy of the employed forecast model parameters. Larger estimates of c_{1t} feed greater market-based trader uncertainty, attenuating market-based trader demand. The dashed lines designated $c_1^{+'}$ and $\tilde{c}_1^{'}$ in Figure 2 reflect the greater stability in the market as each is everywhere above the respective c_1^+ and $\tilde{c}_1$. There is a value $\bar{n}^{'}$ such that for $n_t > \bar{n}^{'}$ the invalid price region does not exist. Thus, regardless of how large is $\hat{c}_{1t} > R$, there exists a positive market clearing price. Similarly, there is a value of $n^{+'}$ such that for $n_t > n^{+'}$ the market-based strategy is always profitable.

3 Simulations

Computational analysis reveals properties of the market model not accessible using analytical tools. Simulations facilitate analysis of the impact of different implementations of bounded rationality.

Let p_t^1, the strong-form efficient market price, be the standard against which the market

Table 1: Simulation parameter settings

Simulation	Figures	n_t process	δ or ρ	λ	μ	$\hat{c}_{2t}$	$\hat{\sigma}_t^k$
1	4 & 5	DCD	10	$1/t$	$1/t$	$c_2^*(\hat{c}_{1t})$	exp
2	6	RD	0.01	$1/t$	$1/t$	$c_2^*(\hat{c}_{1t})$	exp
3	7	RD	1	$1/t$	$1/t$	$c_2^*(\hat{c}_{1t})$	exp
4	8	RD	0.05	$1/t$	1	$c_2^*(\hat{c}_{1t})$	$\sigma^*(n_t)$
5	9	RD	0.01	0.01	$1/t$	$c_2^*(\hat{c}_{1t})$	exp
6	10	RD	0.01	$1/t$	$1/t$	estimate	exp
Shared Parameters: $R = 1.02$, $\phi = 0.5$, $\sigma_\epsilon = \sigma_e = 1 \Rightarrow \beta = 1/2$							

price is evaluated. Let $|p_t - p_t^1|$ be the measure of market inefficiency. In general

$$p_t - p_t^1 = (b_{1t} + \phi b_{2t} - \phi/(R-\phi))d_t + (b_{2t} - 1/(R-\phi))\varepsilon_{t+1}. \tag{42}$$

Equation (42) reveals two sources of deviation from efficiency. The condition $b_{1t} + \phi b_{2t} = \phi/(R-\phi)$, producing zero for the first term, requires only $\hat{c}_{2t} = c_2^*(\hat{c}_{1t})$ without constraint on $\hat{c}_{1t}$. To generate $b_2^*(n_t) \to 1/(R-\phi)$ in the second term additionally requires $\hat{c}_{1t} = c_1^*(n_t)$ and $n_t \to 0$. A non-zero value in the first term indicates market-based trader error induced mispricing of public and private information. Non-zero values in the second term reflects a failure by the market to set price to fully reflect d_{t+1}, allowing d_t to enter into price determination. As reference, consider a market populated by only fundament traders. In this case,

$$p_t^F - p_t^1 = -\left(\frac{1-\beta}{R-\phi}\right)\epsilon_{t+1}$$

All simulations share a starting value, $n_0 = 0.75$ and the parameter values $R = 1.02$, $\phi = 0.5$, and $\sigma_\epsilon = \sigma_e = 1$ so that $\beta = 1/2$. Pre-simulation learning on the market-based model takes place on 200 observations generated using a fixed $n_t = n_0$. At these parameters, $\text{Stdev}(p_t^F - p_t^1) = 0.96$.

Figures 4 through 10 display examples of the evolution of endogenous parameters typical of the treatment. To aid direct comparison, each figure is based on the same underlying randomly generated $\{d_t\}$ series.

3.1 Level population dynamics

The local stability of the fixed point under the DCD implementation of a LPD is assured if the traders employ $\mu_t = 1/t$ in their performance updating. Figure 4 shows the early convergence of the system towards the fixed point values of the respective parameter. Figure 5 shows the asymptotic properties of the convergence. The early evolution includes periods of high volatility in the pricing error each time $\hat{c}_{1t} > c_1^+(n_t)$. Asymptotically, the variance of the pricing error appears uniform as $\hat{c}_{1t} \to c_1^*(n^{fp})$.

Increasing ρ has two effects on the asymptotic position of the market. Increasing ρ decreases n^{fp}. This moves the point of attraction deeper into a region, increasing the price impact of market-based model error, as reflected by the narrowing of the distance between $c_1^*(n_t)$ and $\tilde{c}_1(n_t)$ in Figure 2. Also, the greater sensitivity to differences in performance increases the magnitude of the swings in n_t around n^{fp} as a consequence simply of the random dividend process. The combination increases the time it takes for the market to converge on the fixed point.

Similar to increasing ρ, shortening memory of past performance also generates large swings in n_t. The difference is that the swings do not decrease with the accumulation of experience. In order to obtain convergence to the fixed point with $\mu_t = 1$ requires dampening the population's response to observed performance differentials, accomplished here with $\rho < 0.5$. Otherwise, without the tempering of response that comes with the accumulation of knowledge, the traders are incapable of preventing the low realizations of n_t that put the market in the invalid price region.

3.2 Innovation population dynamics

The RD process offers a point of attraction at $n_t = 0$ and $\hat{c}_{1t} = R$ but not a fixed point. If the system were able to travel along $c_1^*(n_t)$ as $n_t \to 0$, then the market would produce increasing price efficiency with $p_t \to p_t^1$.

3.2.1 Baseline

The baseline setting imposes the REE(n_t) solution parameter constraints, $\hat{c}_{0t} = 0$ and $\hat{c}_{2t} = c_2^*(\hat{c}_{1t})$ on market-based trader behavior. The baseline also employs the long memory of least-

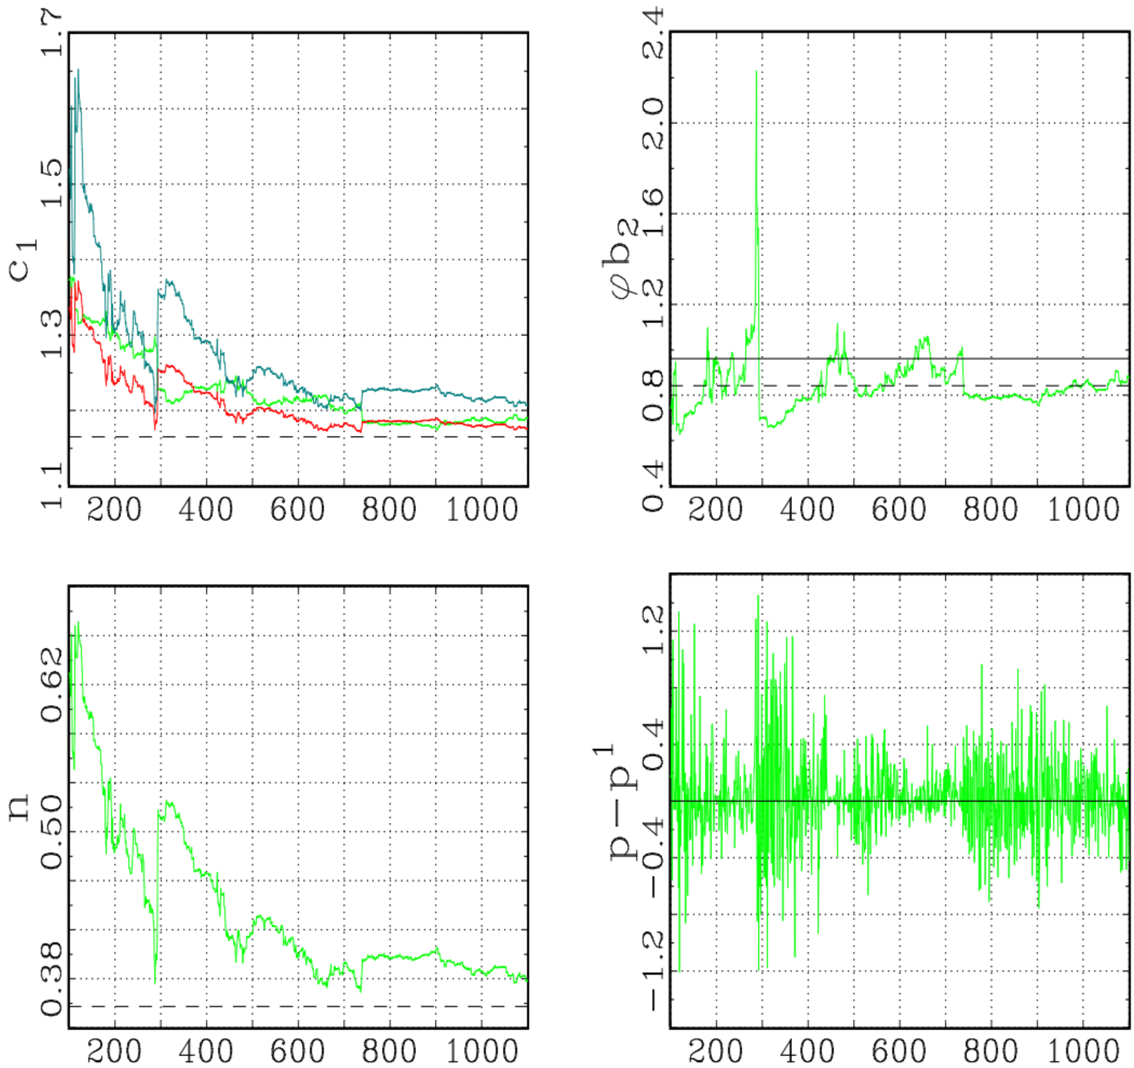

Figure 4: DCD produces convergence towards a REE(n^{fp}) fixed point with $n^{fp} = 0.357$. Top left plots c_{1t} (green), $c_1^*(n_t)$ (red), and $c_{1t}^+(n_t)$ (cyan). Top right plots ϕb_{2t} with a solid line at $\phi b_{2t} = \phi/(R - \phi)$. Lower left plots n_t. Lower right plots $p_t - p_t^1$. In all frames, dashed lines indicate fixed point values, $c_1^*(n^{fp})$, $b_2^*(n^{fp})$, and n^{fp} as appropriate.

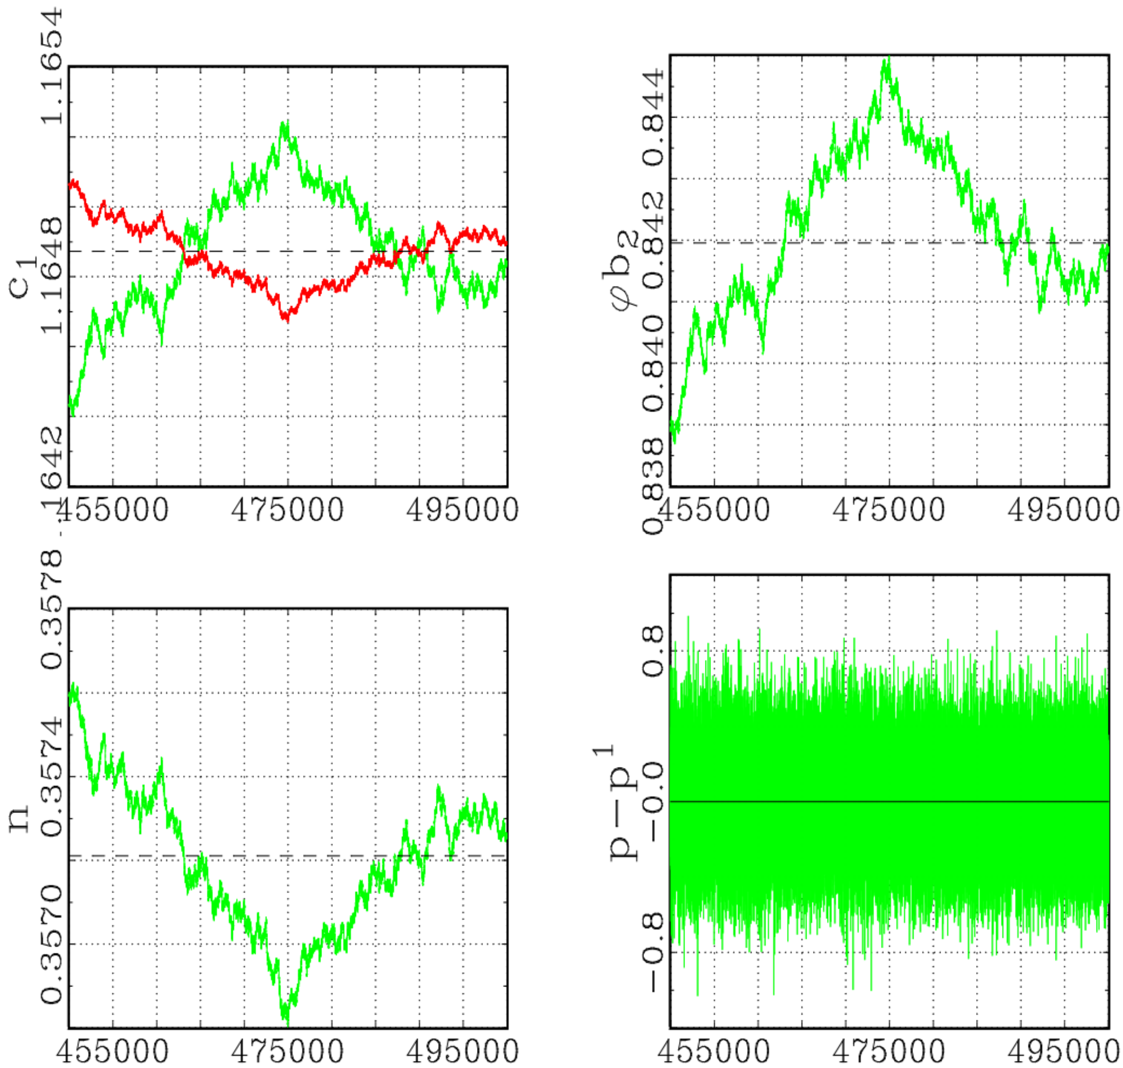

Figure 5: DCD asymptotic behavior around a REE(n^{fp}) fixed point with $n^{fp} = 0.357$. Top left plots $\hat{c}_{1t}$ (green) and $c_1^*(n_t)$ (red). Top right plots ϕb_{2t} with a solid line at $\phi b_{2t} = \phi/(R-\phi)$. Lower left plots n_t. Lower right plots $p_t - p_t^1$. In all frames, dashed lines indicate fixed point values, $c_1^*(n^{fp})$, $b_2^*(n^{fp})$, and n^{fp} as appropriate.

squares learning with λ_t, μ_t, and θ_t all set to $1/t$. Additionally, a low δ produces a slow evolution in the population towards the higher performing strategy. These features make the baseline setting conducive to asymptotic convergence towards the point of attraction. Observed in Figure 6, the system adheres closely to $c_1^*(n_t)$ as $n_t \to 0$. The estimated $\hat{c}_{1t}$ remains well below $\tilde{c}_1(n_t)$, also included in the plot. Despite this apparent success in convergence, the system fails to produce $p_t \to p_t^1$. The system instead generates clustered volatility in the pricing error with no indication of increased accuracy over time. The magnitude of the pricing errors coincide with the magnitude of deviation in ϕb_{2t} from $\phi/(R-\phi)$. The b_{2t} deviations, driven by deviations in $\hat{c}_{1t}$ from $c_1^*(n_t)$, are not independent across time but instead produce a time-series with a highly persistent estimated AR(1) coefficient of 0.998.

3.2.2 Responsiveness

Increasing δ generates fairly regular oscillations in n_t while preserving an underlying process of convergence towards the point of attraction. Figure 7 captures this phenomenon. In contrast to the baseline, the swings in $c_1^*(n_t)$ are larger than the variance in c_{1t}. The highly responsive population produces exaggerated changes in the relative popularity of the two strategies. These changes outpace the slow improvement in $\hat{c}_{1t}$ as t becomes large.

Major price disruptions occur when there is a mis-match between $\hat{c}_{1t}$ and n_t. The swings produced by a large δ would seem to invite such outcomes but instances fail to materialize. Convergence continues without major price disruptions because the system self-regulates the rate of decline in n_t so that it does not outpace the rate of adjustment in $\hat{c}_{1t}$. Also, the incremental changes in n_t within the cycles are relatively small, so that the system delivers feedback of an inconsistency between believes and n_t through the profits awarded to the fundamental traders without n_t ever getting too far out of line with beliefs.

3.2.3 Memory in performance

The simulation presented in Figure 8 substitutes $\mu_t = 1$ for the long memory of $\mu_t = 1/t$. With this change, the treatment closely resembles the model analyzed using the phase space presented in Figure 2.[11] Short memory halts the convergence of n_t towards zero, with n_t instead

[11] To make the comparison complete, and with no discernible impact on the simulations, the traders measure conditional variance using (25) and (26) rather than the experience-driven (36).

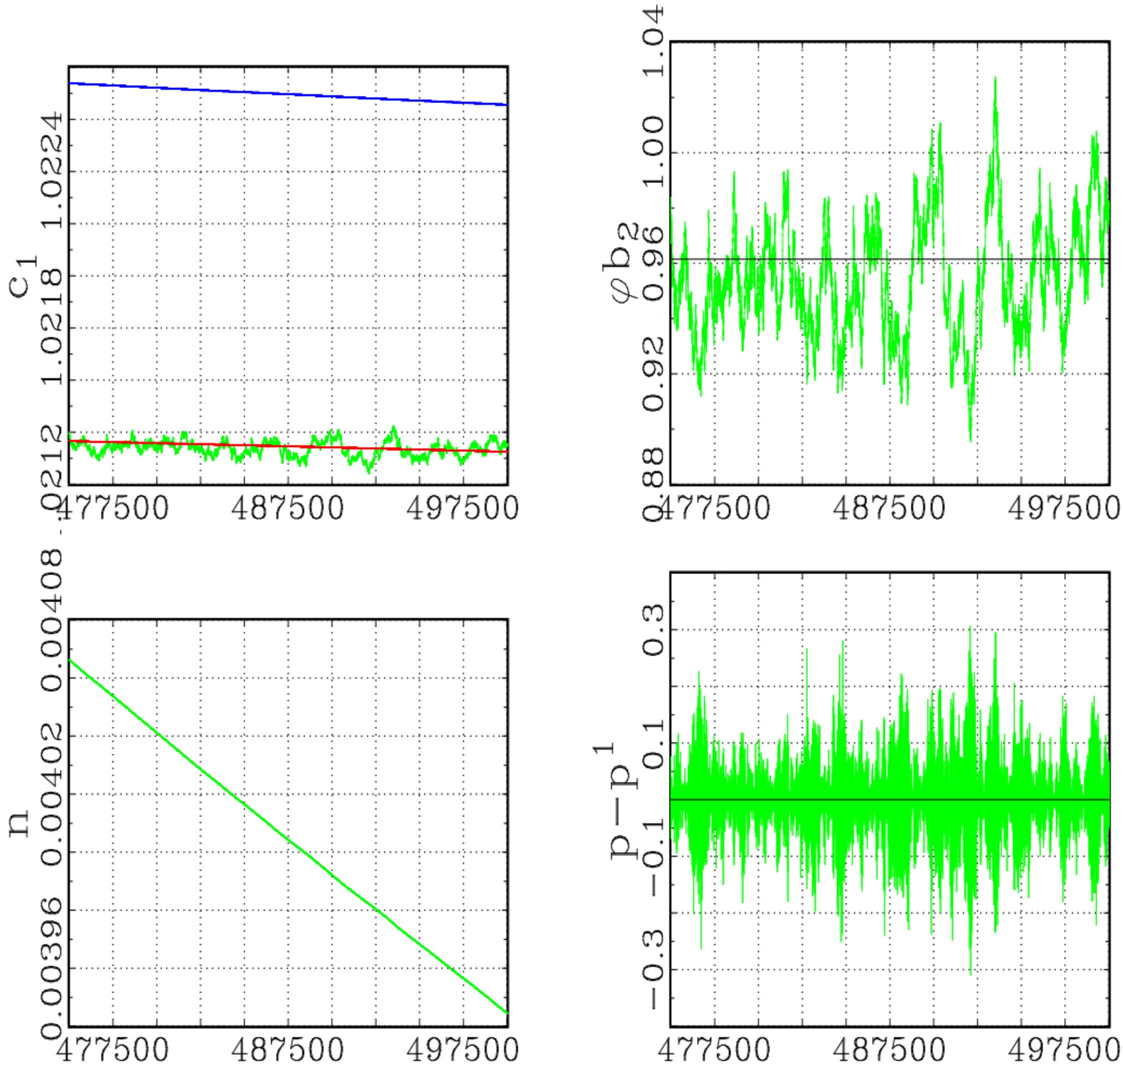

Figure 6: Baseline RD produces smooth convergence $n_t \to 0$ with $\hat{c}_{1t} \to c_1^*(n_t)$. Top left plots $\hat{c}_{1t}$ (green), $c_1^*(n_t)$ (red), and $\tilde{c}_{1t}(n_t)$ (blue). Top right plots ϕb_{2t} with a solid line at $\phi b_{2t} = \phi/(R - \phi)$. Lower left plots n_t. Lower right plots $p_t - p_t^1$.

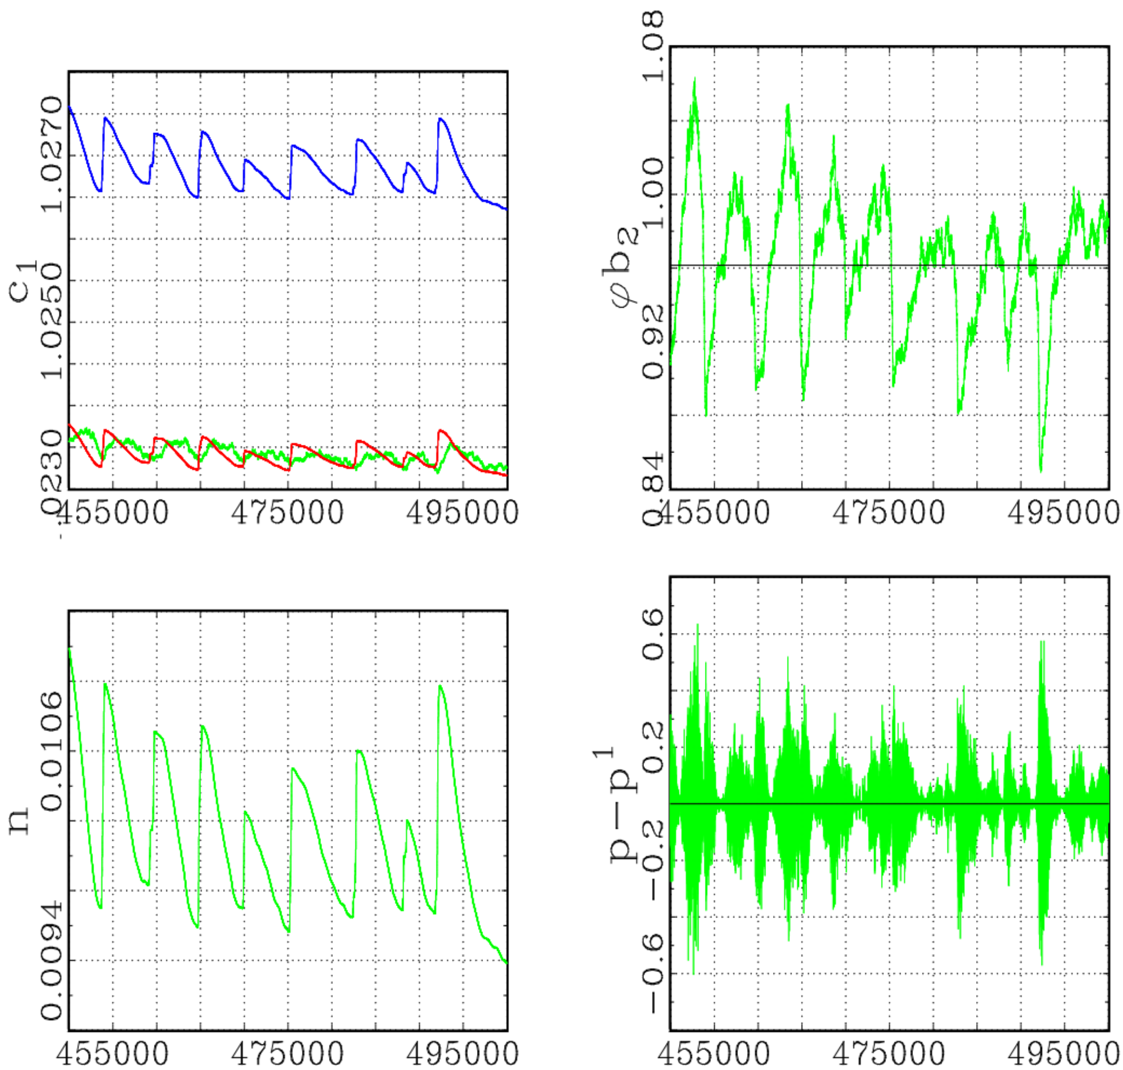

Figure 7: RD with high sensitivity to performance, with $\delta = 1$, produces oscillations in n_t overlaying its general decreasing trend. Top left plots $\hat{c}_{1t}$ (green), $c_1^*(n_t)$ (red), and $\tilde{c}_{1t}(n_t)$ (blue). Top right plots ϕb_{2t} with a solid line at $\phi b_{2t} = \phi/(R - \phi)$. Lower left plots n_t. Lower right plots $p_t - p_t^1$.

hovering around 0.39. At time t, the d_{t+2} component of p_{t+1} remains unpredicted by the market. Payoffs are thus not perfectly forecastable. With a short memory, profit realizations from the unpredictable component of p_{t+1} generate large incremental movement in n_t. This movement undermine the learning of $\hat{c}_{1t}$ and ultimately the convergence of n_t. With large jumps in n_t, incompatibility between n_t and $\hat{c}_{1t}$ arise without prior performance feedback that could prevent the over-use of the market-based model.

While both increasing δ and shortening memory for performance tend to increase swings in n_t, each acts differently on the system. As observed, increasing δ generates cycles in n_t that impact pricing but do not threaten large pricing errors, asymptotically. Short memory produces sudden jumps in n_t rather than smooth swings. Substantial price errors occur when a suddenly low n_t is incompatible with the concurrent $\hat{c}_{1t}$. With short memory a permanent feature of the population, these instances of pricing error do not decline with accumulated experience and thus are a permanent hinderance to convergence.

The stabilization of the system around a fixed n gives the system the appearance of possessing a fixed point in the underlying dynamic system, similar to that produced by a LPD process.

3.2.4 Memory in model

A constant gain in the updating of the market-based model parameters ensures persistence of error in the market-based model. With $\lambda_t = 0.01$, Figure 9 reveals that after a period of learning, $\hat{c}_{1t}$ settles into a stable distribution relative to $c_1^*(n_t)$, moving over time to track the slow evolution in $c_1^*(n_t)$. For sufficiently large n_t, the narrow distribution in $\hat{c}_{1t}$ favors the market-based model. The constant gain becomes a liability as n_t converges towards zero, the distribution in $\hat{c}_{1t}$ is at some point too wide to remain between c_1^- and c_1^+. The resulting mispricing rewards the fundamental model, reversing the progress in n_t. The measured $\hat{\pi}^F - \hat{\pi}^M$ remains positive for some time after the return to near-fundamental pricing until the accumulation of small profits earned by the market-based strategy outweighs the substantial fundamental trader profit earned during the period of mispricing.

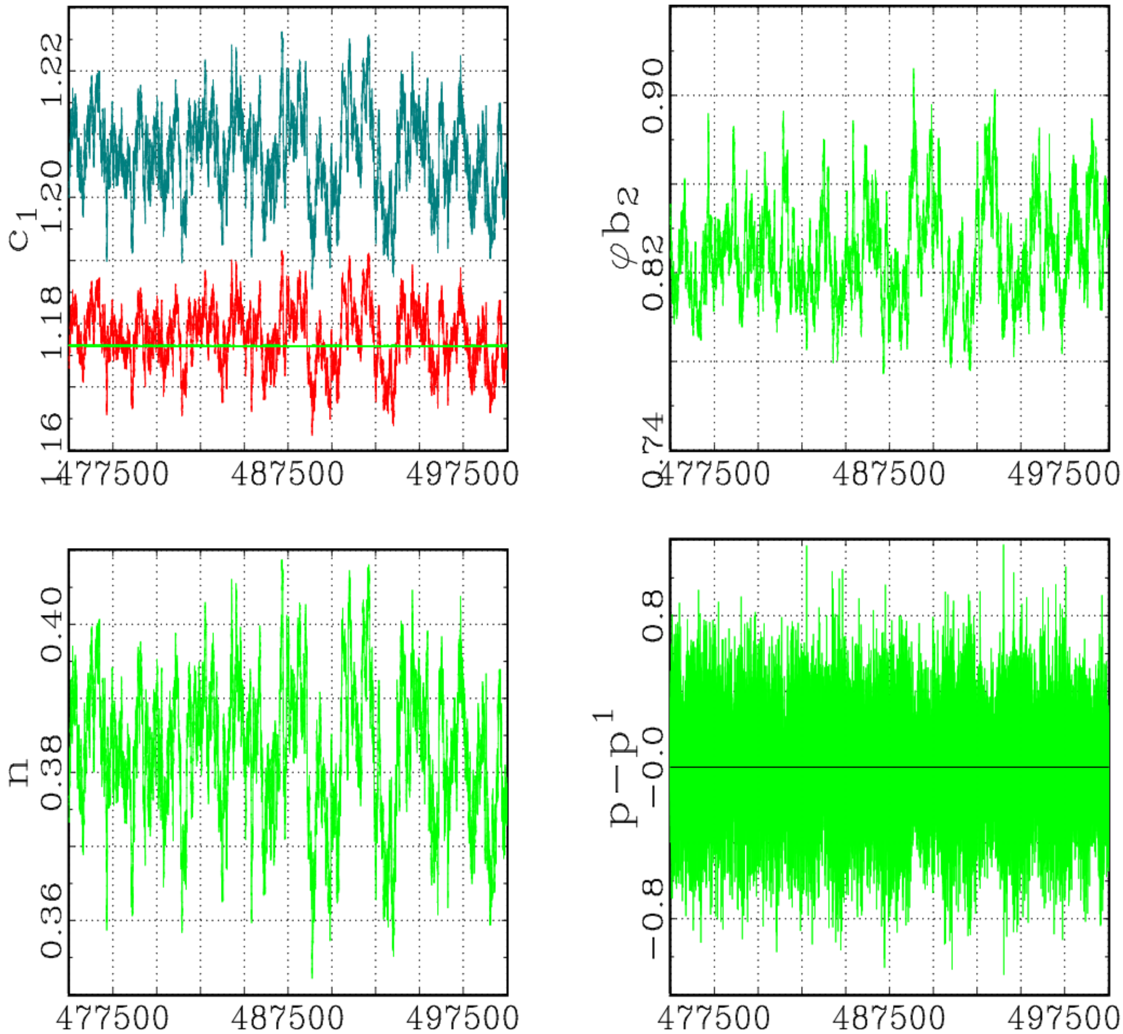

Figure 8: RD with $\mu_t = 1$ stabilizes n at a value above zero. The estimate $\hat{c}_{1t}$ is stable over time while $c_1^*(n_t)$ fluctuates rapidly with the fluctuations in n_t. Top left plots $\hat{c}_{1t}$ (green), $c_1^*(n_t)$ (red), and $c_{1t}^+(n_t)$ (cyan). Top right plots ϕb_{2t} with a solid line at $\phi b_{2t} = \phi/(R-\phi)$. Lower left plots n_t. Lower right plots $p_t - p_t^1$.

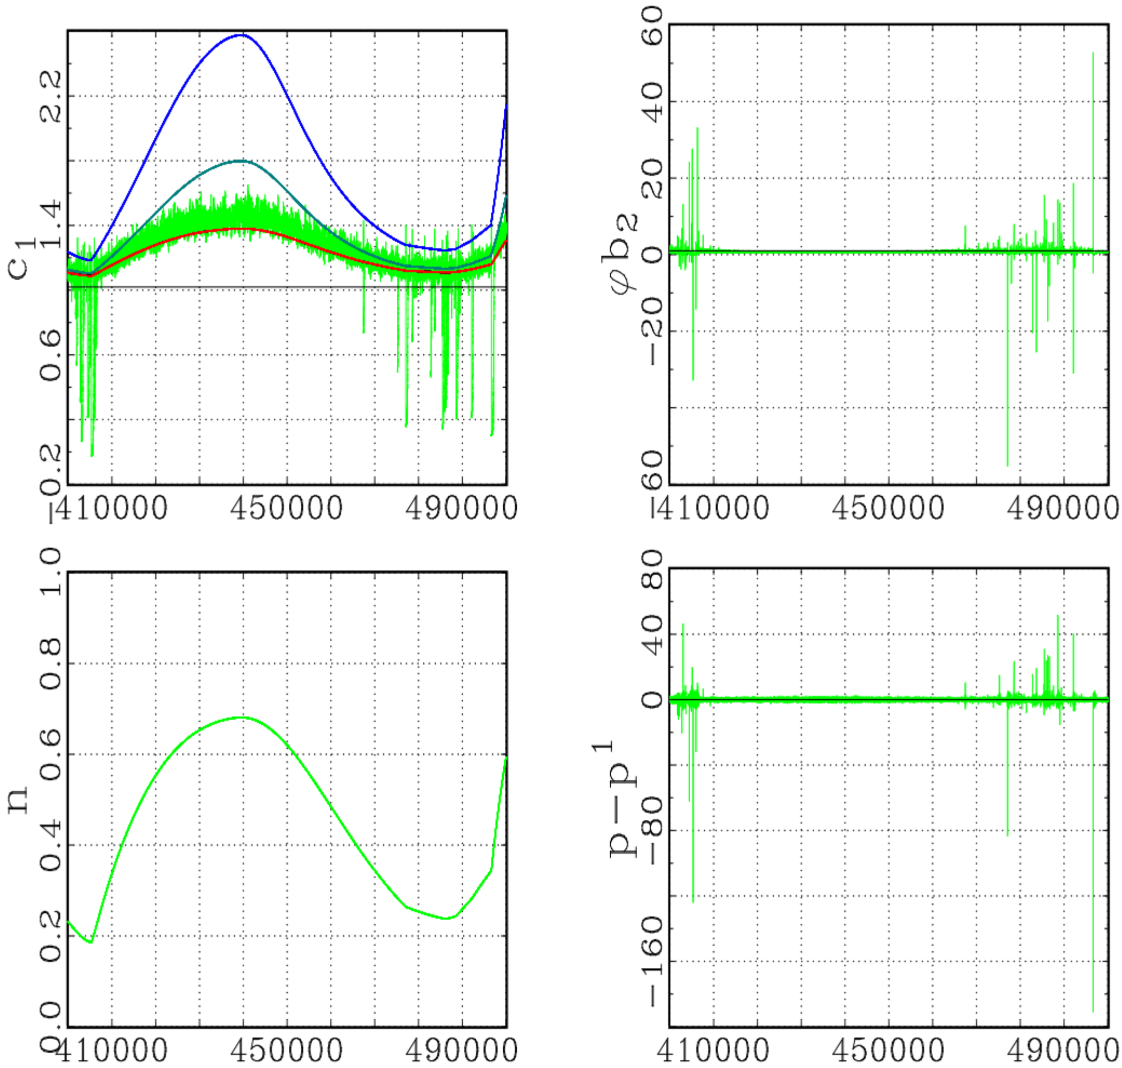

Figure 9: RD with $\lambda_t = 0.01$ generating small but consistent model error that produces long periods of near efficient pricing with inevitable bursts of mispricing. Top left plots $\hat{c}_{1t}$ (green), $c_1^*(n_t)$ (red), $c_{1t}^+(n_t)$ (cyan), and $\tilde{c}_{1t}(n_t)$ (blue). Top right plots ϕb_{2t} with a solid line at $\phi b_{2t} = \phi/(R-\phi)$. Lower left plots n_t. Lower right plots $p_t - p_t^1$.

3.2.5 REE(n_t) conditions

Relax the rationality of the trader by decoupling $\hat{c}_{2t}$ from $c_2^*(\hat{c}_{1t})$ so that the market-based strategy estimates both c_{1t} and c_{2t} through the learning process of (34) and (35). The consequences are twofold. Possible inconsistency in trader estimation such that $c_{2t} \neq c_2^*(c_{1t})$ is a source of error to the market-based model, resulting in a nonzero coefficient on d_t in (18), that feeds back into the price such that $b_{1t} + \phi b_{2t} \neq \phi/(R - \phi)$. The result is a non-zero coefficient on the first term of the price deviation equation, (42). That $b_{1t} + \phi b_{2t} \neq \phi/(R - \phi)$ also undermines the fundamental traders' understanding of price formation as they rely on this feature of price in forming expectations. The consequence of $b_{1t} + \phi b_{2t} \neq \phi/(R - \phi)$ is also reflected in the conditional forecast error in (17) producing non-zero coefficient on ϕd_t.

Figure 10 includes frames for $\hat{c}_{2t}$ on the left and of $b_{1t} + \phi b_{2t}$ on the right. The setting preserves the steady decline in n_t, suggesting continued improved accuracy in the market-based model despite the need to estimate two parameters with only minimal hindrance. Relative to the base simulation, price deviations from the efficient price show greater volatility with greater clustering that coincides with deviations from $b_{t1} + \phi b_{2t} = \phi/(R - \phi)$ and $\phi b_{2t} = \phi/(R - \phi)$.

4 Conclusion

The developed market model contains a tension between a strategy of relying on imperfect fundamental information and that of seeking to optimally exploit the information content of market phenomena. Unique to this model, the market-based alternative to the fundamental information is capable of earning a profit while improving market efficiency. The finding is that the market cannot support exclusive use of a single information source in equilibrium. Employing both fundamental and market information is supported, either as equilibrium behavior that tolerates unequal return performance or in a perpetual state of disequilibrium produced by profit-chasing behavior. The latter is found to be capable of producing substantial pricing error, depending on the behavior of the market participants. Data overcomes deficiencies in trader knowledge when traders rely on increasingly long histories to inform their decisions, producing a well-behaved, though not necessarily efficient, market. Regimes in which traders place greater emphasis on more recent outcomes undermine market efficiency and allow other deficiencies to affect the price. Responding to recent performance differentials also masks the differences in

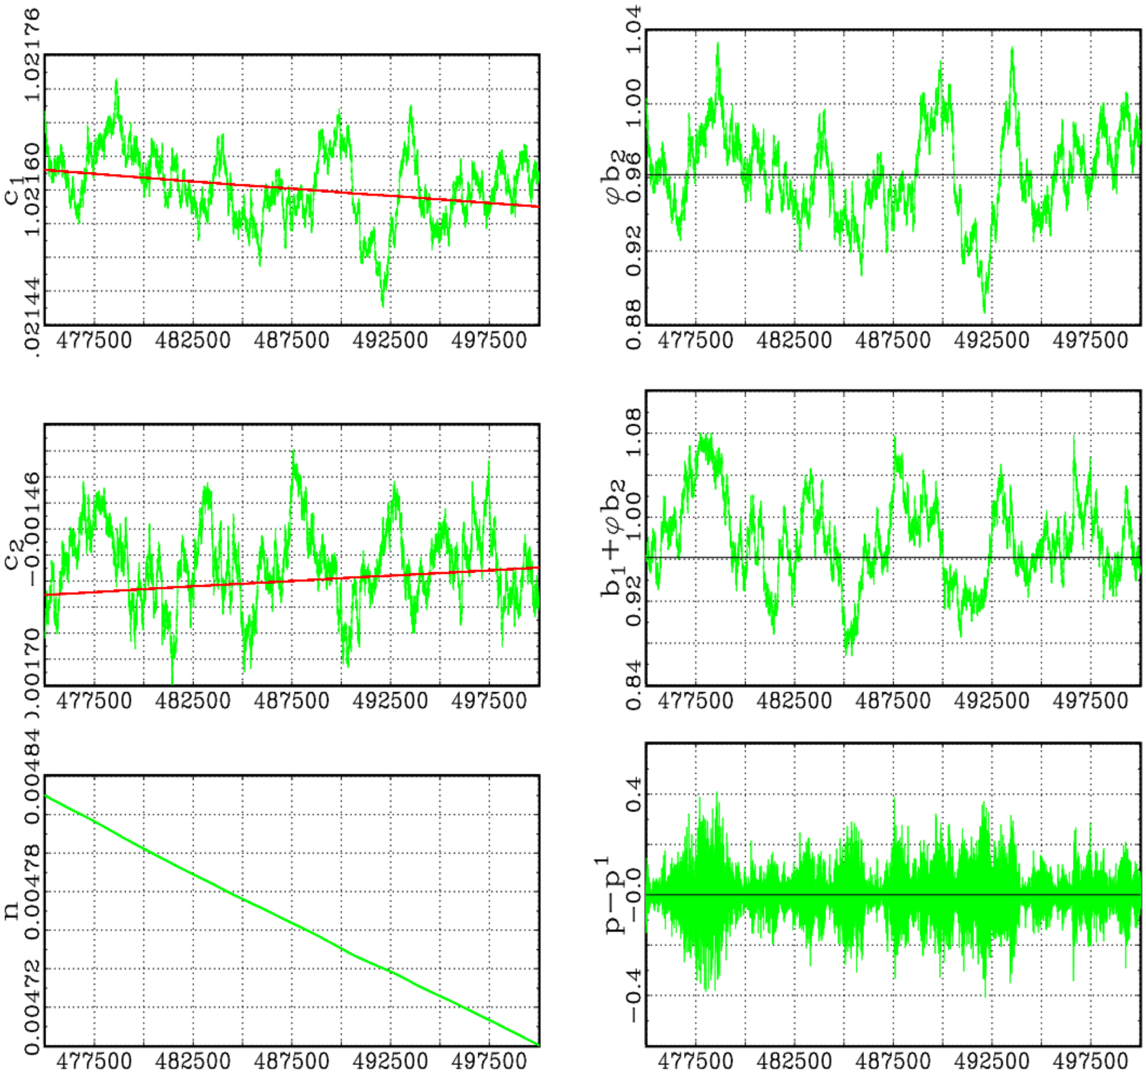

Figure 10: RD with $\hat{c}_{2t}$ allowed to differ from $c_2^*(\hat{c}_{1t})$ producing error in the pricing of observable and unobservable components of price. Top left plots $\hat{c}_{1t}$ (green) and $c_1^*(n_t)$ (red). Top right plots ϕb_{2t} with a solid line at $\phi b_{2t} = \phi/(R-\phi)$. Middle left plots $\hat{c}_{2t}$ (green) and $c_2^*(n_t)$ (red). Middle right plots $b_1 + \phi b_{2t}$. Lower left plots n_t. Lower right plots $p_t - p_t^1$.

the underlying process driving popularity as the system seems to settle around a fixed mix of fundamental and market-based traders independent of the process.

How closely agents adhere to the constraints imposed by rationality alters the behavior of the market. Greater rationality among the market-based traders improves their ability to extract information from the price and improves market performance. Limits on learning may improve near-term performance in a given state but leads to evolution in the state that prevents the market from achieving anything approaching asymptotic efficiency. Greater rationality among the fundamental traders has the potential to undermine market performance to the extent that it builds on inconsistencies already present in the market-based trader behavior.

References

Branch, W., Evans, G. 2006. Intrinsic heterogeneity in expectation 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 & Organization, 65, 224–244.

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.

Brock, W. A., LeBaron, B. D. 1996. A dynamic structural model for stock return volatility and trading volume. The Review of Economics and Statistics, 78, 94–110.

Bullard, J., Duffy, J. 2001. Learning and excess volatility. Macroeconomic Dynamics, 5, 272–302.

Bullard, J. 1994. Learning equilibria. Journal of Economic Theory, 64, 468 – 485.

Bullard, J., Duffy, J. 1999. Using genetic algorithms to model the evolution of heterogeneous beliefs. Computational Economics, 13, 41–60.

Chiarella, C., Dieci, R., Gardini, L. 2006. Asset price and wealth dynamics in a financial market with heterogeneous agents. Journal of Economic Dynamics & Control, 30, 1755–1786.

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

Chiarella, C., He, X.-Z. 2003. Dynamics of beliefs and learning under AL-processes—the heterogeneous case. Journal of Economic Dynamics and Control, 27, 503–531.

De Grauwe, P., Dewachter, H., Embrechts, M. 1993. Exchange rate theory: Chaotic models of foreign exchange markets. Blackwell, Cambridge, MA.

De Long, J. B., Shleifer, A., Summers, L. H., Waldmann, R. J. 1990a. Noise trader risk in financial markets. Journal of Political Economy, 98, 703–738.

De Long, J. B., Shleifer, A., Summers, L. H., Waldmann, R. J. 1990b. Positive feedback investment strategies and destabilizing rational speculation. The Journal of Finance, 45, 379–395.

Droste, E., Hommes, C., Tuinstra, J. 2002. Endogenous fluctuations under evolutionary pressure in cournot competition. Games and Economic Behavior, 40, 232–269.

Evans, G. W., Ramey, G. 1992. Expectation calculation and macroeconomic dynamics. The American Economic Review, 82, 207–224.

Evans, G., Honkapohja, S. 2001. Learning and expectations in macroeconomics. Princeton University Press.

Fama, E. F. 1970. Efficient capital markets: A review of theory and empirical work. The Journal of Finance, 25, 383–417.

Farmer, J., Joshi, S. 2002. The price dynamics of common trading strategies. Journal of Economic Behavior & Organization, 49, 149–171.

Frankel, J. A., Froot, K. A. 1990. Chartists, fundamentalists, and trading in the foreign exchange market. The American Economic Review, 80, 181–185.

Giardina, I., Bouchaud, J. 2003. Bubbles, crashes and intermittency in agent based market models. The European Physical Journal B-Condensed Matter and Complex Systems, 31, 421–437.

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

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

Goldbaum, D., Panchenko, V. 2010. Learning and adaptation's impact on market efficiency. Journal of Economic Behavior & Organization, 76, 635–653.

Goldbaum, D. 2003. Profitable technical trading rules as a source of price instability. Quantitative Finance, 3, 220–229.

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

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

Hellwig, M. F. 1980. On the aggregation of information in competitive markets. Journal of Economic Theory, 22, 477–498.

LeBaron, B., Arthur, W., Palmer, R. 1999. Time series properties of an artificial stock market. Journal of Economic Dynamics and Control, 23, 1487–1516.

Lux, T. 1995. Herd behaviour, bubbles and crashes. The Economic Journal, 105, 881–896.

Lux, T. 1998. The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fat tails of return distributions. Journal of Economic Behavior & Organization, 33, 143–165.

Marcet, A., Sargent, T. 1989a. Convergence of least-squares learning in environments with hidden state variables and private information. The Journal of Political Economy, 97, 1306–1322.

Marcet, A., Sargent, T. 1989b. Convergence of least squares learning mechanisms in self-referential linear stochastic models. Journal of Economic Theory, 48, 337–368.

Panchenko, V., Gerasymchuk, S., Pavlov, O. V. 2013. Asset price dynamics with heterogeneous beliefs and local network interactions. Journal of Economic Dynamics and Control, 37, 2623–2642.

Parke, W. R., Waters, G. A. 2007. An evolutionary game theory explanation of arch effects. Journal of Economic Dynamics and Control, 31, 2234–2262.

Routledge, B. R. 1999. Adaptive learning in financial markets. Review of Financial Studies, 12, 1165–1202.

Sciubba, E. 2005. Asymmetric information and survival in financial markets. Economic Theory, 25, 353–379.

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

A Appendix: Proof of Proposition 3

Proof. Under the regularity conditions (see Marcet and Sargent (1989b), p342-343), the stability of the learning process with $\lambda_t = 1/t$ can be established from the stability of $T(c) - c$ where $T(c)$ maps c into the projection coefficients. From (22) and (23),

$$c_1 = \frac{R}{R-\phi}\frac{1}{b_2}$$

and

$$c_2 = \frac{\phi}{R-\phi}(R - c_1)$$

so that, according to (13),

$$T(c_1) = \frac{nR + (1-n)(R-c_1)\kappa^2}{n\beta}$$

and

$$T(c_2) = -\frac{\phi}{R-\phi}\left(\frac{nR(1-\beta) + (1-n)(R-c_1)\kappa^2}{n\beta}\right)$$

The eigenvalues of the Jacobian matrix, $\frac{\partial[T(\mathbf{c})-\mathbf{c}]}{\partial \mathbf{c}}$, are $\left\{-1,\ -1-\frac{1-n}{n}\frac{1}{\beta}\kappa^2\right\}$, which are both less than zero. The learning process is thus locally stable so that $\Pr(|\mathbf{c}_t - \mathbf{c}^*| > \psi) \xrightarrow{a} 0$ for $\psi > 0$. $\square$