Discussion Paper Series

Prices, fundamental values and learning

By

Michele Berardi

Centre for Growth and Business Cycle Research, Economic Studies,
University of Manchester, Manchester, M13 9PL, UK

September 2015
Number 214

Prices, fundamental values and learning*

Michele Berardi
The University of Manchester

September 29, 2015

Abstract

In this paper we show how uncertainty and learning about fundamental values can lead to excess volatility in prices and to volatility clustering in returns, as observed on real markets. The key assumption is that agents use prices, besides an exogenous signal on long run dividends, to infer fundamental values: as the relative weight on the two signals changes endogenously through learning, price dynamics are affected. In particular, periods of high volatility are periods where agents rely more heavily on prices in predicting fundamentals.

Key words: uncertainty, asset prices, learning, herding.

JEL classification: D83, D84, G12.

* Corresponding author: Dr. Michele Berardi, Economics, School of Social Sciences, The University of Manchester; Tel.: +44 (0)1612754834; Email: michele.berardi@manchester.ac.uk.

Prices, fundamental values and learning

1 Introduction

Shiller (1981) famously showed that asset prices display a level of volatility unexplained by the underlaying fundamentals. Numerous studies - for example, see Cont (2001) - have shown that returns are characterized by volatility clustering, a feature that is unexplained by patterns of volatility in exogenous shocks. We show in this work that both excess volatility in prices and volatility clustering in returns can be explained by uncertainty about fundamental values and learning.

One of the key issues in asset pricing is represented by the relationship between market prices and the underlaying fundamental value of an asset. This issue is of utmost importance both for policymakers and for investors. For example, in terms of policymaking, former Federal Reserve Chairman Ben Bernanke wrote (Bernanke and Gertler, 1999):

"Trying to stabilize asset prices per se is problematic for a variety of reasons, not the least of which is that it is nearly impossible to know for sure whether a given change in asset values results from fundamental factors, nonfundamental factors, or both."

It is just as difficult for investors as it is for policymakers to infer fundamental values, and it is on this aspect, and its consequences for asset prices, that we focus in this work. To this end, we merge two lines of literature: one on signal extraction and Bayesian equilibrium, and one on adaptive learning. From the first, we take the basic building blocks to model prices as an endogenous signal for agents, that summarizes the opinion of the market about the value of an asset. From the second, we take the key insight that agents can only learn from observables. Unfortunately, fundamental values are not observable, even after decisions have been made, demands posted and prices realized: we therefore assume that theory helps agents in guiding their learning activity.

Angeletos and Werning (2006) show that, in a coordination game, agents can use prices as an endogenous signal to better predict other agents' actions. In our model, agents use prices as a signal as they summarize the view of other agents about fundamental values. We assume agents are only concerned about the fundamental value of their portfolio, relative to what they paid for it, and so they are not trying to profit from exploiting short term capital gains. Still, because of the uncertainty about fundamental values and the weight to give to different signals, deviations of prices from equilibrium emerge. In particular, when agents put too much weight on prices in the prediction of fundamental values, large non-fundamental movements in prices arise. It is important to stress that our economy is populated only by fundamentalists and there are no speculators or noise traders here trying to gain from short-term trading strategies that have often been indicated as destabilizing and a cause for excess volatility (e.g., De Long et al., 1990).

Most work on adaptive learning and asset prices has focused on uncertainty about future prices. Uncertainty about fundamental values, on the other hand, has been largely neglected, though it seems a key element for investment strategies of fundamentalists, who should want to buy assets that are underpriced and sell assets that are overpriced with respect to their fundamental value. A key issue, of course, is that the fundamental value of an asset is not known and it can only be inferred, or guessed, using observables such as past dividends and prices as signals.

We will assume that agents receive a noisy exogenous signal about the fundamental value of the stock: this could be thought of as news about long run dividends and other information affecting the value of a firm. Besides this exogenous signal, agents also use prices to infer fundamental values. The reason is that prices summarize information of other market participants: if prices are above what an agent expects, given her own private information, she will think that other people have different information about the fundamental value of the stock, and revise her own beliefs accordingly.

Bayesian theory provides us with the optimal relative weight on the two signals: in the first part of the paper we derive such value and discuss its implications for asset prices.

We then depart from the assumption that agents know a priori such optimal weight, which depends on the relative precision of the signals, and instead require agents to learn it through experience. A key issue here is that fundamental values are not observables, so there can not be direct feedback guiding the learning activity. Instead, agents rely on a mix of theory and evidence: we assume they know that the optimal weight depends on the relative variance of the two signals, and learn about such values.

We first prove that, under decreasing gain, learning converges to the Bayesian equilibrium. This is not an obvious result: because prices are endogenous and depend on agents' beliefs, it could be that higher beliefs lead to higher prices and thus to even higher beliefs, in a self-reinforcing destabilizing loop. This doesn't happen, though, because when the variance of prices increases the relative weight on prices is revised downwards, thus helping stabilizing the system.

We then substitute the decreasing gain with a constant gain, which captures the idea that agents are unsure about the stationarity of the environment and allow for changes. This opens the door to the possibility that the estimated weight used by agents fluctuates over time and never settles down to its optimal value. In particular, because of the feedback from prices to beliefs, we observe that at times there is herding in beliefs,[1] where excessive weight is given to prices in the formation of beliefs about fundamental values: it is at such times that volatility in prices increases and actual prices deviate significantly from equilibrium.

As noted before, our economy is populated only by fundamentalists, who trade on differences between actual prices and (expected) fundamental values. We do not introduce chartists, as it is instead common in the literature, to generate non fundamental movements in prices. All deviations of prices from equilibrium come instead from uncertainty and learning about fundamentals.

It is also worth noting that our model is a model of volatilities. We therefore do not try, and are not able, to explain things like bubbles and crashes, which relate to trends, since there are no trends in our model. All disturbances are i.i.d. random processes, fundamental values are constant and there is no source of persistence artificially built in. We leave for future work to enrich our framework with features that can generate bubble-like events.

The plan of the paper is as follows: after a literature review in the remainder of this Section, in Section 2 we introduce the model, present the information structure and derive the Bayesian equilibrium. In Section 3 we introduce adaptive learning: first we establish convergence under decreasing gain, and then present simulations under constant gain. In Section 4 we discuss the main results and insights, and Section 5 concludes. Some technical details are confined to the Appendix.

1.1 Literature review

Volatility properties of prices and returns have often been explained in the finance literature by appealing to behavioral models of interacting agents. Lux (2009) provides a review of works based on such ideas: in particular, interaction between fundamentalists and chartists, models of local interaction between traders, and field effects have been proposed and analyzed. Contrary to this line of enquiry, we do not assume heterogeneity in agents' degree of sophistication or trading strategies.

A growing literature has also been studying the impact of expectations, bounded rationality and learning on asset prices. Brock and Hommes (1998) analyze the impact of evolutionary dynamics in price predictors on price fluctuations; Branch and Evans (2010) consider a

[1] We define precisely the concept of herding in our setting later on in the paper.

setting where agents predict prices by choosing between two underparameterized models and show that multiple equilibria emerge and the model can reproduce regime-switching returns and volatilities similar to those observed in real data; Branch and Evans (2011) propose a model where agents learn about risk and show that escape dynamics from the fundamental price emerge; Hommes and Zhu (2011) use the concept of stochastic consistent expectations equilibrium to explain excess volatility in stock prices; finally, Adam et al (2008) show how adaptive learning on future prices can generate excess volatility. The main difference between our work and this line of literature is that in our model agent are uncertain about the fundamental value of an asset instead of its future price and learn about a signal extraction problem on exogenous and endogenous information. To the best of our knowledge, it is the first attempt to merge Bayesian equilibrium and adaptive learning to understand asset price dynamics.

In terms of signal extraction and Bayesian learning, there is a large literature on global games, where agents face a coordination problem with heterogeneous information: agents usually receive exogenous public and private signals, and need to extract information in order to solve their coordination problem. An analysis instead of the coordination problem with endogenous signal is provided by Morris and Shin (2006). A notable application of this idea to asset prices is Angeletos and Werning (2006), who consider a model where asset prices act as an endogenous signal in a two stage game where agents need to decide whether or not to carry out a speculative attack: the first stage of that model is similar to our setting, though they do not consider the effect on price dynamics of an endogenously changing weight on the signals. Ozdenoren and Yuan (2007) instead derive excess volatility in a setting where asset prices, by affecting firms' cash flow, impact on fundamental values and therefore create a coordination problem for investors. Finally, Berardi (2015) considers a coordination problem with signal extraction and learning, but in that setting both private and public signals are exogenous.

2 The model

We assume the fundamental value of an asset is constant, given by θ. We can think of it as representing a measure of the present discounted value of future dividends. Assuming a constant flow of dividends d over the infinite future

$$\theta = \sum_{i=0}^{\infty} \beta^i d = \frac{d}{1-\beta},$$

where $0 < \beta < 1$ is the discount factor.

Agents maximize their utility, which is a concave function of the value of their portfolio, expressed as difference between the fundamental value of an asset and the price paid for it. They are mean variance maximizers, i.e., try to maximize the mean with a penalty for the variance of their portfolio.

Their problem is thus to choose the number of shares (k) such that

$$\max_k E_t W_t - \frac{\gamma}{2} Var\left(W_t\right)$$

where γ is the coefficient of risk aversion and

$$W_t = k\left(\theta - p_t\right).$$

It follows that optimal demand is

$$k_t^* = \frac{E_t\theta - p_t}{\gamma\sigma_w^2}, \tag{1}$$

where, with θ unknown and $E_t p_t = p_t$,

$$\begin{aligned} \sigma_w^2 &= E_t\left[\left(\theta - p_t\right) - E_t\left(\theta - p_t\right)\right]^2 = \\ &= E_t\left[\left(\theta - E_t\theta\right)\right]^2. \end{aligned}$$

We assume an exogenous and stochastic supply of shares $s = \varepsilon_t$, which follows a normal distribution with zero mean and variance σ_ε^2. We will see that this noise term will prevent prices from being fully revealing. Such assumption has been used before (see, e.g., Branch and Evans 2011), in order to capture variations in the availability of publicly tradable shares (asset float). Departing from their modelling choice, though, we do not assume that supply becomes endogenous at low prices: instead, in order to avoid prices to fall below zero, we impose a non-negativity constraint. It is worth noting that, for the main parameterization reported in the simulation exercise below, such constraint never becomes binding.

Equilibrium condition therefore implies

$$p_t = E_t\theta - \gamma\sigma_w^2\varepsilon_t. \tag{2}$$

Agents need to form an expectation about the unobservable θ and its conditional variance σ_w^2 in order to implement this strategy. With no uncertainty, prices are constant at the fundamental value, i.e., $p_t = \theta$.

2.1 Information structure and equilibrium

We now introduce uncertainty in the model. We assume that there is a continuum of agents of unit mass, indexed by $i \in [0,1]$. Throughout the paper, we will follow the convention that for every time-varying, agent-specific variable z, z_t^i represents a sequence of measurable functions $z_t(i) : [0,1] \to R$, indexed by t, mapping the set of agents at each time t into a real number. Moreover, for a given t, each function $z_t(i)$ is assumed to be continuous and bounded in i. Aggregating over agents, $Z_t = \int_i z_t^i d\Psi_t^z(i)$, where $\Psi_t^z(i)$ is the marginal cumulative distribution for z_t^i.

Agents observe two signals on the fundamental value: one, endogenous and public, from prices (p_t) and one, exogenous and private, from news (x_t^i). We can think of it as a subjective interpretation of news about dividends and other things that affect the long run value of an asset.

The exogenous signal on dividends is represented as

$$x_t^i = \theta + v_{i,t},$$

where $v_{i,t}$ is an i.i.d. random variable, normally distributed with zero mean and variance σ_v^2.

Because signals are normally distributed and conditionally independent, the expected fundamental value conditional on the two signals, denoted by $\tilde{\theta}_t^i \equiv E_t\left[\theta \mid x_t^i, p_t\right]$, is equal to

$$\tilde{\theta}_t^i = \alpha x_t^i + (1-\alpha)\, p_t, \tag{3}$$

where (see Appendix) the optimal value for α (denoted α^*) is given by the solution to

$$\alpha^* = \frac{E_t p_t^2 - E_t p_t x_t^i}{E_t p_t^2 + E_t\left(x_t^i\right)^2 - 2E_t p_t x_t^i} = \frac{\sigma_p^2}{\sigma_p^2 + \sigma_v^2}, \tag{4}$$

where σ_p^2 represents the variance of prices. It is important to note that it is optimal for agents to put some weight on prices, together with the exogenous signal, when forming beliefs about fundamental values (i.e., in general $\alpha^* \neq 1$).

Individual demand is then given by

$$k_{i,t}^* = \frac{\alpha^*\left(x_t^i - p_t\right)}{\gamma \tilde{\sigma}_w^2},$$

where $\tilde{\sigma}_w^2$ is the portfolio variance conditional on x and p (see Appendix).

Aggregate demand is then given by

$$K_t^* = \int \frac{\alpha^* \left(x_t^i - p_t\right)}{\gamma \tilde{\sigma}_w^2} d\Psi_t^x(i) = \frac{\alpha^* \left(\theta - p_t\right)}{\gamma \tilde{\sigma}_w^2},$$

where $\Psi_t^x(i)$ is the standard normal distribution, and prices evolve according to

$$p_t = \theta - \frac{\gamma \tilde{\sigma}_w^2}{\alpha^*} \varepsilon_t. \tag{5}$$

Substituting in for $\tilde{\sigma}_w^2$ (see Appendix for details) we then get the price equation

$$p_t = \theta - \gamma \sigma_v^2 \varepsilon_t, \tag{6}$$

with

$$\sigma_p^2 = \gamma^2 \left(\sigma_v^2\right)^2 \sigma_\varepsilon^2. \tag{7}$$

As already noted by Angeletos and Werning (2006), we can see from (7) that public information improves with private information. Note also that in this setting equilibrium is unique, defined by the optimal value of α: excess volatility will therefore be derived not from switching between multiple equilibria, but from deviations from the unique equilibrium induced by uncertainty and learning.

An important feature of equilibrium prices is that they do not depend explicitly on α^*. This is due to the fact that, in deriving (6), we have used the optimal (from a Bayesian point of view) values of α and $\tilde{\sigma}_w^2$ (which are themselves function of structural parameters and volatilities of exogenous shocks). When instead α is not at its optimal level, it becomes a parameter that affects prices and their volatility: in particular, the more weight is put on prices in the signal extraction problem (lower α), the higher is the volatility of prices. In fact, we can see from (5) that, for arbitrary α,

$$var(p_t) = \left(\frac{\gamma \tilde{\sigma}_w^2}{\alpha}\right)^2 \sigma_\varepsilon^2$$

and therefore

$$\frac{\delta var(p_t)}{\delta \alpha} < 0.$$

Proposition 1 *In an economy where prices act as endogenous signals for fundamental values, the higher is the weight put on prices in the signal extraction problem, the higher is the volatility of prices.*

Note that if the exogenous signal was public information and everyone was thus observing the same signal x_t, (4) would become

$$\alpha^* = \frac{E_t p_t^2 - E_t p_t x_t}{E_t p_t^2 + E_t x_t^2 - 2E_t p_t x_t} = \frac{\sigma_p^2 - \sigma_v^2}{\sigma_p^2 + \sigma_v^2 - 2\sigma_v^2} = 1 : \tag{8}$$

the noise in the exogenous public signal would be transferred to prices, and thus prices would be completely useless as a signal for the fundamental, since they would encompass both the noise from the exogenous signal and the noise from supply: optimal α would thus be one. In other words, in order for prices to have any informational content above and beyond that provided by the idiosyncratic signal, it must be that the aggregation process that generates prices averages out some of the noise in the information about fundamentals.

Instead, with x_t^i a private signal, we have

$$\alpha^* = \frac{\sigma_p^2}{\sigma_p^2 + \sigma_x^2} = \frac{\gamma^2 \left(\sigma_v^2\right)^2 \sigma_\varepsilon^2}{\sigma_v^2 + \gamma^2 \left(\sigma_v^2\right)^2 \sigma_\varepsilon^2} = \frac{\gamma^2 \sigma_v^2 \sigma_\varepsilon^2}{1 + \gamma^2 \sigma_v^2 \sigma_\varepsilon^2} \tag{9}$$

and therefore

$$\begin{aligned} \lim_{\sigma_\varepsilon^2 \to 0} \alpha^* &= 0; \lim_{\sigma_\varepsilon^2 \to \infty} \alpha^* = 1 \\ \lim_{\sigma_v^2 \to 0} \alpha^* &= 0; \lim_{\sigma_v^2 \to \infty} \alpha^* = 1. \end{aligned}$$

If the variance of supply goes to zero, then prices are fully revealing and only prices are used to infer fundamental values. If instead it goes to infinity, then only the exogenous signal is used as prices loose all information content regarding fundamental values.

Note also that if the variance of the idiosyncratic noise in the exogenous signal goes to zero, α^* goes to zero, as can be easily seen from (9): no weight is put on the private signal and only prices are used. This might seem at first counter-intuitive, and looking at (4) one might actually mistakenly think that α^* goes instead to one when $\sigma_v^2 \to 0$. The reason why this does not happen is that as $\sigma_v^2 \to 0$, the variance of prices goes to zero faster than that of the private signal:[2] this is due to the fact that demand (in absolute value) increases as the precision of the information on the fundamental increases and, as $\sigma_v^2 \to 0$, it goes to (plus or minus) infinity. To avoid this result, we could introduce an additional, aggregate, noise in the exogenous signal so that when the idiosyncratic noise disappears, there would still be uncertainty about the fundamental value of the asset. We avoid this complication here, since we will stay away from the limiting case.

[2]The variance of prices is in fact quadratic in σ_v^2: see (7).

While it is optimal for agents to use prices, together with the exogenous signal, to form beliefs about fundamental values, it is perhaps too much to assume that they know the exact value of the optimal weight to put on each signal. It is for this reason that we now turn to adaptive learning.

3 Learning about variances

One essential element of any learning analysis is to decide what aspects of their decision problem agents should be learning about. Put it another way, how much a priori knowledge should they be granted? How much information should they have about the environment they live in? We take the view that agents know their own preferences, and know how to solve relatively simple problems of signal extraction. This means that we assume they can work out the way the optimal weight depends on the relative precision of the two signals. What they don't know a priori, and need to learn about, are the statistical properties of the environment they live in, and thus of the relevant variables involved in their decision problem. Both prices and news contain information about the fundamental: news in addition have idiosyncratic volatility due to noise, and prices have added volatility due to supply variation. The aim of agents is to disentangle these sources of volatility to extract information about the fundamental.

In particular, agents will need to learn about mean and variances of the two signals. While prices are publicly observable, the exogenous signal is instead idiosyncratic: it follows that, under learning, agents will hold heterogeneous beliefs about the statistical properties of the exogenous signal while, given homogeneous initial beliefs and updating algorithms, they will hold instead homogeneous beliefs about prices.

We denote by $s^i_{x.t}$ agents' heterogeneous estimates about $(x^i_t)^2$ (i.e., $s^i_{x.t} \equiv E^i_t\,(x^i_t)^2$), by $s_{p,t}$ agents' homogeneous estimates about $(p_t)^2$ (i.e., $s_{p.t} \equiv E_t\,(p_t)^2$), by $s^i_{px,t}$ agents' heterogeneous estimates about $p_t x^i_t$ (i.e., $s^i_{px,t} \equiv E^i_t p_t x^i_t$) and by $s^i_{w,t}$ agents' heterogeneous estimates about $\tilde{\sigma}^2_w$ (i.e., $s^i_{w,t} \equiv E^i_t\left[(\theta - E^i_t\theta)\right]^2$). We also define $\tilde{x}^i_t$ as agent i's time t estimate of the mean of x, $\tilde{p}_t$ as agents' time t (homogeneous) estimate of the mean of p and $cov^i_t(x,p)$ as agent i's time t estimate of the covariance between x and p.

An agent i at time t is therefore defined by the signal she receives, x^i_t, and her idiosyncratic beliefs system $[s^i_{x.t}, s^i_{px.t}, s^i_{w,t}, \tilde{x}^i_t, cov^i_t(x,p)]$. The beliefs system is completed by the two homogeneous elements $s_{p.t}$ and $\tilde{p}_t$. The state space Ω is therefore defined as $\Omega = \mathbb{R}^8$, with an element $\omega = \left(x^i_t, s^i_{x.t}, s^i_{px.t}, s^i_{w,t}, \tilde{x}^i_t, cov^i_t(x,p), s_{p.t}, \tilde{p}_t\right)$. There is a joint probability distribution

$\Psi_t(i)$ that defines the cumulative distribution of agents at each point in time, and we define as $\Psi_t^z(i)$ the marginal distribution for the generic variable z.

Individual demands are

$$k_t^i = \frac{\alpha_t^i \left(x_t^i - p_t\right)}{\gamma s_{w,t}^i}, \tag{10}$$

where

$$\alpha_t^i = \frac{E_t p_t^2 - E_t p_t x_t^i}{E_t p_t^2 + E_t x_t^i - 2E_t p_t x_t^i} = \frac{s_{p,t} - s_{px,t}^i}{s_{p,t} + s_{x,t}^i - 2s_{px,t}^i} \tag{11}$$

and[3]

$$\begin{aligned} s_{w,t}^i &= var_t^i(\alpha_t^i x_t^i + \left(1-\alpha_t^i\right) p_t) & (12)\\ &= \left(\alpha_t^i\right)^2 \left(s_{x,t}^i - \left(\tilde{x}_t^i\right)^2\right) + \left(1-\alpha_t^i\right)^2 \left(s_{p,t} - (\tilde{p}_t)^2\right) + 2\alpha_t^i \left(1-\alpha_t^i\right) cov_t^i(x,p). & (13) \end{aligned}$$

We then have that aggregate demand is given by

$$K_t = \int \frac{\alpha_t^i \left(x_t^i - p_t\right)}{\gamma s_{w,t}^i} d\Psi_t(i). \tag{14}$$

From the market clearing condition $K_t = \varepsilon_t$, it follows that prices evolve according to

$$p_t = \frac{\int \frac{\alpha_t^i x_t^i}{\gamma s_{w,t}^i} d\Psi_t(i) - \varepsilon_t}{\int \frac{\alpha_t^i}{\gamma s_{w,t}^i} d\Psi_t(i)}, \tag{15}$$

or, since x_t^i does not covariate with the other state variables (see below the updating rule for $s_{x,t}^i$, $s_{xp,t}^i$, $s_{p,t}$ and $s_{w,t}^i$, which all depend on time $t-1$ information) and $\int x_t^i d\Psi_t^x(i) = \theta$ (where again $\Psi_t^x(i)$ is the standard normal distribution)

$$p_t = \theta - \frac{\varepsilon_t}{\int \frac{\alpha_t^i}{\gamma s_{w,t}^i} d\Psi_t(i)}. \tag{16}$$

We now define how beliefs are updated, according to recursive algorithms. Agents learn about raw second moments according to

$$\begin{aligned} s_{x,t+1}^i &= s_{x,t}^i + g_t\left[\left(x_t^i\right)^2 - s_{x,t}^i\right] & (17)\\ s_{p,t+1} &= s_{p,t} + g_t\left[(p_t)^2 - s_{p,t}\right] & (18)\\ s_{px,t+1}^i &= s_{px,t} + g_t\left[p_t x_t^i - s_{px,t}^i\right] & (19) \end{aligned}$$

[3] Here we are using the identity $Var(x) = E[x^2] - [Ex]^2$.

and about the central moment $s^i_{w,t}$ using

$$s^i_{w,t} = \left(\alpha^i_t\right)^2 \left(s^i_{x,t} - \left(\tilde{x}^i_t\right)^2\right) + \left(1-\alpha^i_t\right)^2 \left(s_{p,t} - \left(\tilde{p}^i_t\right)^2\right) + 2\alpha^i_t \left(1-\alpha^i_t\right) cov^i_t(x,p). \tag{20}$$

Moreover, means and covariance of the two signals are updated according to

$$\begin{aligned} \tilde{x}^i_{t+1} &= \tilde{x}^i_t + g_t\left[x^i_t - \tilde{x}^i_t\right] && (21)\\ \tilde{p}_{t+1} &= \tilde{p}_t + g_t\left[p_t - \tilde{p}_t\right] && (22)\\ cov^i_t(x,p) &= s^i_{px,t} - \tilde{x}^i_t \tilde{p}_t. && (23) \end{aligned}$$

The learning gain g_t, which controls how new information gets factored into the estimates, is set equal to $1/t$ for the decreasing gain case, and to a small constant (more on this later) in case of constant gain.

Note that p_t depends on $s_{p,t}$, as

$$p_t = \theta - \frac{\varepsilon_t}{\int \frac{\alpha^i_t}{\gamma s^i_{w,t}} d\Psi_t(i)} = \theta - \frac{\varepsilon_t}{\int \frac{s_{p,t} - s^i_{px,t}}{s_{p,t} + s^i_{x,t} - 2s^i_{px,t}} \frac{1}{\gamma s^i_{w,t}} d\Psi_t(i)} \tag{24}$$

so (18) and (19) become

$$\begin{aligned} s_{p,t+1} &= s_{p,t} + g_t\left[\left(\theta - \frac{\gamma\varepsilon_t}{\int \frac{s_{p,t} - s^i_{px,t}}{s_{p,t} + s^i_{x,t} - 2s^i_{px,t}} \frac{1}{s^i_{w,t}} d\Psi_t(i)}\right)^2 - s_{p,t}\right] && (25)\\ s^i_{px,t+1} &= s_{px,t} + g_t\left[\left(\left(\theta - \frac{\gamma\varepsilon_t}{\int \frac{s_{p,t} - s^i_{px,t}}{s_{p,t} + s^i_{x,t} - 2s^i_{px,t}} \frac{1}{s^i_{w,t}} d\Psi_t(i)}\right) x^i_t\right) - s^i_{px,t}\right]. && (26) \end{aligned}$$

The system (17), (20), (21)–(23), (24)-(26) is a non-linear stochastic recursive algorithm, whose long-run properties can be analyzed using stochastic approximation techniques.

3.1 Convergence results

We analyze convergence results for the decreasing gain case $g_t = 1/t$, which reduces the learning algorithms in instances of recursive least squares estimates. Since agents are estimating only constants (means, variances, covariances), the algorithms are greatly simplified: care has still to be taken, though, when such values are endogenous, because of the feedback effect from estimates to actual values.

Convergence of (17) is easy to establish, since x_t^i is exogenous and agents are effectively estimating the value of an exogenous constant (the variance of the distribution). Using stochastic approximation techniques (see Evans and Honkapohja (2001) for details), we can represent the evolution, in notional time τ, of estimated variables through an ordinary differential equation (ODE) derived as follows

$$\frac{ds_x^i}{d\tau} = \lim_{t\to\infty} E\left(\left(x_t^i\right)^2 - s_{x,t}^i\right) = \theta^2 + \sigma_v^2 - s_x^i, \tag{27}$$

and therefore in equilibrium $s_x^i = \theta^2 + \sigma_v^2$, $\forall i$. Moreover, this fixed point is stable since the derivative of the ordinary differential equation represented by (27) is equal to -1.

Similarly for (21), which gives rise to the ODE

$$\frac{d\tilde{x}^i}{d\tau} = \lim_{t\to\infty} E\left(x_t^i - \tilde{x}\right) = \theta - \tilde{x}^i \tag{28}$$

with stable fixed point $\tilde{x}^i = \theta$, $\forall i$. Moreover, since ε_t is a zero mean i.i.d. process independent from all time-t beliefs, we have, from (18),

$$\frac{d\tilde{p}}{d\tau} = \lim_{t\to\infty} E\left(\theta - \frac{\varepsilon_t}{\int \frac{s_{p,t}-s_{px,t}^i}{s_{p,t}+s_{x,t}^i-2s_{px,t}^i}\frac{1}{\gamma s_{w,t}^i} d\Psi_t\left(i\right)} - \tilde{p}_t\right) = \theta - \tilde{p} \tag{29}$$

and therefore the stable fixed point is $\tilde{p} = \theta$.

To establish convergence of (25) and (26), we rewrite (20) as

$$s_{w,t}^i = \left(\frac{s_{p,t} - s_{px,t}^i}{s_{p,t} + s_{x,t}^i - 2s_{px,t}^i}\right)^2 \sigma_v^2 + \left(1 - \frac{s_{p,t} - s_{px,t}^i}{s_{p,t} + s_{x,t}^i - 2s_{px,t}^i}\right)^2 \left(s_{p,t} - \theta^2\right) + 2\alpha_t^i\left(1-\alpha_t^i\right) cov_t^i(x,p) \tag{30}$$

so that our system is given by (25) and (26) plus (24) and (30). Defining $\phi_t = \begin{bmatrix} s_{p,t} \\ s_{px,t}^i \end{bmatrix}$, we can write this system as

$$\phi_{t+1} = \phi_t + t^{-1}\mathcal{H}\left(t, s_{p,t}, s_{px,t}^i, s_{x,t}^i, s_{w,t}^i, cov_t^i(x,p), \varepsilon_t\right). \tag{31}$$

and again, using stochastic approximation theory, we have that asymptotically the dynamics of (31) are governed by the ODE

$$\frac{d\phi}{d\tau} = h(\phi)$$

where

$$h(\phi) = \lim_{t\to\infty} E\mathcal{H}\left(t, s_{p,t}, s^i_{px,t}, s^i_{x,t}, s^i_{w,t}, cov^i_t(x,p), \varepsilon_t\right).$$

We can then decompose $h(\phi)$ into

$$h_p = \lim_{t\to\infty} E\left(\left(\theta - \frac{\gamma\varepsilon_t}{\int \frac{s_{p,t}-s^i_{px,t}}{s_{p,t}+s^i_{x,t}-2s^i_{px,t}}\frac{1}{s^i_{w,t}}d\Psi_t(i)}\right)^2 - s_{p,t}\right) \tag{32}$$

$$h_{px} = \lim_{t\to\infty} E\left(\left(\theta - \frac{\gamma\varepsilon_t}{\int \frac{s_{p,t}-s^i_{px,t}}{s_{p,t}+s^i_{x,t}-2s^i_{px,t}}\frac{1}{s^i_{w,t}}d\Psi_t(i)}\right) x^i_t - s^i_{px,t}\right). \tag{33}$$

Since x^i_t does not covariate with the other state variables, (33) reduces to

$$h_{px} = \theta^2 - s^i_{px} \tag{34}$$

and therefore $s^i_{px} = \theta^2\ \forall i$: this fixed point is stable since the derivative of ODE (34) is equal to -1.

As for (23), given that x^i_t is exogenous and independent from prices, results about s^i_{px}, $\tilde{x}^i$ and $\tilde{p}$ imply that $s^i_{px} = \tilde{x}^i\tilde{p} = \theta^2$ and therefore in equilibrium $cov^i_t(x,p) = 0$.

Using $s^i_{px} = \theta^2$ and $s^i_x = \theta^2 + \sigma^2_v$, we have that (32) reduces to

$$h_p = \lim_{t\to\infty} E\left(\left(\theta - \frac{\gamma\varepsilon_t}{\int \frac{s_{p,t}-\theta^2}{s_{p,t}+\theta^2+\sigma^2_v-2\theta^2}\frac{1}{s^i_{w,t}}d\Psi_t(i)}\right)^2 - s_{p,t}\right)$$

and substituting in

$$s^i_{w,t} = \left(\frac{s_{p,t}-\theta^2}{s_{p,t}-\theta^2+\sigma^2_v}\right)^2\sigma^2_v + \left(1 - \frac{s_{p,t}-\theta^2}{s_{p,t}-\theta^2+\sigma^2_v}\right)^2\left(s_{p,t}-\theta^2\right)$$

we have

$$h_p = \lim_{t\to\infty} E\left(\left(\theta - \gamma\sigma^2_v\varepsilon_t\right)^2 - s_{p,t}\right)$$

and therefore

$$h_p = \theta^2 + \gamma^2\left(\sigma^2_v\right)^2\sigma^2_\varepsilon - s_p. \tag{35}$$

It follows that the equilibrium point is $s_p = \theta^2 + \gamma^2\left(\sigma^2_v\right)^2\sigma^2_\varepsilon$ which leads to the Bayesian equilibrium found above. Moreover, the equilibrium is stable under learning (E-stable) since the derivative of ODE (35) is equal to -1. The key result here is that agents can learn about the variance of prices, which is endogenous and changes with agents' beliefs. Such result might seem particularly surprising, since the value of the derivative of the ODE implies that convergence does not depend on the specific parameterization assumed: the reason is that, although prices are endogenous, opposing forces between α_t^i and $s_{w,t}^i$ in the demand schedule ensure stability of prices. In fact, at equilibrium

$$\frac{\alpha_t^i}{s_{w,t}^i} = \frac{\frac{s_{p,t}-\theta^2}{s_{p,t}-\theta^2+\sigma_v^2}}{\left(\frac{s_{p,t}-\theta^2}{s_{p,t}-\theta^2+\sigma_v^2}\right)^2 \sigma_v^2 + \left(\frac{\sigma_v^2}{s_{p,t}-\theta^2+\sigma_v^2}\right)^2 \left(s_{p,t}-\theta^2\right)} = \frac{1}{\sigma_v^2}.$$

Finally, convergence in all the learning algorithms ensures that α_t^i also converges to its equilibrium value, i.e., from (4),

$$\lim_{t\to\infty} \alpha_t^i = \lim_{t\to\infty} \frac{s_{p,t} - s_{px,t}^i}{s_{p,t} + s_{x,t}^i - 2s_{px,t}^i} = \frac{\gamma^2\sigma_v^2\sigma_\varepsilon^2}{1+\gamma^2\sigma_v^2\sigma_\varepsilon^2}, \ \forall i.$$

Proposition 2 *Consider the system of ODEs (27),(28), (29), (34), (35), defining learning in notional time. All ODEs are stable and agents learn the means, variances and covariances of the relevant exogenous and endogenous variables: it follows that $\alpha_t^i \to \alpha^*$ and the system converges to the Bayesian equilibrium with prices given by (6).*

3.2 Constant gain

Though the decreasing gain case allows us to establish convergence of the model under learning to the Bayesian equilibrium, from an applied perspective it seems more relevant the case with constant gain. It might be argued, in fact, that agents trying to learn the fundamental value of an asset would use a constant gain algorithm, to allow for time variation in the fundamentals. A growing literature in applied macroeconomics has used constant gain learning to explain a range of features, from the rise and fall of U.S. inflation in the 70s and 80s (in particular, the seminal work of Sargent (1999) and Sargent at el. (2006)) to the causes of business cycles (e.g., Milani (2011) and Eusepi and Preston (2011)).

Though for simplicity we model the underlying fundamental as a constant, in reality the fundamental value of an asset is likely to change over time as conditions and profitability of a firm change, and with them the expected future cash flow. We choose to model the fundamental as constant because it simplifies the analysis and it allows us to clearly identify the effects of the time variation in α on prices and volatilities. A time-varying fundamental would add to the volatility of the model, thus making it harder to disentangle the effect that comes from learning. Note, however, that such feature would likely enhance even further our conclusions on the effects of learning on prices and volatilities, as agents would then need to learn about a time-varying variable.

It is well known that, with constant gain, beliefs can not converge point-wise, but they can still converge in distribution. Provided some regularity conditions hold (See Evans and Honkapohja (2001), pp. 162-166), parameter estimates converge in distribution to their equilibrium values, and the variance of the distribution is proportional to the gain.

We resort here to simulations to investigate the transient dynamics of the system with constant gain.

3.2.1 Simulations

Simulations of the model under real-time learning are run with a finite number of agents ($n = 1000$). Results reported are obtained with $\gamma = .9$, $\sigma_v^2 = 1$, $\sigma_\varepsilon^2 = 1$, $g = .005$, but we will discuss the impact of different parameter values on results later on.

We first present the evolution of prices and of average weight in Figure 1.[4] We can see that, because of variations in the relative weight on signals, prices show excess volatility with respect to their equilibrium values. In particular, periods of high volatility coincide with periods of low α: as agents put more weight on prices in an attempt to predict fundamentals, herding (to be defined precisely below) emerges and leads prices away from equilibrium. While α deviates away from its equilibrium value for considerable time, both above and below it, because equilibrium is unique and stable under learning deviations sooner or later tend to be corrected and the system moves back towards equilibrium. These dynamics are driven by a mix of herding and contrarian behavior, as explained in the next Section.

In Figure 2 we then show the evolution of average beliefs about means and variances of the relevant variables. We can see that there is quite some variability in these values, driven by the constant gain learning algorithm. The combination of s_p, s_{px} and s_x contribute to determine the evolution of α, while $\tilde{x}$, $\tilde{p}$ and $cov(x,p)$ contribute to determine demand (and thus prices) through s_w.In Figure 3 we then report returns from prices generated in a Bayesian equilibrium and returns under learning. It can be seen that under learning returns show clear patterns of volatility clustering that are missing in a Bayesian equilibrium: the learning-induced variations in α account for this feature. In particular, periods of high

[4] While we report here only the average value of α, simulations show that there is also quite some dispersion across agents in its value, driven by the idiosyncratic shocks in x_t^i.

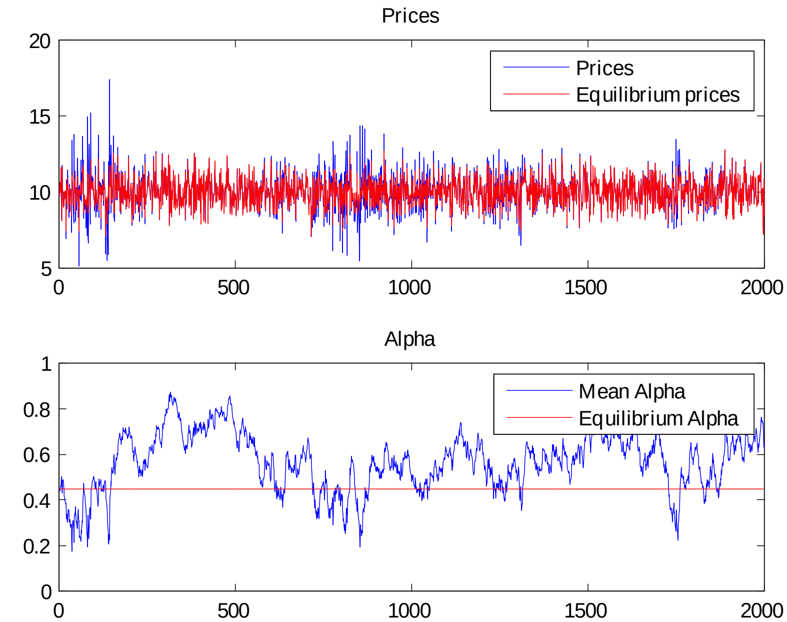

Figure 1: Prices and α.

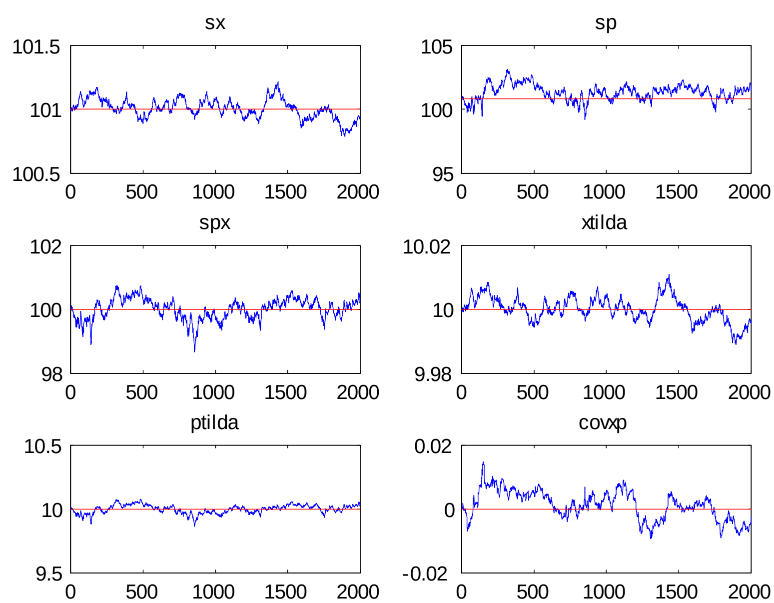

Figure 2: Average beliefs

volatility in returns coincide with periods of low α.

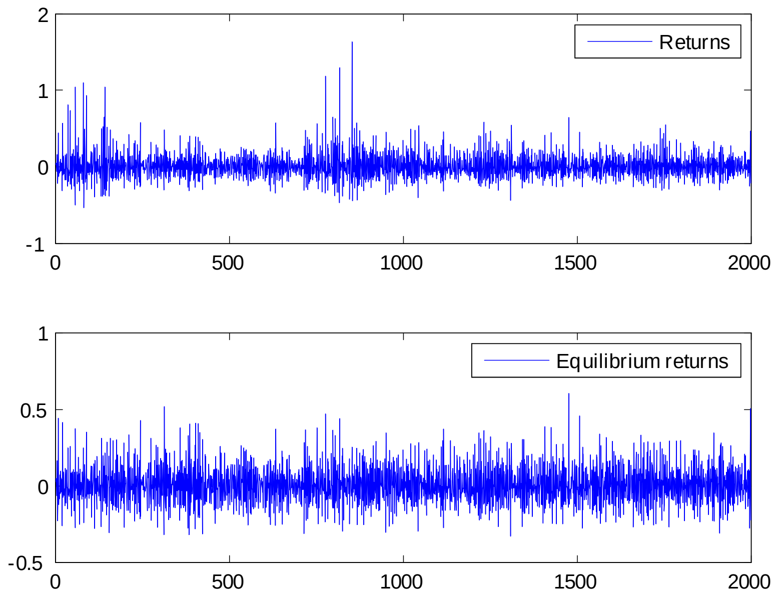

Figure 3: Returns

In terms of our choice of parameter values, we run some sensitivity analysis to verify the impact of different values on results. It emerges that, as expected, the volatility of α and of prices is proportional to g: a smaller gain leads to limited variation in α and therefore reduced volatility in prices; clustering in returns volatility also reduces. As a comparison, for a gain equal to 0.001, average α fluctuates roughly between 0.3 and 0.6, a much smaller range than that shown in Figure 1. In terms of γ, this parameter must be large enough for changes in α to be reflected in prices, but not too high to become destabilizing: a good range found to work is roughly between 0.3 and 2.[5] Finally, the variances σ_v^2 and σ_ε^2 impact in rather different ways. The variance of the supply shock must not be too high to avoid prices becoming negative at times, when large shocks realize. The variance of the noise in the private information instead affects the estimates of s_w^i, and through this channel demand and prices: higher σ_v^2 leads, other things equal, to larger s_w^i, and thus smaller demand and lower price variability.

[5]In particular, for a given σ_ε^2, increasing γ is more likely to lead to occasional negative prices, when large supply shocks hit the system.

4 Discussion

In our model fundamental values are constant, there is no persistence in any of the state variables and shocks are i.i.d., stationary and homoscedastic. In addition, persistence is not built in ad hoc, for example by imposing past values of the state variables in agents's perceived laws of motion. As a result, and as already noted before, our model is a model of volatilities and not of trends.

In the Bayesian equilibrium, prices are normally distributed with mean θ and variance $\gamma^2 \left(\sigma_v^2\right)^2 \sigma_\varepsilon^2$. Uncertainty about the signal extraction problem generates additional volatility in prices, which is greater the more weight is put on prices in the formation of beliefs about fundamental values. In fact, the correlation coefficient between the average of α_t^i and $|p_t - p_t^{eq}|$ is -0.5718: when α decreases, prices become more volatile, since more weight is put on the endogenous signal (prices). A herd-like behavior emerges, where agents put more weight on information that depends on other agents' actions rather than on their own private information.

Results are not driven by changes in perception of risk, like in Branch and Evans (2011): in fact, even if we fix $s_{w,t}^i$ to its equilibrium value, $s_{w,t}^i = \tilde{\sigma}_w^2 = \alpha\sigma_v^2$, volatility in α still emerges and impacts on prices. The key element in driving results here is instead the change in weight on the two signals, exogenous and endogenous, due to changes in beliefs about variances and covariances.

From individual demand under uncertainty (10), it follows

$$\frac{\delta k_t^i}{\delta p_t} = -\frac{\alpha_t^i}{\gamma s_{w,t}^i}.$$

We can see that α contributes to determine the downward sloping curve of individual demand, that is, the extend to which higher prices lead to lower demand. A smaller α means that such effect is dampened, since higher prices are taken to imply higher fundamental values.

Looking at the dynamics of α, we can identify periods of herding and periods of anti-herding or contrarian behavior. A low α leads to herding in agents' behavior: agents disregard part of their own information and follow others in their demand. High α, on the contrary, implies contrarian behavior: agents put more weight on their private information than what would be optimal in deriving their demand schedule.

For the purpose of this paper, we define herding as periods where $\alpha_t^i < \alpha^*$. The effects on demand depend on the sign of the numerator of equation (10) and are as follows:

$$\text{if } \alpha_t^i < \alpha^* \text{ and } x_t^i > p_t \text{ then } k_t^i < k_{i,t}^*;$$
$$\text{if } \alpha_t^i < \alpha^* \text{ and } x_t^i < p_t \text{ then } k_t^i > k_{i,t}^*.$$

In the first case, agents herd into low demand because they put too much weight on a low price signal (compared to the private signal); in the second case instead agents herd into high demand because they put too much weight on a high price signal.

We then define contrarian behavior as periods where $\alpha_t^i > \alpha^*$, with the following impact on demand:

$$\text{if } \alpha_t^i > \alpha^* \text{ and } x_t^i < p_t \text{ then } k_t^i < k_{i,t}^*;$$
$$\text{if } \alpha_t^i > \alpha^* \text{ and } x_t^i > p_t \text{ then } k_t^i > k_{i,t}^*.$$

In the first case, agents' demand is too low because they put too much weight on their own, low, private signal; in the second case, agents's demand is too high because they put too much weight on their own, high, signal.

In other words, there is herding when agents demand less (more) than what would be optimal because they put too much weight on a low (high) public price signal, which leads them to revise downwards (upwards) their beliefs about fundamentals. Similarly, there is contrarian behavior when agents demand less (more) than what would be optimal because they put too much weight on a low (high) private signal, which leads them to revise downwards (upwards) their beliefs about fundamentals.

We then define a herding index, h, as the fraction of time there is aggregate herding in our economy (i.e., when $mean(\alpha_t^i) < \alpha^*$): that is, when, on average, agents put an excessive weight on each other's opinion (as conveyed by prices) when assessing fundamentals. Simulations show that, on average, herding happens around 15% to 20% of the time.

Similarly, we define a contrarian index, c, as the fraction of time there is aggregate contrarian behavior in our economy (i.e., when $mean(\alpha_t^i) > \alpha^*$): that is, when, on average, agents put too little weight on each other's opinion (as conveyed by prices) when assessing fundamentals. Simulations show that around 80% to 85% of the time there is contrarian behavior in our setup.

Note that herding is destabilizing for prices, as it induces lower (higher) demand when prices are low (high), thus reinforcing price deviations from equilibrium. Contrarian behavior instead is stabilizing, as it induces lower (higher) demand when prices are high (low), thus correcting price deviations from equilibrium. It is the combination of these two phenomena that lead to excess volatility in prices, a finding already highlighted by Park and Sabourian (2011) in a different setting.

5 Conclusions

We have proposed a model where uncertainty and learning about fundamental values can lead to excess volatility in prices and volatility clustering in returns. The key insight is that when agents use prices as an endogenous signal for fundamentals, herding and contrarian behavior in the use of information can emerge that lead to prices and returns to deviate from equilibrium. This work thus presents a new mechanism able to generate volatility patterns in prices and returns similar to those observed in real markets: instead of focusing on speculative forces, we highlight the role played by uncertainty and learning about assets' fundamental values.

6 Appendix

6.1 Derivation of α^* with private signal

The optimal weight on the two signals, α^*, is obtained by solving the problem

$$\alpha^* = \arg\min_{\alpha} E_t \left(\theta - \tilde{\theta}_t^i\right)^2 \tag{36}$$

with

$$\tilde{\theta}_t^i = \alpha x_t^i + (1-\alpha)\, p_t. \tag{37}$$

Minimizing (36) subject to (37) leads to the FOC

$$E_t \left(\theta - \tilde{\theta}_t^i\right) \left(p_t - x_t^i\right) = 0,$$

whose solution implies

$$\alpha^* = \frac{E_t p_t^2 - E_t p_t x_t^i}{E_t p_t^2 + E_t \left(x_t^i\right)^2 - 2 E_t p_t x_t^i}. \tag{38}$$

Given that prices and exogenous signal do not covariate (since the noise in the signal is averaged out by aggregation), this reduces to

$$\alpha^* = \frac{\sigma_p^2}{\sigma_p^2 + \sigma_v^2}. \tag{39}$$

Individual demand is then given by

$$k_{i,t}^* = \frac{\alpha^* \left(x_t^i - p_t\right)}{\gamma \tilde{\sigma}_w^2}, \tag{40}$$

where $\tilde{\sigma}_w^2$ is the portfolio variance conditional on x and p, which is equal to the conditional variance of the fundamental and given by

$$\begin{aligned} \tilde{\sigma}_w^2 &= E_t \left[\left(\theta - \tilde{\theta}_t^i \right)^2 \mid x_t^i, p_t \right] \\ &= \alpha^2 \sigma_v^2 + (1-\alpha)^2 \sigma_p^2. \end{aligned} \tag{41}$$

Aggregating individual demand (40) and equating it with supply we obtain the price equation

$$p_t = \theta - \frac{\gamma \tilde{\sigma}_w^2}{\alpha^*} \varepsilon_t \tag{42}$$

which implies

$$\sigma_p^2 = \left(\frac{\gamma \tilde{\sigma}_w^2}{\alpha^*} \right)^2 \sigma_\varepsilon^2. \tag{43}$$

It follows from (39) and (41) that

$$\tilde{\sigma}_w^2 = \left(\frac{\sigma_p^2}{\sigma_p^2 + \sigma_v^2} \right)^2 \sigma_v^2 + \left(\frac{\sigma_v^2}{\sigma_p^2 + \sigma_v^2} \right)^2 \sigma_p^2,$$

which leads to

$$\tilde{\sigma}_w^2 = \frac{\left(\sigma_p^2\right)^2 \sigma_v^2 + \sigma_p^2 \left(\sigma_v^2\right)^2}{\left(\sigma_v^2 + \sigma_p^2\right)^2} = \frac{\left(\sigma_v^2 + \sigma_p^2\right) \sigma_p^2 \sigma_v^2}{\left(\sigma_v^2 + \sigma_p^2\right)^2} = \frac{\sigma_p^2 \sigma_v^2}{\sigma_v^2 + \sigma_p^2} = \alpha^* \sigma_v^2. \tag{44}$$

Substituting (44) into (42) we obtain

$$p_t = \theta_t - \gamma \sigma_v^2 \varepsilon_t$$

and therefore

$$\sigma_p^2 = \gamma^2 \left(\sigma_v^2\right)^2 \sigma_\varepsilon^2.$$

References

[1] Adam, K., Marcet, A., Nicolini, J.P., 2008. Stock market volatility and learning. European Central Bank Working Paper Series n. 862.

[2] Angeletos, G.-M., Werning, I., 2006. Crises and Prices: Information Aggregation, Multiplicity, and Volatility. American Economic Review 96, 1720-1736.

[3] Berardi, M., 2015. Learning and coordination with dispersed information. Journal of Economic Dynamics and Control 58, 19-33.

[4] Bernanke, B., Gertler, M., 1999. Monetary Policy and Asset Price Volatility. Presented at New Challenges for Monetary Policy, a symposium sponsored by the Federal Reserve Bank of Kansas City, Jackson Hole, Wyoming, August 26-28, 1999. Available at https://www.kansascityfed.org/publicat/sympos/1999/S99gert.pdf.

[5] Branch, W.A., Evans, G.W., 2010. Asset return dynamics and learning. Review of Financial Studies 23, 1651-1680.

[6] Branch, W.A., Evans, G.W., 2011. Learning about risk and return: a simple model of bubbles and crashes. American Economic Journal: Macroeconomics 3, 159-191.

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

[8] Cont, R., 2001. Empirical properties of asset returns: stylized facts and statistical issues. Quantitative Finance 1, 223-236.

[9] De Long, J.B., Shleifer, A., Summers, L.H., Waldmann, J., 1990. Noise trader risk in financial markets. Journal of Political Economy 98, 703-38.

[10] Eusepi, S., Preston, B., 2011. Expectations, learning, and business cycle fluctuations. American Economic Review 101, 2844–2872.

[11] Hommes, C., Zhu, M., 2011. Learning under misspecification: a behavioral explanation of excess volatility in stock prices and persistence in inflation. CeNDEF Working Papers 11-04.

[12] Lux, T., 2009. Stochastic behavioral asset pricing models and the stylized facts. In T. Hens and K. Schenk-Hoppé eds., Handbook of Financial Markets: Dynamics and Evolution. Amsterdam, North-Holland, 161-215.

[13] Milani, F., 2011. Expectation shocks and learning as drivers of the business cycle. The Economic Journal 121, 379–401.

[14] Morris, S., Shin, H.S., 2006. Endogenous public signals and coordination. Mimeo.

[15] Ozdenoren, E., Yuan, K., 2007. Feedback effects and asset prices. Journal of Finance 63, 1939-75.

[16] Park, A., Sabourian, H., 2011. Herding and contrarian behavior in financial markets. Econometrica 79, 973–1026.

[17] Sargent, T.J., 1999. The Conquest of American Inflation. Princeton, NJ: Princeton University Press.

[18] Sargent, T.J., Williams, N., Zha, T., 2006. Shocks and government beliefs: the rise and fall of American inflation. American Economic Review 96, 1193-1224.