Learning to Agree: A New Perspective on Price Drift

by

Andrea Giusto
Dalhousie University

Working Paper No. 2014-02

March 2014

DALHOUSIE UNIVERSITY
Inspiring Minds

DEPARTMENT OF ECONOMICS

DALHOUSIE UNIVERSITY
6214 University Avenue
PO Box 15000
Halifax, Nova Scotia, CANADA
B3H 4R2

Learning to Agree: A New Perspective on Price Drift.

December 29, 2013

Abstract

This paper introduces statistical learning in an asset pricing model of differences of opinions. I show that the model converges globally 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, Kaniel, and Kremer (2009). Furthermore, I obtain a number of new empirically testable hypotheses related to price drift from a series of Monte Carlo explorations of a model of differences of opinion with learning.

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). The notion of belief heterogeneity puts a strain on the representative agent fiction, and it increases the complexity of the models since it becomes necessary to deal with the problem of "forecasting the forecasts of others" as the intuition originally formulated by Keynes (1936) was put by Townsend (1983).

Currently, there are two alternative modeling paradigms to incorporate belief heterogeneity in asset pricing models: on the one hand, the rational-expectations approach stresses the presence of private noisy information that is gradually incorporated into the asset's prices, and on the other hand, the differences-of-opinion approach that assumes that the traders disregard public information and focus on their private information: they simply "agree to disagree."

These approaches are largely viewed as incompatible by the current literature, yet both present advantages and limitations. The rational-expectations model of Allen, Morris, and Shin (2006) shows that in a dynamic setting with repeated trading the presence of beliefs heterogeneity breaks the martingale property of expected prices, thus generating a drift in the price of an asset towards its liquidation value. Importantly, this kind of drift is shown to exist in an *ex-post* sense, i.e. it is *conditional* on the liquidation value of the asset. As argued by Banerjee et al. (2009), the drift that properly explains the empirical regularities, must be in an *ex-ante* sense (i.e. it must not be conditional on future observables) and they show that a rational-expectations approach *cannot* produce this outcome. On the other hand, models of differences-of-opinions will deliver price drift, under certain conditions, but their tractability requires very strong and unrealistic assumptions. In particular, in such models it is typically assumed that (i) the private signals are common knowledge in the traders' population and (ii) the traders stick to their beliefs, irrespective of their forecasting performance. These assumptions challenge the very idea of rationality: if other traders' beliefs are fully known and they lead to *better* predictions, then rational traders should immediately adopt the most accurate belief systems.[1]

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 –in a statistical sense– and discard the poorer ones to instead adopt the stronger beliefs. I also show that, under certain conditions, the price dynamics are characterized by inertia in the *ex-ante* sense of Banerjee et al. (2009). 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

[1] A strong argument in favor of this view is based on evolutionary grounds: those traders that do not adopt the "correct" expectational model will be impoverished and driven away from the trading floor – see Friedman (1953) for an early and cogent exposition this concept.

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.[2] 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) present 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 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), providing a basic framework exploited by the model of this paper, and to De Long, Shleifer, Summers, and Waldmann (1990), considering non-rational traders. While in this paper there is no distinction between rational and "noise" traders, the traders of the model presented in Section 2 will not arbitrage away all deviations of prices from their subjective expected values, very much based on the same intuition uncovered by De Long et al. (1990): as long as the traders perceive a risk that others may disagree with them, they will be willing to only take finite positions in the risky asset.

The paper proceeds as follows: Section 2 contains a mathematical formulation of the model, Section 3 presents the convergence results, and Section 4 contains the local results on price drift.

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

Section 5 provides an in-depth comparison with the relevant existing literature and Section 6 extends the analysis of the model through a series of Monte Carlo experiments. Section 7 concludes.

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.[3] 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

$$\max E_k \left[-\exp\left(-\frac{1}{2}\gamma W_{t+1} \right) \right]$$

$$W_{t+1} = 1 + \lambda_{k,t} \left(p_{t+1} - p_t \right)$$

where W_{t+1} is terminal wealth, E_k is the expectational operator relative to ϕ_k, γ is the risk aversion parameter, and the trader's endowment has been normalized to one. Exploiting the normality of

[3]Random occurrences of bankruptcies can be fit into this modely by assuming that the supply of the risky asset is stochastic, and that the losses and gains from such occurrences are shared equally by all the traders active in the risky asset market. Analytical tractability requires a non-random supply of the risky asset, however, I will consider this possibility explicitly in Section 6.

the beliefs, the demand for the risky asset of a trader with beliefs k is

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

Equation (1) shows that as long as $\sigma_k > 0$ trader k will not be willing to arbitrage away *all* the deviations from the mean μ_k as in De Long et al. (1990). Imposing the market clearing condition yields

$$\sum_{k=1}^{K} x_{k,t} \frac{\mu_k - p_t}{\gamma \sigma_k^2} = 0 \tag{2}$$

which implicitly solves for the current price of the risky asset p_t. The market clearing price, as defined by (1), is a linear function of x_t which will be denoted by $p : \Delta \mapsto \mathbb{R}_+$. 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.[4]

Before retirement, trader $i \in I$, endowed with beliefs $\phi_k \in \phi$ chooses whether the newly observed beliefs ϕ_j are preferable by calculating the likelihood of the price p_t. The original beliefs are maintained if and only if

$$\zeta_i \phi_k(p_t) \geq \phi_j(p_t) \tag{3}$$

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 (3) 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.

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

The Population Dynamics. The unconditional probability that agent $i \in I$, having beliefs $\phi_k \in \phi$ adopts beliefs $\phi_j \in \phi$ is

$$\rho^i_{t,k\to j} = x_{t,j}\Phi\left(\frac{\phi_j(p_t)}{\phi_k(p_t)}\right)$$

which, given the uncountable number of traders, is independent of i . Therefore while the adoption of competing beliefs is stochastic at the individual level, the aggregate dynamics are nevertheless deterministic. The flows of traders over the various elements of Φ are governed by the following system of difference equations (derivation in the appendix)

$$x_{k,t+1} = \frac{\phi_k(p_t)}{\sum_j x_{j,t}\phi_j(p_t)}\, x_{t,k}, \quad k = 1, \ldots, K. \tag{4}$$

which is the discrete version of the well-known replicator dynamics. These equations, together with the pricing equation (2) define a map from and to the K-dimensional unit simplex $F : \Delta \mapsto \Delta$ 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.[5] Furthermore, the zero net-supply for the risky asset necessarily implies that the representative agent demand for it is zero at the equilibrium. 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: traders may disagree at the equilibrium, as long as their disagreement leads to trading positions that are

[5] 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.

observationally equivalent to a rational-expectations equilibrium with a representative trader. The following lemma is a straightforward consequence of equations (1), (2), and (4).

Lemma 1. *Each of the K vertices of Δ is both a steady-state for the population dynamics (4) and an REE.*

proof: Suppose that population state $\bar{x}_{t_0}$ is such that $\bar{x}_{k,t_0} = 1$ for some $k \in \{1, \dots, K\}$ and t_0. Clearly, equations (4) imply that $x_{j,t} = 0$ for all $j \in \{1, \dots, K\}, j \neq k$ and for all $t > t_0$, hence $\bar{x}_{t_0}$ is a steady-state. Equation (2) implies that at $\bar{x}_{t_0}$ the market clearing price is μ_k and equation (1) implies that $\lambda_{k,t} = 0$ for all $t > t_0$.

□

This Lemma, despite its simplicity, has important implications for the logic of what follows. First, 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.[6] Furthermore, according to Lemma 1, 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. However, there are quite a few other possibilities that may arise from the dynamics implied by equations (4). First of all, there may be steady-states located not on a vertex of Δ, provided that the population shares and their trading positions produce a price that implies that all beliefs are equally likely. Such situations, in general, do not constitute an REE according to Definition 1: in the interior of the simplex where belief heterogeneity persists, the individual demands for the risky asset are in general different from zero. Finally, there could exist cycles for F which may either be REE or not. The objective of the rest of this section is to show that the only asympototically stable situation is the REE in which agents learn to agree.

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*

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

[6]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.

Vertex and Interior Rest Points. I first consider the non-cyclical instances of the steady states of the population dynamics F, and I prove that only the REE equilibria located on particular vertices of the simplex are asymptotically stable, while interior rest points are not.

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

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

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 shown by equation (2). 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 $[F(x)]_k$, equations (2) and (4) imply 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 extended this argument to situations with more than two beliefs present in the population and this concludes the proof.

□

The intuitive reasoning behind the proof of Proposition 1 is illustrated in Figure 1: the thick line is a belief system that does not belong to the set $\widetilde{\phi}$ since it not maximal around its own mean.

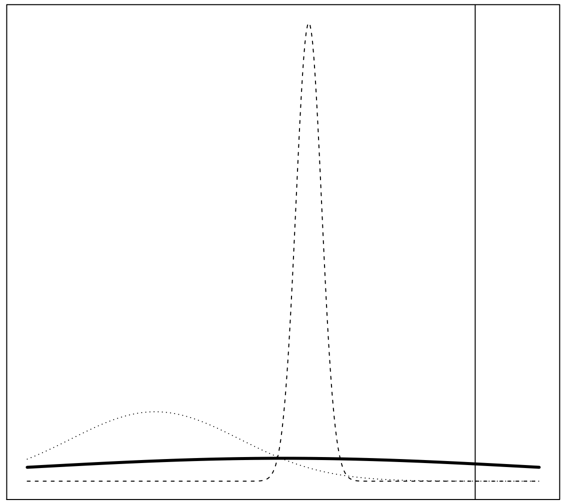

Figure 1: An illustration of the intuition behind Proposition 1.

As the other two distributions depicted in this Figure intersect exactly at its peak, this distribution is not asymptotically stable. On the converse, the dotted and dashed distributions both belong to $\widetilde{\phi}$. Furthermore, while the distribution that is not in $\widetilde{\phi}$ is maximal in other regions of the real line, this has no bearing on its asymptotic stability. In fact, the trading that results from each belief system pushes the market price towards its mean. For this reason, in the conditions illustrated in Figure 1, if the initial price was at the point marked by the vertical line, the population's mass would converge to one of the two beliefs in $\widetilde{\phi}$.

Figure 1 also provides an illustration of the second part of Proposition 1. Suppose that the intersection of the dotted and the dashed beliefs is a rest point of the population dynamics. Clearly any generic change in the population composition, however small, will either increases or decreases the price, and therefore it will unavoidably tilt the balance between the two beliefs, causing the population to amass on either one of them.

Cycles. A straightforward extension of the argument in the second part of Proposition 1 shows that cycles are not asymptotically stable in this model.

Proposition 2. *Let n be a positive finite integer. If there exists an n-cycle of F, it is not asymptotically stable.*

proof: Consider an n-cycle for finite integer $n \geq 2$, such that $x^1 = F(x^n) = F(x^{n-1}) = \cdots = F(x^1)$. Because the dynamics implied by F imply that for any $k = 1, \ldots, K$ the condition $x_{k,t} = 0$ implies $x_{k,t+\tau} = 0$ for $\tau = 1, 2, \ldots$ a non-degenerate cycle cannot involve the vertices of Δ. The second clause of Proposition 1 is therefore directly applicable to each of the population states $x^1, \ldots, x^n$ thus concluding the proof.

□

Propositions 1 and 2 imply that –provided that there are no strange attractors for F– the population dynamics are globally convergent to the REE in which traders learn to agree.[7] This is the first main result of the paper: a parsimonious model of different opinions, without the common knowledge assumption and with learning, converges globally to the REE model with agreement. The section that follows shows that together with this desirable outcome, the model is also consistent with predictable drift in prices.

4 Price Drift

Banerjee et al. (2009) define price drift in terms of conditional expectations. According to their definition, a stochastic price process exhibits drift if $E[p_{t+1} - p_t | p_t - p_{t-1}]$ is increasing in $p_t - p_{t-1}$. Clearly, this definition is based on the understanding that prices tend to increase. However, there is no obvious theoretical reason to exclude that prices may be drifting *downward* if short sales are possible and just as costly as outright purchases, as assumed here. For this reason, I adapt the above mentioned stochastic definition to the deterministic setup considered here and include the possibility of downward drift.

Definition 3. *Prices exhibit predictable drift if larger price changes at t are followed by larger changes at $t + 1$, in the same direction.*

[7] The numerical analysis of Section 6 confirms that chaos is not likely to arise in this model. See Section 5 for a discussion.

Clearly this definition is just a rephrasing of that of Banerjee et al. (2009) for a non-stochastic price process. In Section 6, where I introduce exogenous supply shocks, I will use directly the stochastic definition.

Proposition 3. *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 $\phi_k \in \tilde{\phi}$. By Lemma 1, $\bar{x}$ is an REE with agreement. Let x be in a neighborhood of $\bar{x}$ and furthermore let $x \to \bar{x}$. 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 (4) imply that

$$[F^n(x)]_k > [F^{n-1}(x)]_k > \cdots > [F(x)]_k > x_k \tag{5}$$

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 (5) 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.

$\square$

According to this proposition prices follow a predictable sequence of adjustments of decreasing magnitude toward the REE. The early elements of this sequence are both larger and followed by larger changes than elements that are further along this sequence. Therefore, this sequence constitutes price drift according to Definition 3. Crucially, the notion of predictable drift imposes no constraints of causality: according to Definition 3, there is predictable price drift even if price changes of larger magnitude do not *cause* larger future changes, as long as they are happening in a predictable temporal sequence. This is also true for the stochastic definition of Banerjee et al. (2009) which by requiring that $E[p_{t+1} - p_t | p_t - p_{t-1}]$ is increasing in $p_t - p_{t-1}$ impose constraints exclusively on the correlations produced by the model. In a deterministic setup these correlations are perfect, and Proposition 3 shows that they are caused by iterations of F: belief heterogeneity with learning implies that larger changes in prices are preceded by even larger ones.

5 Discussion

This section contains a detailed comparison of the results of Sections 3 and 4 with the existing literature. First, Proposition 3 indirectly confirms that prices do not drift in an REE, as argued by Banerjee et al. (2009). The model of section 2 does not yield drift at an asymptotically stable equilibrium, however, drift is present in a possibly small neighborhood of the equilibrium. I will illustrate numerically in Section 6 that the limitations of local analysis are not too restrictive here, by showing that there are situations in which drift exists on large regions of Δ. Banerjee et al. (2009) prove that models of difference of opinions are consistent with drift in asset prices only under specific conditions about the disagreement among traders. More specifically, they define two orders of disagreement for traders: the first order of disagreement requires traders to share beliefs about the distribution of the private information received by the agents, but to disagree on the joint implications of these signals for the asset price. Disagreement of higher-degree requires that traders have different beliefs about the distribution of the private information. According to these definitions, the model of this paper is characterized by first-order disagreement and according to Proposition 2 of Banerjee et al. (2009), prices should not drift in it. However, the contradiction implied by Proposition 3 is only apparent, and it is explained by the endogeneity of public signal (price) produced by the model of Section 2. In a neighborhood of the REE, the market price is close to the value considered maximally likely by some belief ϕ_k, which imply that those traders that are given knowledge of ϕ_k find it compelling. When the share of the population believing in ϕ_k increases, the price moves even closer to the mean of ϕ_k and this causes endogenous dependence of the price on its previous value. This is a possibility explicitly considered by Banerjee et al. (2009) who agree that price drift may be a result of serial correlation in the public signal. However, they only mention the case in which this serial correlation is introduced exogenously through the noise process rather than produced endogenously by the beliefs selection process.

The dynamics of Brock and Hommes (1998), Branch and McGough (2008), and Waters (2009) are motivated differently from the model of this paper, but they all are similar both in spirit and mathematically to (4). The common thread linking these papers is that all these models display very interesting bifurcations as a certain key parameter of the model increases. Using the terminology of Brock and Hommes (1998), this key parameter is called "intensity of imitation" and it controls how aggressively are the traders seeking to discover the best predictor of future prices, and how keen they are to adopt them. While the role of this parameter has not been explicitly highlighted this far, in this paper it corresponds to the expected value of the random variables ζ_i. It is natural here to assume that on average, the traders are unbiased judges of the likelihood of their beliefs relative to the alternative they observe. However it would be possible to analyze the changes in the dynamics of the model when agents have a bias towards the new beliefs and are more likely to adopt them. While interesting, an analysis of the bifurcations in this model is nevertheless outside the scope of this paper.

There is another important difference from the cited literature concerned with bifurcations in similar models. The traders in this paper condition their choice of beliefs exclusively on the most recent price level, and this is crucial for the analytical tractability of the model. The motivation is the assumption that agents understand that the process generating the price data is not stationary during the learning dynamics, and hence its ergodicity is not guaranteed. Brock and Hommes (1998), Branch and McGough (2008), and Waters (2009) assume that the traders judge the performance of the various beliefs based on a *history* of past returns rather than only on the latest data point. Prompted by this literature, I will also consider the case in which agents use longer data sets to choose beliefs in Section 6.4 where I show the results of several Monte Carlo experiments that suggest that this assumption is not likely to drive the results of this paper.

The concept of a misspecification equilibrium in Branch and Evans (2006) corresponds to an equilibrium located in the interior of the simplex using the terminology of the present model. I find that these equilibria are not asymptotically stable but Branch and Evans (2006) show that there are instances of misspecification equilibria that are learnable when agents use econometric models of learning. In terms of the structure developed here, only by assuming an uncountably infinite set of beliefs ϕ, this model could nest this kind of learning, and most likely, the finiteness of ϕ is instead determinant for the asymptotic instability of the interior steady-states. A more general analysis of the conditions in which equilibria with heterogeneous beliefs are learnable is an interesting topic for future research.

6 Numerical Analysis

In this section I use numerical methods to analyze in further detail the learning model of Section 2. The analytical results of Sections 3 and 4 are local and as such, they provide no indication on whether the neighborhoods in which prices drift are large enough to be empirically relevant. Furthermore, for the sake of tractability, I have not considered so far the presence of random shocks to the price dynamics. However, randomness is a ubiquitous feature of asset pricing models, and therefore I will introduce in the model two separate sources of exogenous shocks. First, consistently with much of the existing literature, I will modify the model to have a random net supply of the risky asset. As suggested in the introduction, it is possible to interpret these shocks as the result of idiosyncratic bankruptcies in the presence of a risk sharing mechanism that redistributes equally any capital gain or loss associated with these events. Second, I will also consider exogenous shocks to the population composition. Such shocks can be interpreted as failures of the laws of large numbers to deliver perfectly deterministic dynamics. Accordingly, one can interpret the magnitude of these shocks as inversely related to the number of traders present in the population. Alternatively, one can think of these shocks as originating from random inflows and outflows to and from the population of traders.

In the remainder of this Section I will seek to answer the following questions: (1) is price-drift a global feature of this model? (2) Do random shocks affect the probability and magnitude of price drift? (3) Is price drift also present if agents have robustness concerns about their portfolios? (4) Is drift more likely and/or stronger when there is more disagreement in the market? (5) Is drift present when traders use longer time series to rank and choose beliefs? (6) Is the belief selection process consistent with the maximization of expected return or is it possible to find more profitable alternatives?

6.1 Non-Local Drift Analysis

To investigate the prevalence of price drift in this model, I conduct a Monte Carlo experiment with 10 different beliefs. Each element $\phi_k \in \phi$, $k = 1, 2, \ldots, 10$, is randomly chosen by drawing μ_k from a uniform distribution on $[1, 3]$ and σ_k from a uniform distribution on $[0.2, 1]$. After choosing a random initial point of the simplex Δ, I iterate on the dynamics (4) until the price for the risky asset reaches a steady state. The following AR(1) functional form is then fit to the data produced by the simulation

$$p_{t+1} - p_t = \alpha_0 + \alpha_1(p_t - p_{t-1}) \tag{6}$$

as to obtain a linear expectation function for $p_{t+1} - p_t$ conditional on $p_t - p_{t-1}$. The average for the parameter α_1 over 10,000 repetitions of these calculations is 0.7685 which implies that drift is a common feature of the asset prices in this model. In the vast majority (94.67%) of the replications the coefficient α_1 is positive thus indicating that under this parameterization drift is globally pervasive feature of this model.

The first row of Table 1 contains these numbers together with the results from experiments that include normally distributed random shocks. The first pair of columns present the results relative to the shocks in the net supply of the risky asset in correspondence of their standard deviation.[8] Clearly, prices drift predictably in the presence of moderate shocks to the net supply of the risky asset, but as these grow in magnitude, both the importance and the prevalence of price drift decrease. The percentage of Monte Carlo replications for which the parameter α_1 is positive and significant at the 5% level drops from 89.02% to 38.23% when the standard deviation of the supply shocks goes from 0.01 to 0.10.

	Supply Shocks		Population Shocks	
$\sigma_{p/x}$	$\bar{\alpha}_1$	$\alpha_1 > 0$	$\bar{\alpha}_1$	$\alpha_1 > 0$
0.00	0.7685	94.67%	0.7685	94.67%
0.01	0.7732	89.02%	0.6108	75.54%
0.05	0.4781	64.86%	-0.0332	15.61%
0.10	0.1761	38.23%	-0.2349	2.13%

Table 1: Global prevalence of price drift. Each row identifies the standard deviation of the shocks hitting the market in each period and each pair of columns reports (i) the average value for the parameter α_1 from equation (6) and (ii) the percentage of Monte Carlo replications for which this parameter is positive and significant at the 5% level.

The second set of columns in Table 1 report the effects of random shocks to the population state in the simplex Δ. These shocks are assumed to displace the current population state x_t in a random direction by a normally distributed random amount.[9] Clearly, the picture conveyed by these results is consistent with that pertaining to supply shocks, although more extreme: when the

[8]The relative magnitude of the standard deviation of these shocks is best framed in relation to the unitary population size under the interpretation that a random supply shock can be taken to represent a share of of the population to default on their obligations.

[9]Obviously, these shocks require that the population state is not pushed outside the unit simplex Δ. I enforce this constraint by assuming that the boundaries of the simplex act as absorbing barriers.

shocks to the population composition are large enough, prices stop drifting on average and only in a small region of the simplex do they drift predictably.

There are two main conclusions that can be drawn from these results. First, these results illustrate that the local nature of Proposition 3 may not be a severe limitation for this model to produce drift. Second, Table 1 suggests two empirically testable hypotheses: prices of assets characterized by frequent changes in their supply of public equity are less likely to drift than prices of assets with more stable ownership. Furthermore, if one interprets the shocks to the population state as failures of the laws of large numbers to deliver deterministic aggregate dynamics it is possible to see that this model predicts that prices of assets with many market participants are more likely to drift than markets in which the number of traders is smaller.

6.2 Robustness and Risk Aversion

In the model of this paper the traders understand that their subjective beliefs about the distribution of the next-period asset price are subject to revisions. Accordingly, the traders may want to adopt a more conservative approach in their trading by making choices that are robust to model misspecification in the sense of Hansen and Sargent (2008). Maenhout (2004, proposition 1) shows that in the environment studied here, robustness concerns take the form of an additive term to the risk aversion parameter. Hence, it is straightforward to investigate the effect of robust portfolio allocations on predictable drift by increasing the value of the parameter γ and adopting a Monte Carlo strategy as in the previous section.

	No Shocks		Population Shocks					
			$\sigma_x = 0.01$		$\sigma_x = 0.05$		$\sigma_x = 0.10$	
γ	$\bar{\alpha}_1$	$\alpha_1 > 0$	$\bar{\alpha}_1$	$\alpha_1 > 0$	$\bar{\alpha}_1$	$\alpha_1 > 0$	$\bar{\alpha}_1$	$\alpha_1 > 0$
2	0.7736	94.68%	0.6099	75.44%	-0.0337	15.40%	-0.2348	2.23%
5	0.7722	94.89%	0.6135	75.58%	-0.0356	16.19%	-0.2361	2.14%
10	0.7651	94.40%	0.6080	75.24%	-0.0281	16.48%	-0.2320	2.25%

Table 2: Prevalence of drift with varying degrees of risk aversion in the presence of no shocks and of shocks to the composition of the population. Each pair of columns reports (i) the average value for the parameter α_1 from equation (6) and (ii) the percentage of Monte Carlo replications for which this parameter is positive and significant at the 5% level.

γ	$\sigma_p = 0.01$ $\bar{\alpha}_1$	$\alpha_1 > 0$	$\sigma_p = 0.05$ $\bar{\alpha}_1$	$\alpha_1 > 0$	$\sigma_p = 0.10$ $\bar{\alpha}_1$	$\alpha_1 > 0$
2	0.7094	83.93%	0.1883	39.43%	-0.1333	13.47%
5	0.4707	63.71%	-0.2297	6.70%	-0.4125	0.17%
10	0.1821	39.05%	-0.4108	0.18%	-0.4756	0.00%

Table 3: Prevalence of drift with varying degrees of risk aversion in the presence of supply shocks. Each pair of columns reports (i) the average value for the parameter α_1 from equation (6) and (ii) the percentage of Monte Carlo replications for which this parameter is positive and significant at the 5% level.

Table 2 shows that robustness concerns on the part of the traders do not have an effect on predictable price drift, in the case of no random shocks and of shocks to the distribution of beliefs in the traders' population. The percentages of iterations in which the autoregressive parameter α_1 is positive and its average do not change significantly as the risk aversion parameter is increased. On the converse, Table 3 shows that the joint presence of robust portfolio allocations and shocks to the supply of the risky asset are sufficient for eliminating predictable drift from this model's parameterization, even for the case of small shocks. This suggests yet another empirically testable hypothesis: according to this parameterization of this model, shocks to the supply of a stock should interact significantly and negatively with risk aversion in a regression predicting price drift.

6.3 Disagreement and the Intensity of Drift

Verardo (2010) finds that the dispersion of opinions of financial analysts is a significant predictor of the returns to a particular stock when interacted with past returns. It is by no means obvious that in the theoretical framework of Section 2 the intensity of the drift should be positively correlated with a measure of disagreement: the constraints imposed by the model's edifice do not rule out the possibility that *fewer* beliefs produce a stronger drift. This section explores this issue numerically to show that the model proposed here is consistent with this empirical regularity.

I measure the disagreement between beliefs ϕ_k and ϕ_j by the Hellinger distance between the two normal distributions, defined as

$$\delta(\phi_k, \phi_j) = \sqrt{1 - \sqrt{\frac{2\sigma_k \sigma_j}{\sigma_k^2 + \sigma_j^2}} \exp\left(-\frac{1}{4}\frac{(\mu_k - \mu_j)^2}{\sigma_k^2 + \sigma_j^2}\right)}$$

This metric is easily computed and it depends on the distributions' variances and averages, which are crucial in the model, unlike other measures such as differential entropy. Using δ as measure of disagreement between a pair of traders, it is possible to build a measure of disagreement in the general population by calculating the expected Hellinger distance of two traders randomly drawn from the population. At any given time t a trader with beliefs ϕ_k has an expected disagreement with another randomly chosen trader from the population of $\delta_t^k = \sum_{j\in\phi} x_{j,t}\delta(\phi_k, \phi_j)$. Therefore, any two random traders, will have an expected disagreement given by

$$\delta_t = \sum_{k\in\phi} x_{k,t}\delta_t^k = \sum_{k\in\phi}\sum_{j\in\phi} x_{k,t}x_{j,t}\delta(\phi_k, \phi_j) \tag{7}$$

I now perform a Monte Carlo exercise similar to the ones detailed above, but this time I fit the following functional form to the process

$$p_{t+1} - p_t = \alpha_1(p_t - p_{t-1}) + \alpha_2\delta_t + \alpha_3(p_t - p_{t-1}) \times \delta_t + \varepsilon_t \tag{8}$$

where the main interest is relative to α_3 which must be positive if the model conforms to the empirical data. Table 4 shows that the disagreement metric increases monotonically as the

Number of Beliefs	$\bar{\delta}$	$\bar{\alpha}_1$	$\bar{\alpha}_2$	$\bar{\alpha}_3$
5	0.3347	-0.2177	0.0022	2.1994
6	0.3543	-0.3103	0.0009	2.3971
7	0.3692	-0.3987	-0.0040	2.5167
8	0.3798	-0.4290	0.0006	2.5725
9	0.3884	-0.4717	0.0011	2.6606
10	0.3943	-0.5134	0.0029	2.7083
11	0.3997	-0.5163	-0.0000	2.6801
12	0.4039	-0.5248	-0.0012	2.6533
13	0.4077	-0.5225	0.0026	2.6786
14	0.4114	-0.5382	0.0018	2.6554
15	0.4138	-0.5441	0.0018	2.6527

Table 4: Dependence of drift on cardinality of the set ϕ. Each row reports the average over 10,000 Monte Carlo replications of the initial measure of disagreement (δ_1) and of the coefficients of (8).

cardinality of the set of beliefs increases. Furthermore according to this table, the interaction term between past price movements and disagreement is positive and it increases monotonically up to 10 different beliefs, at which point it reaches a plateau. The variation in the average of α_3 past this point is consistent with the uncertainty inherently associated with the Monte Carlo nature of these calculations. These results indicate that this parameterization is consistent with the empirical results of Verardo (2010). Furthermore they suggest a new empirically testable hypothesis: increasing the amount of disagreement at "moderate" levels should be expected to yield price drift that is more persistent, while similar increase should not produce the same effects if disagreement is already high in the market.

6.4 More Persistent Dynamics

Analytical tractability has so far imposed the assumption that the traders use only the most recent price of the risky asset to evaluate beliefs. A more realistic assumption regarding this aspect of the model would be to impose that traders have a finite history of past prices available and that they evaluate the joint likelihood of their sample. Table 5 shows that increasing the sample size used by the traders does not alter significantly the picture painted so far.

Sample size	No Shocks		Supply Shocks		Populations Shocks	
	$\bar{\alpha}_1$	$\alpha_1 > 0$	$\bar{\alpha}_1$	$\alpha_1 > 0$	$\bar{\alpha}_1$	$\alpha_1 > 0$
3	0.4772	0.8477	0.5878	0.7747	0.4464	0.6602
5	0.4545	0.8482	0.5728	0.7676	0.4530	0.6614
7	0.4690	0.8536	0.5864	0.7783	0.4503	0.6647

Table 5: Effects of larger samples on drift. The supply and Population shocks columns refer to Monte Carlo experiments with normal shocks with a standard deviation of 0.01.

6.5 Investing With Drift

The portfolio decisions of the traders take drift into account only indirectly, through the likelihood-based belief selection process. However, it is natural to question if a trader that puts more weight on ex-ante drift could earn higher returns, on average. To explore, this issue I implement a series of Monte Carlo experiments in which I track the profits of a randomly selected trader under two alternative heuristic modifications of the optimal portfolio implied by (1).[10] Under the first

[10]The parameterization is the same as in the previous subsections.

modification a trader uses equation (1) to calculate their demand for the risky asset, but they substitute μ_k with $\tilde{\mu}_k = \mu_k + p_{t-1} - p_{t-2}$. The second modification instead has the trader replace μ_k with $\tilde{\mu}_k$ and also σ_k^2 with $\tilde{\sigma}_k^2 = \sigma_k^2 + (p_{t-1} - p_{t-2})^2$. The experiment is repeated 10,000 times and the average capital gain for the trader using the optimal portfolio allocation is zero by definition. However the two modified portfolio allocation rules yield average capital losses of -0.27% and -0.22% on average.[11]

7 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.

The numerical explorations of Section 6 also suggest some new hypotheses that may be tested empirically. First, Table 1 suggests that the prevalence of drift is diminished by randomness in the population dynamics. Under the interpretation that a small number of traders active on a market is the cause of such randomness, the model suggests that drift should be more pervasive for heavily traded assets than for thinner ones. Second, Table 3 suggests that disturbances to the supply of assets should interact with the degree of risk aversion of the typical investor and acting towards reducing the prevalence and magnitude of drift. Finally, according to Table 4 it appears

[11] This result is robust to the presence of small supply and population shocks. However, larger shocks produce more unpredictable results, since the presence of drift is weakened by these shocks, as shown in Table 1.

that the effect of beliefs heterogeneity on price drift is nonlinear and that only at moderate levels of disagreement, an increase in belief heterogeneity is associated with increases in drift.

References

Allen, F., Morris, S., & Shin, H. S. (2006). Beauty contests and iterated expectations in asset markets. *Review of Financial Studies*, *19*(3), 719–752.

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

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

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

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

De Long, J., Shleifer, A., Summers, L., & Waldmann, R. (1990). Noise trader risk in financial markets. *Journal of Political Economy*, *98*(4), 703–738.

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

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

Friedman, M. (1953). *The case for flexible exchange rates.* University of Chicago Press.

Hansen, L. P., & Sargent, T. J. (2008). *Robustness.* 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., & Titman, S. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of finance*, 65–91.

Keynes, J. (1936). *The general theory of interest, employment and money.* London: Macmillan.

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

Maenhout, P. J. (2004). Robust portfolio rules and asset pricing. *Review of financial studies*, *17*(4), 951–983.

Townsend, R. M. (1983). Forecasting the forecasts of others. *The Journal of Political Economy*, 546–588.

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.

Appendices

Derivation of 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 *abandon* ϕ_k for some other belief. The probability that a trader with beliefs $\phi_i \in \phi$ adopts ϕ_k is $\rho_{t,i\to k} = x_{t,k}\Phi\left(\frac{\phi_k(p_t)}{\phi_i(p_t)}\right)$ so that

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

$$x_{t+1,k} - \cancel{x_{t,k}} = \sum_{i=1}^{K} x_{t,i} x_{t,k} \Phi\left(\frac{\phi_k(p_t)}{\phi_i(p_t)}\right) - \cancel{x_{t,k}} \underbrace{\sum_{i=1}^{K} \rho_{t,k\to i}}_{1}$$

Exploiting the linearity of Φ yields

$$x_{t+1,k} = x_{t,k}\phi_k(p_t) \sum_{i=1}^{K} x_{t,i} \Phi\left(\frac{1}{\phi_i(p_t)}\right) \tag{A1}$$

Changing the index k to j and summing over it yields

$$1 = \sum_{j=1}^{K} x_{t+1,j} = \sum_{j=1}^{K} x_{t,j}\phi_j(p_t) \sum_{i=1}^{K} x_{t,i} \Phi\left(\frac{1}{\phi_i(p_t)}\right) \tag{A2}$$

Dividing (A1) by (A2) yields the population dynamics (4):

$$x_{t+1,k} = \frac{x_{t,k}\phi_k(p_t) \cancel{\sum_{i=1}^{K} x_{t,i} \Phi\left(\frac{1}{\phi_i(p_t)}\right)}}{\sum_{j=1}^{K} x_{t,j}\phi_j(p_t) \cancel{\sum_{i=1}^{K} x_{t,i} \Phi\left(\frac{1}{\phi_i(p_t)}\right)}} .$$

Population Dynamics With More Persistence

Suppose that the traders use a series of the past T prices to evaluate beliefs. The probability of the sample assigned by beliefs $\phi_i \in \phi$ is $\phi_i(p_{t-T+1}) \times \cdots \times \phi_i(p_t)$ and therefore a trader with these beliefs will adopt ϕ_k instead if and only if

$$\zeta \phi_k(p_{t-T+1}) \times \cdots \times \phi_k(p_t) \geq \phi_i(p_{t-T+1}) \times \cdots \times \phi_i(p_t)$$

where ζ is a trader-specific idiosyncrasy with uniform cumulative distribution Φ having a positive support centered on one. By defining

$$\tilde{\phi}_i(\{p_{t-\tau}\}_{\tau=0}^{T-1}) = \phi_i(p_{t-T+1}) \times \cdots \times \phi_i(p_t)$$

and using an identical derivation as before yields

$$x_{t+1,k} = \frac{x_{t,k}\tilde{\phi}_k(\{p_{t-\tau}\}_{\tau=0}^{T-1})}{\sum_j x_{t,j}\tilde{\phi}_j(\{p_{t-\tau}\}_{\tau=0}^{T-1})} .$$