Volume 35, Issue 1

Learning to Agree: A New Perspective on Price Drift.

Andrea Giusto
Dalhousie University

Abstract

This paper introduces statistical learning in an asset pricing model of differences of opinions. I show that the model converges to the unique rational-expectation equilibrium and furthermore I show that asset prices drift predictably in its neighborhood. Accordingly, the model offers a unifying perspective between two so-far mutually exclusive strands of the asset pricing literature. Learning preserves all the desirable features offered by the rational-expectations hypothesis (i.e. the traders use efficiently both the private and the public information available) while yet implying asset prices that drift predictably in the ex-ante sense of Banerjee et al. (2009).

1   Introduction

Asset prices are inertial: an increase in the price of a stock in the past is predictive of future increases, and this fact is often referred to as "price drift." This well-known empirical regularity was first documented by Jegadeesh and Titman (1993) and in order to explain this finding, the theoretical literature has emphasized the slow aggregation process of heterogeneous beliefs held by the market participants. The empirical literature confirms that belief heterogeneity is associated with drift in prices (see Verardo (2010) and Hommes (2011) for example). There exist two competing modeling paradigms to incorporate belief heterogeneity in asset pricing models: first, the rational-expectations approach stresses the presence of private noisy information that is gradually incorporated into the asset's prices, and second, the differences-of-opinion approach that assumes that the traders disregard public information and focus on their private information: they simply "agree to disagree." Both approaches have advantages and limitations. The differences-of-opinion approach is capable of accounting for the existence of price drift, but it requires strong assumptions about the lack of rationality of economic agents. The rational-expectations approach instead has proved unable to explain drift in the empirically relevant *ex-ante* sense defined in Banerjee, Kaniel, and Kremer (2009).

This paper offers a unifying perspective through a model of learning. I show that a model featuring diverse beliefs converges to the the representative-agent rational-expectation equilibrium when the traders learn which beliefs provide better forecasts. I also show that, under certain conditions, the price dynamics are characterized by inertia. The learning process yielding these results is new: at the end of each trading period, the traders collect and use both the public information (the market clearing price) as well as some limited information about the distribution of beliefs in the population of traders. More specifically, random pairs of traders are assumed to mutually disclose their beliefs and use the market clearing price to rank these beliefs based on likelihood ratios. Given initial conditions, the learning process evolves endogenously: the distribution of beliefs in the traders' population determines the market clearing price of the asset, and in turn this determines the new distribution of beliefs. This learning model therefore allows for a significant reduction in the free parameters required by rational-expectations models, since there is no necessity of a myriad of external sources of information emitting the private signals received by the traders.

The learning literature on asset prices has focused on a variety of topics related to this paper.[1] Brock and Hommes (1998), Brock, Hommes, and Wagener (2005), Branch and McGough (2008), Waters (2009), and Brock, Hommes, and Wagener (2009) use evolutionary model selection dynamics that are closely related to the dynamics of this paper. These papers conduct bifurcation analyses of various alternative population dynamics and they all find that apparently simple dynamics of selection of beliefs may reach chaotic complexity. All of these papers assume that beliefs selection dynamics are governed by past profitability, rather than maximum likelihood. De Grauwe and Markiewicz (2013) compare these two alternative approaches in the context of forecasting currency exchange rates. Asset pricing models with statistical learning are also considered by Branch and Evans (2006), who introduce the concept of a misspecification equilibrium: belief heterogeneity may persist in an equilibrium

---

[1] See Evans and Honkapohja (2013) for an extensive survey of this literature.

provided that all of the alternative belief systems are equally misspecified. Branch and Evans (2011) show that in a framework in which agents learn about the risk-return trade-off there are recurring bubbles in the asset prices. LeBaron (2012) shows that several interesting empirical facts about asset prices can be reproduced by a simple model of learning with heterogeneity, but the issue of price drift is not addressed there. Finally this paper is also related to Weibull (1997, Chapter 5), which provides the basic framework exploited by the model of this paper.

2   A Model of Different Opinions with Learning

A continuum of traders $i \in I = [0, 1]$ live for two periods and have a unit exogenous lifetime endowment. At the beginning of the first period a trader chooses a portfolio allocation between a safe and a risky asset so as to subjectively maximize a CARA expected-utility function. The safe asset is the numeraire, it yields a return r normalized to zero in each period, and it is available in unlimited supply. The risky asset is underwritten by the traders themselves and it has the same characteristics of the safe asset. However, this asset is in zero net supply and its price is determined by a market-clearing condition, thus involving the risk of capital gains and losses. The traders do not have complete knowledge of the distribution of beliefs in the population and thus they treat the next-period price of the risky asset as a random variable with a specific Gaussian distribution, which I will refer to as the trader's beliefs. The set of possible beliefs is discrete and finite and it will be denoted by $\phi = \{\phi_1, \ldots, \phi_K\}$, for finite but possibly large K, where ϕ_k denotes a normal probability density function with mean μ_k and variance σ_k^2. The elements of the set ϕ are all distinct, and there is no belief that may lead to a default on the risky asset. Let $x_t = (x_{1,t}, \ldots, x_{K,t})' \in \Delta$ denote the distribution of traders over ϕ, where Δ is the K dimensional simplex. A trader active at time t, having beliefs ϕ_k maximizes the expected utility of wealth, consistent with their belief. The solution to this well-known portfolio problem is given by

$$\lambda_{k,t} = \frac{\mu_k - p_t}{\gamma \sigma_k^2}. \tag{1}$$

where γ is the coefficient of risk-aversion, and $\lambda_{k,t}$ is the demand of a trader with beliefs k for the risky asset when its price is p_t. The market-clearing condition for the risky asset $\left(\sum_{k=1}^K x_{k,t} \frac{\mu_k - p_t}{\gamma \sigma_k^2} = 0 \right)$ defines implicitly the function $p : \Delta \mapsto \mathbb{R}_+$ that yields the market clearing price. This function is linear in the population state x_t. After trading takes place and the market clears, the price is broadcast publicly, while each trader receives a private noisy signal about a belief system randomly selected from the population. At the beginning of their second period the traders sell their contracts and they are replaced by another young trader with the same beliefs.[2] Before retirement, each trader $i \in I$ - endowed with beliefs $\phi_k \in \phi$ - observes beliefs ϕ_j from another trader randomly drawn from the population.[3] The

---

[2] This set up captures the portfolio selection process in a principal–agent framework with asymmetric information.

[3] The assumption that only one alternative belief system is observed by each trader is not very restrictive: allowing for a wider sampling would only complicate the notation and the speed of the beliefs selection process.

relative likelihood of the market price p_t according to each belief determines which one is chosen by trader i and the original beliefs are maintained if and only if

$$\zeta_i \phi_k(p_t) \geq \phi_j(p_t) \tag{2}$$

where ζ_i is a real random variable with uniform cumulative distribution Φ having a positive support centered on one. This random factor can be thought of as an idiosyncratic component in the trader's priors, or alternatively as noise affecting the observation of the performance of the alternative beliefs. From the mathematical stand point, the idiosyncratic shocks ζ_i serve the purpose of removing the discontinuity implied by the choice between pairs of beliefs. The behavioral assumption embodied in equation (2) is a natural choice that is consistent with the emphasis in the statistics literature on maximum likelihood estimation and likelihood ratio tests. Here I posit this a natural principle for revising forecasts.

The Population Dynamics. Fix $\phi_k \in \phi$ and consider the change in the share of traders adopting beliefs ϕ_k between time t and $t+1$. This change is given by the flow of traders that adopt ϕ_k minus the traders that drop ϕ_k to adopt some other belief. The probability that a trader with beliefs $\phi_i \in \phi$ adopts ϕ_k conditional on sampling beliefs k is $\rho_{t,i\to k} = \frac{x_k \phi_k(p_t)}{\sum_{j=1}^{K} x_j \phi_j(p_t)}$. In fact by Bayes law this probability is equal to the unconditional probability $x_{t,k} \Phi\left(\frac{\phi_k}{\phi_i}\right)$ divided by the constant normalizing factor $\sum_{j=1}^{K} x_{t,j} \Phi\left(\frac{\phi_j(p_t)}{\phi_i(p_t)}\right)$. The probability $\rho_{t,i\to k}$ is then obtained by exploiting the linearity of Φ. It follows that

$$x_{t+1,k} - x_{t,k} = \sum_{i=1}^{K} x_{t,i}\rho_{t,i\to k} - \sum_{i=1}^{K} x_{t,k}\rho_{t,k\to i}$$

$$x_{t+1,k} - \cancel{x_{t,k}} = \sum_{i=1}^{K} x_{t,i}\rho_{t,i\to k} - \cancel{x_{t,k}} \underbrace{\sum_{i=1}^{K} \rho_{t,k\to i}}_{1}$$

$$x_{k,t+1} = \frac{\phi_k(p_t)}{\sum_j x_{j,t}\phi_j(p_t)} x_{t,k}, \quad k = 1,\ldots,K. \tag{3}$$

which is the discrete version of the well-known replicator dynamics. These equations, together with the pricing function p define a map $F : \Delta \mapsto \Delta$ which governs the population dynamics and it will be the subject of the rest of the paper.

3   Rational-Expectation Equilibrium and Its Stability

Under the assumption of a representative rational trader, the market equilibrium of this model would most simply imply no trading of the risky asset and its price would be undetermined. This is because the two assets offer the same returns and the agent's risk aversion implies that utility is maximized by trading only in the non-risky

asset.[4] Furthermore, the zero net-supply for the risky asset necessarily implies that the representative agent's net demand for it is zero at the equilibrium. In other words, in this framework there is no fundamental price for the risky asset, but rather only an equilibrium quantity: differently from most asset pricing models, in this model there is no exogenous stochastic process governing the returns offered by the risky asset, and hence one cannot calculate the fundamental price as the present discounted value of dividends. Rather, the returns offered by the risky assets are endogenous to the model. These considerations motivate the definition that follows.

Definition 1. *The representative-agent Rational-Expectation Equilibrium (REE) for the model of differences of opinions with learning is attained at time T if $\lambda_{k,t} = 0$ for all $t \geq T$, and all $k = 1, \ldots, K$ such that $x_{k,t} > 0$.*

The rationale behind this definition is the equivalence in observable market outcomes, since when traders do not take positions in the risky asset the market outcome is observationally identical to that implied by one rational representative trader. An important consequence of this definition is that the exact composition of the set of beliefs K does not matter for the REE: since exists no "true" fundamental price for the risky asset, there is no reason to worry that the set K may not contain an element that is consistent with the price at a rational expectations equilibrium.

Lemma 1. *Each of the K vertices of Δ is both a steady-state for the population dynamics (3).*

   proof: Suppose that population state $\bar{x}_{t_0}$ is such that $\bar{x}_{k,t_0} = 1$ for some $k \in \{1, \ldots, K\}$ and t_0. Clearly, equations (3) imply that $x_{j,t} = 0$ for all $j \in \{1, \ldots, K\}, j \neq k$ and for all $t > t_0$, hence $\bar{x}_{t_0}$ is a steady-state. The market clearing condition implies that at $\bar{x}_{t_0}$ the price is μ_k and equation (1) implies that $\lambda_{k,t} = 0$ for all $t > t_0$. $\square$

Remark 1: Lemma 1 does not imply the existence of multiple rational-expectations equilibria. Rather, all the vertices of the K-dimensional simplex represent possible instances of the *same* REE: according to Defintion 1, this multiplicity is *irrelevant* to the market outcome and hence it does not yield equilibria multiplicity.[5]

Remark 2: The REEs on the vertices of Δ are those stable situations in which traders learn to agree, and hence they are of central interest for the purpose of this paper.

Definition 2. *Let d denote the Euclidean metric in the simplex Δ. A population state $x \in \Delta$ is asymptotically stable if there exists a positive real number ε and $y \in \Delta$ such that*

$$d(x, y) < \epsilon \to \lim_{n \to \infty} d(F^n(x), F^n(y)) = 0$$

---

[4] The terminology here seems to lead to a paradox: if there is a representative trader then there could be no disagreement and hence the notion of "risky asset" loses its meaning. This logical difficulty is avoided through the standard assumption that the representative agent does not take into account its own role on the aggregate dynamics. In the specific terms of the model considered here, under this assumption the asset in limited supply bears a risk from the perspective of the representative trader while being effectively riskless at the equilibrium.

[5] The existence of multiple instances of the same equilibrium has its roots in the notion of an equilibrium with a representative and rational agent where there is no trading in the risky asset and its price is indeterminate. This price indeterminacy corresponds to the multiplicity of instances of the same REE implied by Lemma 1.

where F^n denotes the n-th iteration of F.

The following proposition shows that only the REE equilibria located on particular vertices of the simplex are asymptotically stable, while interior rest points are not stable in this sense.

Proposition 1. *Let $\widetilde{\phi} \subseteq \phi$ be the set of pdfs that are maximal in a neighborhood of their respective mean*

$$\widetilde{\phi} := \{\phi_k \in \phi : \phi_k(p) \geq \phi_j(p), \forall \phi_j \in \phi, |p - \mu_k| \leq \epsilon, \text{for some } \epsilon > 0\}.$$

Population states for which $x_k = 1$ are asymptotically stable under F if and only if $\phi_k \in \widetilde{\phi}$. Other non-cyclical population states that are steady-states of F are not asymptotically stable.

proof: Clearly the finiteness of ϕ implies that $\widetilde{\phi}$ is not empty. Denote with $\bar{x}$ a population state in which the whole mass of traders adopts beliefs k, that is $\bar{x}_k = 1$ where $\phi_k \in \widetilde{\phi}$. The risky asset's price that clears the market in this case is μ_k, the mean of the distribution ϕ_k, as seen from the pricing function p. Since the pricing function p is linear, it is possible to find a small number $\epsilon > 0$ such that any population state for which $d(x, \bar{x}) < \epsilon$ implies that $x_k = 1 - \delta$ for some $\delta > 0$, and $p(x)$ is close to μ_k so that ϕ_k is still maximal with respect to all the other beliefs. Hence, by denoting the k-th component of $F(x)$ as $[F(x)]_k$, the map F implies that $[F(x)]_k > x_k$ and $|p(F(x)) - \mu_k| < |p(x) - \mu_k|$. Because ϕ_k is symmetric around μ_k, the new price $p(F(x))$ is still in the region where ϕ_k is maximal, and it follows by induction that $F^n(x) \to \bar{x}$ as $n \to \infty$. On the converse, if $\phi_k \notin \widetilde{\phi}$ an arbitrarily small ϵ that keeps the price where ϕ_k is maximal does not exist and this concludes the first part of the proof.

Consider now, a point $\breve{x}$ not on a vertex of Δ, such that $F(\breve{x}) = \breve{x}$. The necessary condition for this to be an equilibrium is $\phi_k(p(\breve{x})) = \phi_j(p(\breve{x}))$ for all k and j such that $\breve{x}_k, \breve{x}_j > 0$. By the assumption that ϕ_k and ϕ_j are not identical it follows that *any* x in an arbitrarily small neighborhood of $\breve{x}$ is such that either $\phi_k(p(x)) > \phi_j(p(x))$ or $\phi_k(p(x)) < \phi_j(p(x))$. Hence $\breve{x}$ is not stable under F. It is straightforward to extend this argument to situations with more than two beliefs. $\square$

4   Price Drift

According to Banerjee et al. (2009), Prices exhibit predictable drift if the size of time-t price change is positively correlated with the size of the change at $t+1$. The following proposition shows that prices drift in this sense in this model.

Proposition 2. *Prices drift predictably in some neighborhoods of the REE with agreement.*

proof: Let $\bar{x}$ be such that $\bar{x}_k = 1$ for some k satisfying $\phi_k \in \widetilde{\phi}$. By Lemma 1, $\bar{x}$ is an REE with agreement. By Proposition 1 there exists a neighborhood of $\bar{x}$ such that for any x in this neighborhood the population dynamics converge to it $(x \to \bar{x})$. This ensures that the following sequence $\{F^n(x)\}_{n=0}^{\infty}$ of populations states is entirely contained in an interval

where ϕ_k is maximal. Proposition 1 shows that such a point x exists. Again, by denoting the k-th component of $F(x)$ as $[F(x)]_k$, equations (3) imply that

$$[F^n(x)]_k > [F^{n-1}(x)]_k > \cdots > [F(x)]_k > x_k \tag{4}$$

Consequently, as the share of traders using beliefs k increases the price gets closer and closer to μ_k implying $|p(F^n(x)) - \mu_k| < \cdots < |p(x) - \mu_k|$. The quantities within the absolute value operators have all the same sign, since otherwise (4) would be violated. Accordingly, the sequence of prices $\{p(F^n(x))\}_{n=0}^{\infty}$ is a monotone converging sequence. Therefore, along an asymptotically convergent path, if $n_2 > n_1$, price changes at n_2 are both smaller and preceded by smaller price changes than at n_1, so that larger price movements are followed by larger successive changes. $\square$

5   Conclusions

This paper proposes a model of asset pricing that reconciles the differences-of-opinion paradigm with the rational-expectations hypothesis. Models with differences of opinions do not consider learning but this paper shows that filling this gap is desirable. First, in a model with learning, traders can use efficiently all the information they possess and it is no longer necessary to assume that traders simply discard any information beyond their own private "signals." Second, the learning dynamics are consistent with many well-documented empirical facts about asset prices. Models without learning sacrifice either of these *desiderata*.

The main empirical motivation and focus of this paper is price drift, a deeply puzzling regularity in a representative-agent rational-expectations world. The mechanism that this paper proposes to explain is based on the relaxation of the assumption of a representative agent: price drift has the elements of a self-fulfilling prophecy in the model of this paper. Competition for accurate forecasts makes the traders adopt beliefs that better explain the price data, and *at the same time*, prices are the reflection of this selection mechanism and naturally tend to confirm the most popular beliefs, as a result of the trading activity.

References

Banerjee, S., Kaniel, R., and Kremer, I. (2009). Price drift as an outcome of differences in higher-order beliefs. *Review of Financial Studies*, *22*(9), 3707–3734.

Branch, W., and Evans, G. (2006). Intrinsic heterogeneity in expectation formation. *Journal of Economic theory*, *127*(1), 264–295.

Branch, W., and Evans, G. (2011). Learning about risk and return: A simple model of bubbles and crashes. *American Economic Journal: Macroeconomics*, *3*(3), 159–191.

Branch, W., and McGough, B. (2008). Replicator dynamics in a cobweb model with rationally heterogeneous expectations. *Journal of Economic Behavior and Organization*, *65*(2), 224–244.

Brock, W., and Hommes, C. (1998). Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic dynamics and Control*, *22*(8-9), 1235–1274.

Brock, W., Hommes, C., and Wagener, F. (2005). Evolutionary dynamics in financial markets with many trader types. *Journal of Mathematical Economics*, *41*, 7–42.

Brock, W., Hommes, C. H., and Wagener, F. O. O. (2009). More hedging instruments may destabilize markets. *Journal of Economic Dynamics and Control*, *33*(11), 1912–1928.

De Grauwe, P., and Markiewicz, A. (2013). Learning to forecast the exchange rate: Two competing approaches. *Journal of International Money and Finance*, *32*, 42–76.

Evans, G. W., and Honkapohja, S. (2013). *Learning as a rational foundation for macroeconomics and finance* (R. Frydman and P. Edmund S., Eds.). Princeton University Press.

Hommes, C. (2011). The heterogeneous expectations hypothesis: Some evidence from the lab. *Journal of Economic Dynamics and Control*, *35*(1), 1–24.

Jegadeesh, N., and Titman, S. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of finance*, 65–91.

LeBaron, B. (2012). Heterogeneous gain learning and the dynamics of asset prices. *Journal of Economic Behavior & Organization*, *83*(3), 424–445.

Verardo, M. (2010). Heterogeneous beliefs and momentum profits. *Journal of Financial and Quantitative Analysis*, *44*(4), 795.

Waters, G. A. (2009). Chaos in the cobweb model with a new learning dynamic. *Journal of Economic Dynamics and Control*, *33*(6), 1201–1216.

Weibull, J. (1997). *Evolutionary game theory*. The MIT press.