Discussion Paper Series

Herding through learning in an asset pricing model

By

Michele Berardi

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

August 2016
Number 223

Herding through learning in an asset pricing model*

Michele Berardi
The University of Manchester

August 25, 2016

Abstract

In this paper we show how uncertainty and learning can lead to a disconnection between fundamental values and prices in a simple asset pricing model. Agents use prices, besides an idiosyncratic exogenous signal, to infer fundamental values: as agents accumulate information, they put increasing weight on the public signal and in the limit they ignore completely their private information. The Bayesian equilibrium implies that agents end up relying only on prices in their signal extraction problem, an outcome that reminds the rational herding result in sequential decision making. We also consider two extensions that should mitigate this effect, namely constant gain adaptive learning and Bayesian learning with an explicit probability of change in the fundamental. In both cases the problem persists, though somewhat mitigated. As a by-product, we also establish a connection between the constant gain parameter in adaptive learning and the subjective probability of exogenous changes in Bayesian learning.

Key words: uncertainty, signal extraction, adaptive learning, asset prices.

JEL classification: D83, D84, G12.

* I would like to thank participants at the Royal Economic Society 2016 Conference and at the Computations in Economics and Finance 2016 Conference, as well as at various seminars, for useful comments, suggestions and discussions. I claim ownership of all remaining errors.

Herding through learning in an asset pricing model

1 Introduction

The fundamental value of an asset summarizes the future stream of cash flows that the asset entitles to. By definition, such value is never observable, as it is always determined by events in the future. The price of an asset represents what market participants are willing to pay for it. Such value is readily observable. This paper analyses the link between these two elements in a context of imperfect information. To investigate this issue, we merge two lines of literature: one on signal extraction and Bayesian learning, the other on adaptive learning. From the first, we take the basic building blocks to model prices as an endogenous signal for agents, a signal 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: in particular, moments of the relevant distributions need to be estimated from observed data.

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 because prices summarize the view of other agents about the fundamental value. To better isolate the link between fundamentals and prices, we assume agents are only concerned about the fundamental value of their portfolio, relative to its price, and so they do not try to profit from exploiting short term capital gains. In other words, our economy is populated only by fundamentalists and there are no speculators or noise traders 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 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 would seem to be 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 using observables such as dividends and prices as indirect information.

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 dividends and other information affecting the value and profitability of a firm. Besides this exogenous signal, agents also use prices to infer fundamental values, as they summarize the view of other market participants and thus convey important information about the fundamental value.

Bayesian theory provides the optimal weight on the two signals: in the first part of the paper we derive such value in a static setting and discuss its implications for asset prices. We then extend the model to a dynamic framework where agents repeatedly observe signals and can accumulate information. We show that in this case it is optimal for agents to put increasing weight on prices as time goes on, and in the limit the private exogenous signal is completely ignored. This can open up the door for deviations of prices from fundamental values.

Fearing changes in fundamentals, agents could try to respond in two ways. In the tradition of the adaptive learning literature, they could estimate the moments of the distribution of the two signals using a constant gain algorithm, thus weighting most recent observations more heavily. Alternatively, in the Bayesian learning framework, they could explicitly allow for fundamentals to change over time.

We show that in both cases the problem of a disconnection between prices and fundamental values largely persists. We also establish an interesting connection between these two frameworks, and in particular between the constant gain in the adaptive learning algorithm and the subjective probability of changes in fundamentals in the Bayesian learning setting.

When introducing adaptive learning in the Bayesian framework, we depart from the assumption that agents know a priori the statistical properties of their environment and instead we require agents to learn such properties 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 empirical moments. 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 might fear changes in the statistical properties of their environment. The main result is that prices under-react to changes in fundamentals, particularly after a long period of stability: in fact, we will show that, if the change happens late enough in time, prices do not react at all. Higher gains would imply a better tracking performance, but at the cost of increased instability.

We show then that a similar outcome can emerge under Bayesian learning: in this case it a subjective probability of changes that governs a trade-off between the ability to track movements in the fundamental value and the volatility of prices. In particular, because Bayesian updating leads agents to put decreasing weight on the exogenous signal as time goes on, the subjective probability of changes in fundamentals must be much higher than their actual frequency for such changes to get factored into prices.

These results shows that uncertainty about fundamentals can lead agents to over-rely on public endogenous information and discount their private signals in a dynamic setting where information can be accumulated. This effect follows from the fact that the precision of the public signal is endogenous, and it improves over time faster than that of the private one. This finding reminds of the rational herding literature where agents, acting sequentially, end up relying only on information from previous agents, conveyed through their actions, rather than their own private information. In our setting, agents act all simultaneously, rather than sequentially, but repeatedly over time: in the limit, as the relative precision of the two signals changes, they disregard their own private information and all act only on the basis of the aggregate signal represented by prices.

1.1 Literature review

Our work touches upon different streams of literature. 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 a coordination problem with an endogenous signal is provided by Morris and Shin (2006). A notable application of this idea to asset prices is proposed by 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 static setting, though they do not consider the possibility of an endogenously changing weight on signals in a dynamic setting. Angeletos, Hellwig, Pavan (2007) propose instead a dynamic global game, but contrary to our framework, both public and private signals are there exogenous. Finally, Berardi (2015) considers a coordination problem with signal extraction and adaptive learning, but also in that setting both private and public signals are exogenous.

Our paper is also related to the literature on rational herding and informational cascades.

Banerjee (1992) proposes a model of herd behavior where agents follow what others are doing rather than using their own information. Bikhchandani, Hirshleifer and Welch (1992) define an informational cascade as occurring when it is optimal for an individual, after observing the actions of those ahead of him, to follow the behavior of those preceding individuals without regard to his own information. Welch (1992) analyses the rise of informational cascades in sequential sales in the market for initial public stock offerings and Devenow and Welch (1996) propose a review of papers on the economics of rational herding in financial markets.

Less directly related to our work, 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 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 (2014) use the concept of stochastic consistent expectations equilibrium to explain excess volatility in stock prices; finally, Adam et al (2016) show how adaptive learning on future prices can generate excess volatility. A key 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 and adaptive learning to understand the relation between asset prices and fundamental values.

2 Learning from prices

We assume there is an asset available for trade on the market, whose fundamental value is denoted by θ. We can think of such value as representing a measure of the present discounted value of future dividends. Assuming a constant flow of dividends d over the infinite future, for example, we would have

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

where $0 < \beta < 1$ is the discount factor. We will not be concerned here with the derivation of the fundamental value and simply assume it can be represented by θ.

Agents maximize their utility, which is a concave function of the value of their portfolio, expressed as difference between the fundamental value of the 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 prices observable,

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

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., Mele and Sangiorgi (2015) and Branch and Evans (2011)), in order to capture variations in the availability of publicly tradable shares (asset float).

An 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, $E_t\theta = \theta$ and $\sigma_w^2 = 0$ and thus prices would be constant at the fundamental value, i.e., $p_t = \theta$.

2.1 Information structure and equilibrium: static setting

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 $z_t(i) : [0,1] \to \mathbb{R}$, 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, $Z_t = \int_i z_t^i di$.

Nature moves first and draws θ from an improper uniform distribution over the positive real line $\mathbb{R}^+$. Agents don't observe directly θ but observe two signals on it: one, endogenous and public, from prices (p_t) and one, exogenous and private, from news (x_t^i). We can think of this last component as agents receiving different news because accessing different sources of information, or as a subjective interpretation of the same news. News could be about dividends or any other element that is perceived to affect the long term value of an asset.

The exogenous signal on dividends is represented as

$$x_t^i = \theta + v_{i,t}, \tag{3}$$

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

With signals 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

$$\tilde{\theta}_t^i = \alpha x_t^i + (1-\alpha)\, p_t, \tag{4}$$

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

$$\alpha^* = \frac{\sigma_p^2}{\sigma_p^2 + \sigma_v^2}, \tag{5}$$

where σ_p^2 represents the variance of prices. For generic parameterizations, $\alpha^* \in (0,1)$ and it is thus optimal for agents to put some weight on prices, together with the exogenous signal, when forming beliefs about fundamental values.

Individual demand is then given by

$$k_{i,t}^* = \frac{\alpha^*\left(x_t^i - p_t\right)}{\gamma \tilde{\sigma}_w^2}, \tag{6}$$

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

Aggregate demand is then given by

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

where aggregation is over the standard normal distribution, and prices evolve according to

$$p_t = \theta - \frac{\gamma\tilde{\sigma}_w^2}{\alpha^*}\varepsilon_t. \tag{7}$$

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

$$p_t = \theta - \gamma\sigma_v^2\varepsilon_t, \tag{8}$$

with

$$\sigma_p^2 = \gamma^2 \left(\sigma_v^2\right)^2 \sigma_\varepsilon^2, \tag{9}$$

which confirms that prices are normally distributed and conditionally independent from the private signals. As already noted by Angeletos and Werning (2006), we can also see from (9) that public information improves with private information. The linear equilibrium is here unique, defined by the optimal value of α^*.

An important element in the determination of the equilibrium level of α is the aggregation of the noise in the private signal. If the exogenous signal was instead public information and everyone was thus observing the same signal x_t, (5) would become

$$\alpha^* = \frac{E_t p_t^2 - E_t p_t x_t}{E_t p_t^2 + E_t x_t^2 - 2E_t p_t x_t} = \frac{\sigma_p^2 - \sigma_v^2}{\sigma_p^2 + \sigma_v^2 - 2\sigma_v^2} = 1. \tag{10}$$

Because the noise in the exogenous public signal would be transferred to prices, prices would be completely useless as a signal for the fundamental value, as they would encompass both the noise from the exogenous signal and the noise from supply: optimal α would thus be equal to 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 noise.

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

$$\alpha^* = \frac{\sigma_p^2}{\sigma_p^2 + \sigma_x^2} = \frac{\gamma^2 \left(\sigma_v^2\right)^2 \sigma_\varepsilon^2}{\sigma_v^2 + \gamma^2 \left(\sigma_v^2\right)^2 \sigma_\varepsilon^2} = \frac{\gamma^2\sigma_v^2\sigma_\varepsilon^2}{1 + \gamma^2\sigma_v^2\sigma_\varepsilon^2} \tag{11}$$

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 the 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 lose all informational 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 (11): no weight is put on the private signal and only prices are used. This might seem at first counter-intuitive, and looking at (5) one might actually mistakenly think that α^* goes instead to one as $\sigma_v^2 \to 0$. 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:[1] this is due to the fact that the volatility of prices originates from the variance of supply and demand (multiplicatively), and the latter is quadratic in the variance of the portfolio.

An important feature of the Bayesian equilibrium is that α^* is not a free parameter but depends instead on the deep structure of the model, as shown by (11). If α was to be considered as a free parameter, instead of as the outcome of an optimization problem, it would become an element affecting prices and their volatility.

Proposition 1 *In the static setting presented in this Section, the optimal Bayesian weight on private information is given by (11) and prices are defined by (8).*

It is clear, by comparing (8) to the full information equilibrium ($p_t = \theta$), that in a Bayesian equilibrium prices are characterized by excess volatility with respect to the (constant) fundamental value: uncertainty generates volatility.

2.2 Dynamic setting

In the previous section we have considered a static framework where information is only received once. It could be argued, though, that agents trading on markets play in fact a repeated game, in which information can be accumulated and beliefs revised over time. We extend therefore our framework here to account for these features.

Time is discrete and indexed by $t = \{1, 2, 3, ..\}$. Nature draws the fundamental θ at the beginning of time $t = 1$ and agents at every period receive a private and a public signal and thus accumulate information over time. As before, the public signal is represented by prices, and the private signal at each time t is given by (3).

[1] The variance of prices is in fact quadratic in σ_v^2: see (9).

We can write mean and precision of the posterior of θ at each time t conditional on the history of $x_{t,i}$ as

$$\bar{x}_{t,i} = \frac{\beta_{t-1}}{\beta_t}\bar{x}_{t-1,i} + \frac{\sigma_v^{-2}}{\beta_t}x_{t,i} = \frac{1}{t}\sum_{m=1}^{t} x_{m,i} \tag{12}$$

$$\beta_t = \beta_{t-1} + \sigma_v^{-2} = \sum_{\tau=1}^{t}\sigma_v^{-2} = \frac{t}{\sigma_v^2}, \tag{13}$$

with $\beta_0 = 0$ and $\bar{x}_{0,i} = 0$ (that is, $\bar{x}_{1,i} = x_{1,i}$).

Consider now the public endogenous signal that comes from prices. The mean and precision of the posterior of θ at time t, conditional on the history of p_t are given by

$$\bar{p}_t = \frac{\omega_{t-1}}{\omega_t}\bar{p}_{t-1} + \frac{\sigma_{p,t}^{-2}}{\omega_t}p_t = \frac{\sum_{j=1}^{t}\sigma_{p,j}^{-2}p_j}{\sum_{j=1}^{t}\sigma_{p,j}^{-2}} \tag{14}$$

$$\omega_t = \omega_{t-1} + \sigma_{p,t}^{-2} = \sum_{j=1}^{t}\sigma_{p,j}^{-2}, \tag{15}$$

with $\omega_0 = 0$ and $\bar{p}_0 = 0$ (that is, $\bar{p}_1 = p_1$).

Hence, conditional on the two signals, the expected value of the fundamental $E_t\left[\theta \mid \bar{x}_t^i, \bar{p}_t\right]$ is then given by

$$E_t\left[\theta \mid \bar{x}_t^i, \bar{p}_t\right] = \alpha_t\bar{x}_t^i + (1-\alpha_t)\,\bar{p}_t \tag{16}$$

with α_t determined by the relative precision of the two posteriors and given by

$$\alpha_t = \frac{\beta_t}{\beta_t + \omega_t}. \tag{17}$$

Using the demand equation

$$k_t^i = \frac{E_t^i\theta - p_t}{\gamma\sigma_{w,t}^2} \tag{18}$$

with

$$\sigma_{w,t}^2 = \alpha_t^2\beta_t^{-1} + (1-\alpha_t)^2\,\omega_t^{-1} = \frac{1}{\beta_t + \omega_t},$$

aggregating and imposing demand equal supply, we have

$$p_t = (\alpha_t\theta + (1-\alpha_t)\,\bar{p}_t) - \gamma\sigma_{w,t}^2\varepsilon_t. \tag{19}$$

Solving for p_t we then find

$$p_t = \frac{\beta_t}{\beta_t + \omega_{t-1}}\theta + \frac{\omega_{t-1}}{\beta_t + \omega_{t-1}}\bar{p}_{t-1} - \frac{\gamma\varepsilon_t}{\beta_t + \omega_{t-1}}. \tag{20}$$

Note that, with $\omega_0 = 0$ and $\bar{p}_0 = 0$, at time $t = 1$, p_1 from (20) reduces to (8), as in the static case, since there is no previous information to exploit.

From (20), it then follows that the conditional variance of prices is

$$\sigma^2_{p,t} = \frac{\gamma^2\sigma^2_\varepsilon}{(\beta_t + \omega_{t-1})^2}. \tag{21}$$

Note that both $\bar{x}_{t,i}$ and $\bar{p}_t$ are unbiased estimators for θ, since $E\bar{x}_{t,i} = E\bar{p}_t = \theta$.

2.3 Limiting dynamics

We now derive the limiting outcomes in this economy as $t \to \infty$. Starting with the exogenous signal, by the law of large numbers we have

$$\operatorname*{plim}_{t\to\infty} \bar{x}_{t,i} = \theta$$

and clearly

$$\lim_{t\to\infty} \beta_t = \frac{t}{\sigma^2_v} = \infty.$$

That is, the sample mean converges in probability to the mean of the distribution and its variance goes to zero. Consider then $\sigma^2_{w,t} = \frac{1}{\beta_t + \omega_t}$. Since $\lim_{t\to\infty}\beta_t = \infty$, and by definition $\omega_t \geq 0$, it follows that

$$\lim_{t\to\infty} \sigma^2_{w,t} = 0$$

and therefore, from (21),

$$\lim_{t\to\infty} \sigma^2_{p,t} = 0.$$

This last result also implies, from (15),

$$\lim_{t\to\infty} \omega_t = \infty.$$

Finally, since

$$\omega_t = \omega_{t-1} + \sigma^{-2}_{p,t} = \omega_{t-1} + \left(\gamma^2\sigma^2_\varepsilon\right)^{-1}(\beta_t + \omega_{t-1})^2,$$

it follows that, while β_t grows linearly in t, ω_t grows at a quadratic rate and thus

$$\lim_{t\to\infty} \alpha_t = 0,$$

which then, together with results about $\sigma^2_{w,t}$, implies

$$\lim_{t\to\infty} p_t = \bar{p}_t.$$

Given that $\bar{p}_t$ is a weighted average of all p_j, each one centered around θ and with decreasing variance (and that, from (14), the weight on each p_j is inversely proportional to its variance), it follows that

$$\lim_{t\to\infty} \bar{p}_t = \theta$$

and thus p_t converges over time to the fundamental θ.

Note that while both β_t and ω_t tend to ∞, β_t grows linearly while ω_t quadratically: both private and public information become infinitely precise in the limit, but the precision of the public signal improves faster and agents end up relying only on prices to infer fundamental values. Over time the variance of prices goes to zero: accumulated information reduces uncertainty and dampens volatility.

Proposition 2 *In the dynamic setting presented above, the optimal Bayesian weight on private information is given by (17) and prices evolve according to (19). In the limit, α_t converges to* 0 *and p_t converges to the fundamental value θ.*

Because agents try to minimize the variance of their portfolio by relying more on the less volatile signal, and because the variance of prices decreases to zero faster than that of private information, agents end up disregarding their private signals. This result suggests that, if fundamental values change, they might not get factored properly into prices, thus leading to a divergence between the two. In order to investigate this issue properly, though, we need to consider the signal extraction problem in a framework where the fundamental value is explicitly allowed to change.

2.4 Allowing for changes in fundamentals

We have considered so far the case where the environment is stationary and agents accordingly use past information to improve their estimates of the fundamental. Over time, as the precision of the public signal increases faster than the precision of the private one, α_t decreases to zero and agents end up relying only on prices in forming their beliefs: any change in the fundamental would thus be neglected. The Bayesian learning analysis carried out so far, though, was based on the assumption of a fixed fundamental, so there seems to be an inconsistency there in deriving conclusions about what happens if the fundamental changes. In order to shed light on this issue, we consider now a framework where agents allow explicitly for fundamentals to change over time. In particular, given that fundamental values are not observables, agents will never be sure whether a change has actually taken place at any specific time in the past, and only entertain subjective probabilities on such possibility.[2] This means that, with some fixed probability, information from each time in the past might have to be discarded as no longer relevant if the fundamental has since changed.

As before, Nature draws the fundamental from an improper uniform distribution over $\mathbb{R}^+$ at the beginning of time $t = 1$. Nature, though, can now also re-draw, with some fixed but unknown probability, a new value for the fundamental, from the same improper distribution, at the beginning of each time $t > 1$.

Agents are aware of this possibility, and entertain a subjective probability π that Nature re-draws the fundamental each period. Note that it is not important at this point whether π coincides or not with the actual probability by which Nature re-draws the fundamental over time. Later on, we will comment on this issue, and in fact show that π has to be greater than the actual frequency of changes in order to compensate for the decreasing value of α_t over time.

We first define

$$\bar{x}_{t,i}^{j\leq t} = \frac{1}{t-j+1}\sum_{n=j}^{t} x_{n,i},$$

the posterior at time t for $E_t(\theta \mid x)$ if a change in the fundamental had occurred at (the beginning of) time $j \leq t$. Then, the best predictor for the fundamental at time t, conditional on exogenous signals only, is given by

$$\tilde{x}_{t,i} = \sum_{j=1}^{t} a_j^t \bar{x}_{t,i}^j$$

where

$$\begin{aligned} a_1^t &= (1-\pi)^{t-1} \\ a_j^t &= (1-\pi)^{t-j}\pi, t \geq j > 1. \end{aligned}$$

[2]This modelling strategy can be thought of been particularly relevant for relatively small changes in the fundamental, which are more difficult to disentangle from the noise in the signal.

The coefficients a_j^t capture the probability that each (truncated) series $\bar{x}_{t,i}^j$ is the appropriate one for computing the conditional expected value of θ (that is, the perceived probability that Nature re-drew the fundamental at the beginning of time j).

Note that

$$\sum_{j=1}^{t} a_j^t = 1.$$

We can then rewrite $\tilde{x}_{t,i}$ as a weighted sum of current and past values of $x_{t,i}$

$$\tilde{x}_{t,i} = \sum_{j=1}^{t} h_j^t x_{j,i}, \tag{22}$$

where

$$h_1^t = \frac{(1-\pi)^{t-1}}{t} \tag{23}$$

$$h_j^t = h_{j-1} + \frac{(1-\pi)^{t-j}\pi}{t-j+1}, t \geq j > 1. \tag{24}$$

Again, $\sum_{j=1}^{t} h_j^t = 1$. We can then define

$$\tilde{\beta}_t = (var_t(\tilde{x}_{t,i}))^{-1} = \left(var_t\left(\sum_{j=1}^{t} h_j^t x_{j,i}\right)\right)^{-1}$$

and since all the $x_{j,i}$ are i.i.d. over time

$$\tilde{\beta}_t = \sigma_v^{-2}\left(\sum_{j=1}^{t}\left(h_j^t\right)^2\right)^{-1}.$$

As for the public signals, we define similarly the composite signal

$$\bar{p}_t^{j\leq t} = \frac{\sum_{n=j}^{t}\sigma_{p,n}^{-2}p_n}{\sum_{n=j}^{t}\sigma_{p,n}^{-2}}$$

which summarizes the relevant information from prices if the fundamental had changed at the beginning of time j. The best predictor for the fundamental at time t, conditional on endogenous signals only, is then given by

$$\tilde{p}_t = \sum_{j=1}^{t} a_j^t \bar{p}_t^j,$$

with each a_j^t defined as before. This can then be rewritten as a combination of current and past prices as follows

$$\tilde{p}_t = \sum_{j=1}^{t} k_j^t p_j,$$

where

$$\begin{aligned} k_1^t &= \frac{(1-\pi)^{t-1}\sigma_{p,1}^{-2}}{\sum_{m=1}^{t}\sigma_{p,m}^{-2}} \\ k_j^t &= \frac{\sigma_{p,j}^{-2}}{\sigma_{p,j-1}^{-2}} k_{j-1}^t + \frac{(1-\pi)^{t-j}\pi\sigma_{p,j}^{-2}}{\sum_{m=j}^{t}\sigma_{p,m}^{-2}}, t \geq j > 1 \end{aligned}$$

or

$$k_j^t = \sigma_{p,j}^{-2}\left[\frac{(1-\pi)^{t-1}}{\sum_{m=1}^{t}\sigma_{p,m}^{-2}} + \pi\sum_{m=2}^{j}\frac{(1-\pi)^{t-m}}{\sum_{z=m}^{t}\sigma_{p,z}^{-2}}\right] \text{ for } t \geq j > 1, \tag{25}$$

with $\sum_{j=1}^{t} k_j^t = 1$.

Note that as $\pi \to 0$, $k_j^t \to \frac{\sigma_{p,j}^{-2}}{\sum_{m=1}^{t}\sigma_{p,m}^{-2}}$, and as $\pi \to 1$, $k_{j\neq t}^t \to 0$ and $k_t^t \to 1$. That is, as $\pi \to 0$ the framework converges to the dynamic setting analyzed in Section 2.2, since no past information needs to be discounted for its probability of being no longer relevant. As $\pi \to 1$, instead, the framework converges to the static game seen in Section 2.1, since every period agents discount completely the past.[3]

We can then define

$$\tilde{\omega}_t = (var_t(\tilde{p}_t))^{-1} = \left(var_t\left(\sum_{j=1}^{t} k_j^t p_j\right)\right)^{-1}$$

and since prices are conditionally independent, we have

$$\tilde{\omega}_t = \left(\sum_{j=1}^{t}\left(k_j^t\right)^2\sigma_{p,j}^2\right)^{-1}. \tag{26}$$

It can be seen that, for $\pi = 0$, $\tilde{\omega}_t = \omega_t$ from (15).

Consider then the pricing equation

$$p_t = E_t\theta - \gamma\sigma_{w,t}^2\varepsilon_t \tag{27}$$

[3]For $\pi = 1$, there is an indeterminate form of 0^0 for k_1^t at $t = 1$: this is resolved by considering it equal to 1, which is correct if the exponential is solved before substituting out for π. Alternatively, we can take the limit as $\pi \to 1$.

where

$$E_t\theta = \int E_t^i \theta di = \tilde{\alpha}_t \int \tilde{x}_{t,i} di + (1 - \tilde{\alpha}_t)\, \tilde{p}_t$$

and

$$\tilde{\alpha}_t = \frac{\tilde{\beta}_t}{\tilde{\beta}_t + \tilde{\omega}_t}. \tag{28}$$

Since $\tilde{p}_t$ includes current prices, solving (27) for p_t we have

$$\begin{aligned} p_t &= \tilde{\alpha}_t \int \tilde{x}_{t,i} di + (1 - \tilde{\alpha}_t)\, \tilde{p}_t - \gamma \sigma_{w,t}^2 \varepsilon_t \\ &= \frac{\tilde{\alpha}_t}{1 - (1 - \tilde{\alpha}_t)\, k_t^t} \int \tilde{x}_{t,i} di + \frac{(1 - \tilde{\alpha}_t)}{1 - (1 - \tilde{\alpha}_t)\, k_t^t} \sum_{j=1}^{t-1} k_j^t p_j - \frac{\gamma \sigma_{w,t}^2 \varepsilon_t}{1 - (1 - \tilde{\alpha}_t)\, k_t^t}. \end{aligned} \tag{29}$$

The conditional variance of prices, σ_p^2, is endogenous and only implicitly defined by the above equations. From (29), we first derive

$$\sigma_{p,t}^2 = \frac{\gamma^2 \sigma_\varepsilon^2}{\left[\tilde{\beta}_t + \tilde{\omega}_t \left(1 - k_t^t\right)\right]^2}, \tag{30}$$

where both $\tilde{\omega}_t$ and k_t^t depend on $\sigma_{p,t}^2$. We then use (26) and (25) (with $j = t$) to substitute in and solve numerically for $\sigma_{p,t}^2$. Note that, for $\pi = 0$, $\sigma_{p,t}^2$ from (30) reduces to $\sigma_{p,t}^2$ from (21).

In order to understand the behavior of this system, we now simulate it.

2.4.1 Simulations

We run simulations for three cases: $\pi = .05; \pi = .5; = .95$, corresponding to the cases where agents believe there is, respectively, a 5%, 50% and 95% chance of change in the fundamental at every period. For illustrative purposes, the actual change in the fundamental takes place in all three cases at time $t = 50$, when it jumps from $\theta = 10$ to $\theta = 15$. As we will see, despite this being a rather large change (increase by 50%), it still does not get properly factored into prices even for relatively large π.

Figure 1 shows simulations for prices and α_t in the three different scenarios: in all cases simulations are run with the same series of shocks and parameterization, apart from π. In particular, we set $\gamma = .75$ and $\sigma_v^2 = \sigma_\varepsilon^2 = 1$.[4] We see that the only case where the change

[4]We have carried out sensitivity analysis with respect to these parameters and results are robust over large intervals.

in the fundamental value is properly transferred onto prices is for $\pi = .95$, which is also the only case for which α_t remains substantially bounded away from zero. The reason why α_t decreases towards zero even for relatively high values of π is that while β_t^{-1} remains now strictly positive, ω_t^{-1} still decreases towards zero fast: the uncertainty about changes in fundamentals is not high enough to compensate for the rapidly decreasing variance of prices. Moreover, similar results are obtained if the actual change in the fundamental takes place much sooner (or, of course, later) than $t = 50$, as it can be inferred by looking at the behavior of α_t in the three right hand side panels in Figure 1: in particular, for $\pi = .05$, any change in the exogenous signal that takes place any time after the very first few periods is completely neglected, while for $\pi = .5$ things are slightly better, with changes up to around period 15 being factored into prices to some extend.

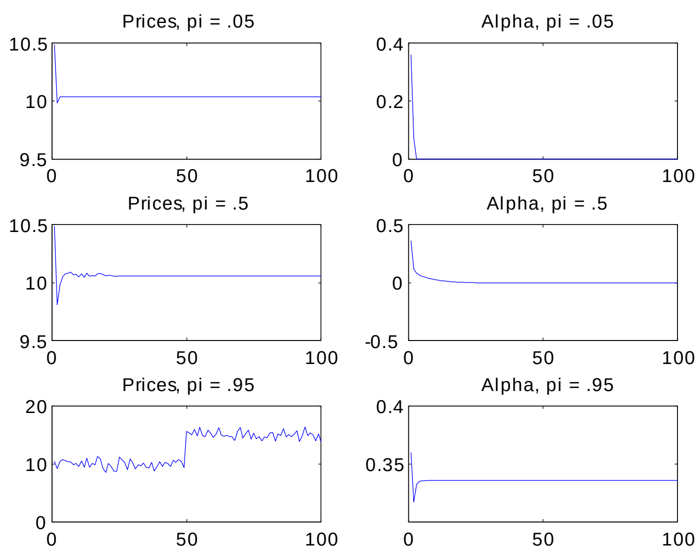

Figure 1: Dynamics of prices and α for different values of π.

With 5% probability of change, movements in fundamentals are expected to happen on average once every 20 periods, more often with higher subjective probabilities: even though agents here are overestimating the probability of changes in all three cases, they still fail to factor them properly into prices most of the times (except for the rather extreme case of almost certainty of changes, with $\pi = .95$): this is due to the fact that very high subjective probabilities are required in order to compensate for the tendency of the optimal weight on the exogenous information to decrease rapidly to zero in these dynamic settings under

Bayesian learning.

It follows that, if instead of modelling π as a subjective probability we were to treat it as the actual probability of changes, for example estimated from frequencies (assuming they were somehow observables to agents), agents would always miss out on movements in fundamentals, unless they were happening almost every period.

Treating π as a free parameter instead introduces a trade-off for agents: high π allows changes in fundamentals to be tracked and properly transferred onto prices, but at the cost of increased volatility in prices, as can be seen in the third raw of Figure 1. There is thus a trade off, governed by π, between volatility in prices and the ability to track changes in fundamentals.

3 Adaptive learning

Bayesian learning assumes agents know the relative precision of the different signals they receive, and use such information optimally. We will now relax the first assumption by merging Bayesian learning with adaptive learning, using the framework of Section (2.2). One essential element of any learning analysis is to decide what aspects of their decision problem agents should be learning about. We take here 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.

In particular, agents will need to learn about mean and variances of the two signals they receive in order to apply optimal weights. While prices are publicly observable, the exogenous signal is instead idiosyncratic: it follows that, under learning, agents will hold homogeneous beliefs about prices (assuming common initial beliefs and updating algorithms) but heterogeneous beliefs about the statistical properties of the exogenous signal.

We first provide convergence results under decreasing gain in a stationary environment, and then consider the case of constant gain, which will allow us to draw a connection with the framework of Section 2.4. Constant gain learning, in fact, places relatively more weight on most recent observations and would seem to be well suited to address the problems arising with changing fundamentals in the signal extraction problem.

In order to implement their optimal demand (18), agents will need to learn about

$\bar{x}_{t,i}, \bar{p}_t, \beta_t$ and ω_t. At each time t, each agent i observes the whole history from 1 to t of his own private signal $x_{t,i}$ and of prices p_t. An agent i at time t is thus defined by the history of signals he receives and his beliefs system $[s^i_{x,t}, \hat{x}^i_t, \hat{\beta}^i_t; s_{p,t}, \hat{p}_t, \hat{\omega}_t, s^i_{w,t}]$, which is the learning counterpart to $[E[x^i_t]^2, \bar{x}^i_t, \beta^i_t; E[p_t]^2, \bar{p}_t, \omega_t, \sigma^2_{w,t}]$. Individual demand is then given by

$$k^i_t = \frac{\hat{\alpha}^i_t \hat{x}^i_t + \left(1 - \hat{\alpha}^i_t\right) \hat{p}_{t-1} - p_t}{\gamma s^i_{w,t}}, \tag{31}$$

where

$$\hat{\alpha}^i_t = \frac{\hat{\beta}^i_t}{\hat{\beta}^i_t + \hat{\omega}_t}, \tag{32}$$

with

$$\hat{\beta}^i_t = \hat{\beta}^i_{t-1} + \left(s^i_{x,t} - \left(\hat{x}^i_t\right)^2\right)^{-1} \tag{33}$$

$$\hat{\omega}_t = \hat{\omega}_{t-1} + \left(s_{p,t} - \left(\hat{p}_t\right)^2\right)^{-1} \tag{34}$$

and

$$s^i_{w,t} = \left(\hat{\alpha}^i_t\right)^2 \left(\hat{\beta}^i_t\right)^{-1} + \left(1 - \hat{\alpha}^i_t\right)^2 \hat{\omega}^{-1}_t = \frac{1}{\hat{\beta}^i_t + \hat{\omega}_t}. \tag{35}$$

Equations (33)-(34) say that agents use estimated means and second raw moments (based on observables) in order to derive the relevant central moments (which can not be estimated directly).

Aggregate demand is thus given by

$$K_t = \int \frac{\left(\hat{\alpha}^i_t \hat{x}^i_t + \left(1 - \hat{\alpha}^i_t\right) \hat{p}_t\right) - p_t}{\gamma s^i_{w,t}} di. \tag{36}$$

From market clearing conditions, it follows that prices evolve according to

$$p_t = \frac{\left(\int \frac{\hat{\alpha}^i_t}{s^i_{w,t}} di\right) \theta + \left(\int \frac{\left(1 - \hat{\alpha}^i_t\right)}{s^i_{w,t}} di\right) \hat{p}_t - \gamma \varepsilon_t}{\int \frac{1}{s^i_{w,t}} di}, \tag{37}$$

or, since $\hat{p}_{t-1}$ is common across agents, x^i_t does not covariate with the other state variables (see below the updating rules for time t variables in the beliefs system, which all depend on time $t-1$ information) and $\int x_t^i di = \theta$

$$p_t = \frac{\int \hat{\beta}_t^i di}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}}\theta + \frac{\hat{\omega}_{t-1}}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}}\hat{p}_{t-1} - \frac{\gamma\varepsilon_t}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}}, \tag{38}$$

which is the learning counterpart of (20).

We now define how beliefs are updated through recursive algorithms that implement adaptive learning. Agents learn about raw second moments according to

$$s_{x,t}^i = s_{x,t-1}^i + g_t\left[\left(x_t^i\right)^2 - s_{x,t-1}^i\right] \tag{39}$$

$$s_{p,t} = s_{p,t-1} + g_t\left[\left(p_t\right)^2 - s_{p,t-1}\right], \tag{40}$$

and the means of the two signals according to

$$\hat{x}_t^i = \hat{x}_{t-1}^i + g_t\left[x_t^i - \hat{x}_{t-1}^i\right] \tag{41}$$

$$\hat{p}_t = \hat{p}_{t-1} + g_t\left[p_t - \hat{p}_{t-1}\right]. \tag{42}$$

Equations (38)-(42), together with (33)-(34), represent the dynamics of the system under adaptive learning. It is a non-linear stochastic recursive algorithm and its long-run properties can be analyzed using stochastic approximation techniques. 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 the case of constant gain.

3.1 Convergence results

We analyze here convergence results for the decreasing gain case $g_t = 1/t$, which reduces the learning algorithms to instances of recursive least squares estimates. Since agents are estimating only means and second moments, 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. Using stochastic approximation techniques (see Evans and Honkapohja (2001) for details), we can represent the evolution, in notional time τ, of estimated variables through ordinary differential equations (ODEs).

Convergence of (39) is easy to establish, since x_t^i is exogenous and agents are effectively estimating the value of an exogenous constant (the second moment of the distribution). The relevant ODE can be derived as follows

$$\frac{ds_x^i}{d\tau} = \lim_{t\to\infty} E\left(\left(x_t^i\right)^2 - s_{x,t-1}^i\right) = \theta^2 + \sigma_v^2 - s_x^i, \tag{43}$$

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 (43) is equal to -1.

Similarly for (41), which gives rise to the ODE

$$\frac{d\hat{x}^i}{d\tau} = \lim_{t\to\infty} E\left(x_t^i - \hat{x}_{t-1}\right) = \theta - \hat{x}^i \tag{44}$$

with stable fixed point $\hat{x}^i = \theta$, $\forall i$.

The two results above imply that $\left[s_{x,t}^i - (\hat{x}_t^i)^2\right] \to \sigma_v^2$ in (33) and thus $\hat{\beta}_t^i \to \infty$.

In terms of the endogenous variables, starting from $\hat{p}_t$, we have that, since ε_t is a zero mean i.i.d. process independent from other variables, the relevant ODE derived from (42) is

$$\begin{aligned}
\frac{d\hat{p}}{d\tau} &= \lim_{t\to\infty} E\left(\frac{\int \hat{\beta}_t^i di}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}}\theta + \frac{\hat{\omega}_{t-1}}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}}\hat{p}_{t-1} - \frac{\gamma\varepsilon_t}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}} - \hat{p}_{t-1}\right) \\
&= \lim_{t\to\infty} E\left(\frac{\int \hat{\beta}_t^i di}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}}\theta + \left(\frac{\hat{\omega}_{t-1}}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}} - 1\right)\hat{p}_{t-1}\right) \\
&= \lim_{t\to\infty} \frac{\int \hat{\beta}_t^i di}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}}\left(\theta - \hat{p}\right)
\end{aligned} \tag{45}$$

whose fixed point is $\hat{p} = \theta$. Stability depends on $\lim_{t\to\infty} -\frac{\int \hat{\beta}_t^i di}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}}$, which is negative and, we will see, converges asymptotically to zero: in the limit, deviations of prices from the fundamental will stop being factored into estimates.

In terms of the second raw moment, we have

$$\begin{aligned}
\frac{ds_p}{d\tau} &= \lim_{t\to\infty} E\left(\begin{array}{c}\left(\frac{\int \hat{\beta}_t^i di}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}}\theta + \frac{\hat{\omega}_{t-1}}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}}\hat{p}_{t-1} - \frac{\gamma\varepsilon_t}{\left(\int \hat{\beta}_{t-1}^i di + \hat{\omega}_{t-1}\right)}\right)^2 + \ldots \\ \ldots - s_{p,t-1}\end{array}\right) \\
&= \lim_{t\to\infty}\left(\begin{array}{c}\left(\frac{\int \hat{\beta}_t^i di}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}}\right)^2\theta^2 + \left(\frac{\hat{\omega}_{t-1}}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}}\right)^2(\hat{p}_{t-1})^2 + \frac{\gamma^2\sigma_\varepsilon^2}{\left(\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}\right)^2} + \ldots \\ \ldots + 2\frac{\int \hat{\beta}_t^i di}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}}\frac{\hat{\omega}_{t-1}}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}}\theta\hat{p}_{t-1} - s_{p,t-1}\end{array}\right)
\end{aligned} \tag{46}$$

and since $\hat{\beta}_t^i \to \infty$ and $\hat{\omega}_t > 0$ we have that $\lim_{t\to\infty} \frac{\gamma^2\sigma_\varepsilon^2}{\left(\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}\right)^2} = 0$ and thus

$$\begin{aligned}
\frac{ds_p}{d\tau} &= \lim_{t\to\infty}\left(\frac{\left(\int \hat{\beta}_t^i di\right)^2 + (\hat{\omega}_{t-1})^2 + 2\left(\int \hat{\beta}_t^i di\right)\hat{\omega}_{t-1}}{\left(\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}\right)^2}\theta^2 - s_{p,t-1}\right) \\
&= \theta^2 - s_p \qquad (47)
\end{aligned}$$

which then implies that $s_p \to \theta^2$: the variance of prices thus converges to zero, since the second raw moment converges to the squared mean.

Previous results imply that $\hat{\omega}_t \to \infty$ and $s_{w,t}^i = \frac{1}{\hat{\beta}_t^i + \hat{\omega}_t} \to 0$. Finally, convergence in all the learning algorithms ensures that α_t^i also converges to its long run equilibrium value, i.e.,

$$\lim_{t\to\infty} \alpha_t^i = \frac{\hat{\beta}_t^i}{\hat{\beta}_t^i + \hat{\omega}_t} = 0, \ \forall i.$$

The economy under learning thus converges to the same equilibrium derived under full information of the statistical properties of the economy. Moreover, the equilibrium is stable under learning (E-stable). We summarize these results in the following proposition:

Proposition 3 *Consider the system of ODEs (43), (44), (45), (47), defining learning in notional time. All ODEs are stable and agents learn the means and variances of the relevant exogenous and endogenous variables: it follows that $\hat{\alpha}_t^i$ converges to 0 as t grows and the system converges to the Bayesian equilibrium defined in Section (2.2).*

3.2 Constant gain

Though the decreasing gain case allows us to establish convergence of the model under adaptive learning to the Bayesian equilibrium, in light of the previous results we are more interested here in the dynamics of the system under constant gain learning. It could be argued, in fact, that agents trying to learn the fundamental value of an asset might want to use a constant gain algorithm, in order 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 works of Sargent (1999) and Sargent at el. (2006)) to the causes of business cycles (e.g., Milani (2011) and Eusepi and Preston (2011)).

The constant gain assumption seems well suited for our purpose of analyzing the dynamics in an environment where structural breaks in the fundamental can happen: by discounting past observations, it allows new information to play a larger role in the determination of beliefs.

We will show that, with constant fundamental value, estimated moments converge to their equilibrium values fairly quickly, both for the endogenous and exogenous variables. After a change in fundamentals, instead, a disconnection between the exogenous and endogenous variables arises: while beliefs for the exogenous variables are quick to re-adjust to the new values, because the constant-gain implies heavier weight on new observations, such new information is not passed on to prices, because of the effect of a low (and decreasing) weight on the exogenous signal (α_t) in the signal extraction problem. This also implies that the under-reaction is more pronounced the later the change in fundamental takes place, since α_t is further out on its convergence path towards zero.

We present now results from simulations. We choose a value for the gain parameter of $g = 0.025$, which is a fairly common value in the macroeconomic applied literature. Though there is no direct evidence of the appropriate value for such parameter, Berardi and Galimberti (2015) provide a thorough discussion of the role and estimates bands for the gain parameter in macroeconomic applications. In general, higher gains imply faster reaction to changes, but more volatile estimates. We let the model run for 100 periods with the baseline parameterization ($\sigma_v^2 = \sigma_\varepsilon^2 = 1; \gamma = .75; g = 0.025$) and then at $t = 100$ a structural break occurs and the fundamental increases from 10 to 15. We then repeat the same exercise (with the same history of disturbances) with the break instead at $t = 200$. Results are reported in Figure 2, where we show the evolution of prices and α_t. We can see that prices largely under-react to the change in the fundamental when the break happen at $t = 100$ (solid line), and they barely move at all if it takes place at $t = 200$. As mentioned before, these results follow from the fact that, as time goes by, agents become more and more relying on prices to predict fundamentals (that is, α_t decreases): when a change in the fundamental occurs, then, it doesn't get factored into prices properly.

In Figure 3 we then show the evolution of beliefs about the moments of the exogenous and endogenous signals for the case with the break at $t = 100$. We can see that, because of the constant gain, agents are quick to learn about the new mean and second moment of the exogenous signal, while mean and second moment of prices under-react, for the reason discussed above. These results show that, while constant gain is effective in adjusting estimates about the exogenous variables after a change, it can not solve the problem of the under-reaction of prices, because the effect of α_t in the signal extraction problem limits the pass-through from exogenous signals to prices.

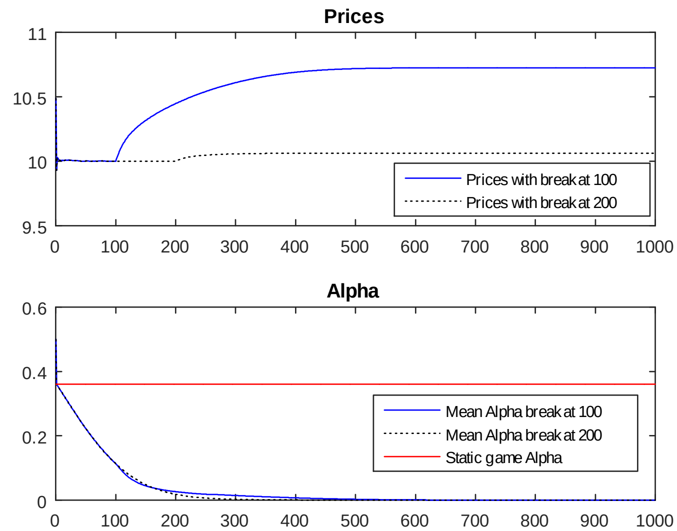

Figure 2: Prices and α with structural breaks in fundamental at $t = 100$ or $t = 200$.

4 Relation between g and π

In light of results so far, it is instructive to analyze the relationship between the adaptive learning gain g and parameter π in the Bayesian learning framework. The gain parameter in an adaptive learning algorithm determines the weight put on past observations: with a decreasing gain $1/t$, all observations receive equal weight; with a constant gain g instead the weight decays exponentially with past observations. A similar interpretation can be given to π, but for different reasons. The parameter π represents the perceived probability of a change in the fundamental at every time. The probability that a time $t-j$ observation is relevant for time t inference is thus $(1-\pi)^{t-j}$: again, the weight decays exponentially as we move back in time.

Formally, if we look at the exogenous private signal (the same considerations hold for the endogenous signal), we have that the updating rule for the adaptive learning scheme, equation (41), with constant gain leads to

$$\begin{aligned}\hat{x}_t^i &= \hat{x}_{t-1}^i + g\left[x_t^i - \hat{x}_{t-1}^i\right] = (1-g)\,\hat{x}_{t-1}^i + gx_t^i \\ &= g\sum_{j=1}^{t}(1-g)^{t-j}\,x_j^i,\end{aligned}$$

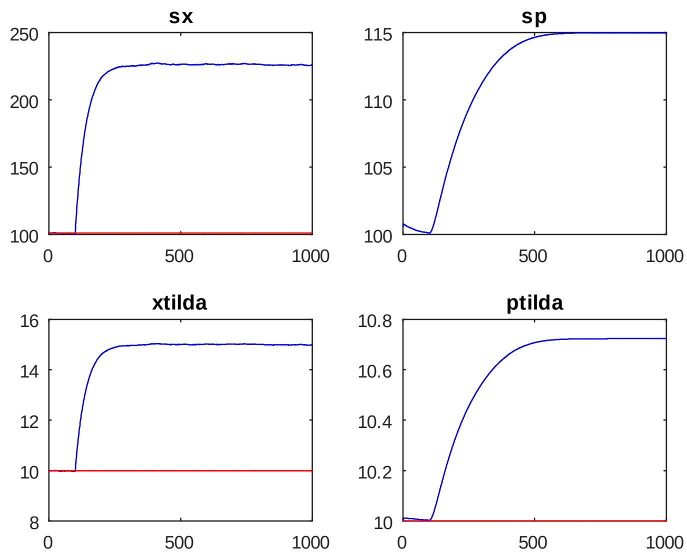

Figure 3: Beliefs about moments of exogenous and endogenous signals.

assuming $\hat{x}_0^i = 0$. We thus define the weight at time t on observation from time j as

$$b_j^t = g\,(1-g)^{t-j}$$

for $j = 1, ..., t$.

In the Bayesian framework, we have that $\tilde{x}_{t,i}$ is given by (22), with weights from (23)-(24), which, in a non-recursive way, can be rewritten as

$$\begin{aligned} h_1^t &= \frac{(1-\pi)^{t-1}}{t} \\ h_j^t &= \frac{(1-\pi)^{t-1}}{t} + \sum_{m=2}^{j} \frac{(1-\pi)^{t-m}\,\pi}{t-m+1}. \end{aligned}$$

While the weighting structure in the Bayesian framework is more convoluted, we can see that both b_j^t and the leading term in h_j^t (represented by $\frac{(1-\pi)^{t-j}\pi}{t-j+1}$) decay exponentially, leading to similar weight profiles on older observations in both cases. In order to provide some more insight into such weighting structures, we draw b_j^t and h_j^t in Figure 4. Curves are computed for $g = 0.025$ and $\pi = 0.01$, with $t = 100$. We can see that, despite being derived in different frameworks and through different assumptions, the shape of the two weighting structures is remarkably similar, leading to similar weighting on past information in the two cases.

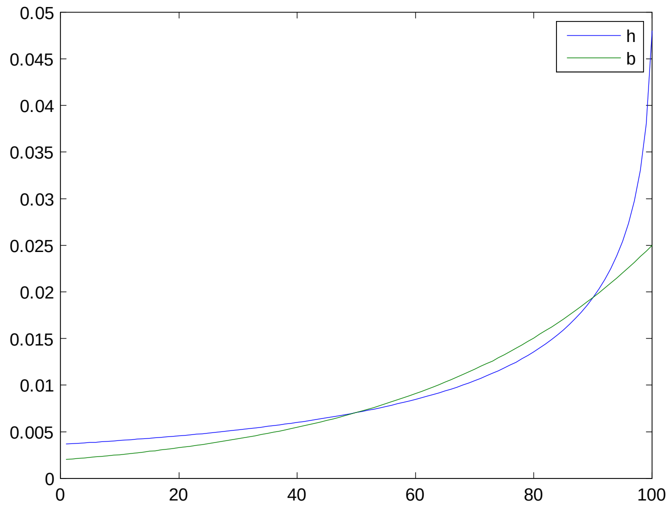

Figure 4: Weights on past observation under constant gain adaptive learning (b) and Bayesian learning (h).

5 Discussion and conclusions

We have proposed in this paper a model of uncertainty and learning about fundamental values. Agents are faced with a signal extraction problem, which in a static setting introduces volatility in prices compared to the full information case, where prices would simply coincide with the fundamental value at all times. The possibility to learn over time in a repeated game drives the volatility of prices to zero, but it also implies that all weight in the limit is put on the public endogenous signal and none on the private exogenous one. As agents end up relying only on prices, a form of rational herding with an informational cascade emerges: private information is disregarded and only the public signal is used.

The literature on informational cascades has usually focused on sequential games, where subsequent agents discard their own private information and base their actions only on information received from previous agents' behavior. Though in our model all agents act instead simultaneously, in a repeated game the outcome is similar: private information gets neglected in favour of public one. Bikhchandani, Hirshleifer and Welch (1992) write: "The problem with cascades is that they prevent the aggregation of information of numerous individuals." In our framework, similarly, as α_t converges to zero, private information of each individual about the fundamental value of the asset is neglected and does not contribute to the determination of prices: aggregation fails.

In order to understand the robustness of this result, we extended the framework in two directions. First, in the Bayesian rational learning setting, we introduced an explicit subjective probability of change in the fundamental value: in this way agents explicitly allow for movements in the variable they are trying to infer through their signal extraction problem. Second, in the adaptive learning setting, we allowed agents to discount past information through a constant gain algorithm. In both cases our original result is shown to persist, though somewhat mitigated.

This paper is the first, to our knowledge, to consider a signal extraction problem in a repeated game with endogenous public information. The fact that public information is endogenous means that its variance changes over time, thus affecting the optimal weight on public versus private signals, which is time varying. In the limit, as the precision of the public signal improves faster, private information is discarded. We have presented the analysis in an asset pricing model with fundamentalist agents, but results can be relevant to any dynamic setting where agents' optimal actions depend on a signal extraction problem and in turn determine the quality of the endogenous public signal being used.

6 Appendix

6.1 Derivation of α^* with private signal

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

$$\alpha^* = \arg\min_{\alpha} E_t \left(\theta - \tilde{\theta}_t^i\right)^2 \tag{48}$$

with

$$\tilde{\theta}_t^i = \alpha x_t^i + (1-\alpha)\, p_t. \tag{49}$$

Minimizing (48) subject to (49) leads to the FOC

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

whose solution implies

$$\alpha^* = \frac{E_t p_t^2 - E_t p_t x_t^i}{E_t p_t^2 + E_t \left(x_t^i\right)^2 - 2 E_t p_t x_t^i}. \tag{50}$$

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

$$\alpha^* = \frac{\sigma_p^2}{\sigma_p^2 + \sigma_v^2}. \tag{51}$$

Individual demand is then given by

$$k_{i,t}^* = \frac{\alpha^* \left(x_t^i - p_t\right)}{\gamma \tilde{\sigma}_w^2}, \tag{52}$$

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

$$\begin{aligned}\tilde{\sigma}_w^2 &= E_t\left[\left(\theta - \tilde{\theta}_t^i\right)^2 \mid x_t^i, p_t\right] \\ &= \alpha^2 \sigma_v^2 + (1-\alpha)^2 \sigma_p^2. \end{aligned} \tag{53}$$

Aggregating individual demand (52) and equating it with supply we obtain the price equation

$$p_t = \theta - \frac{\gamma \tilde{\sigma}_w^2}{\alpha^*} \varepsilon_t \tag{54}$$

which implies

$$\sigma_p^2 = \left(\frac{\gamma \tilde{\sigma}_w^2}{\alpha^*}\right)^2 \sigma_\varepsilon^2. \tag{55}$$

It follows from (51) and (53) 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

$$\tilde{\sigma}_w^2 = \frac{\left(\sigma_p^2\right)^2 \sigma_v^2 + \sigma_p^2 \left(\sigma_v^2\right)^2}{\left(\sigma_v^2 + \sigma_p^2\right)^2} = \frac{\left(\sigma_v^2 + \sigma_p^2\right) \sigma_p^2 \sigma_v^2}{\left(\sigma_v^2 + \sigma_p^2\right)^2} = \frac{\sigma_p^2 \sigma_v^2}{\sigma_v^2 + \sigma_p^2} = \alpha^* \sigma_v^2. \tag{56}$$

Substituting (56) into (54) 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., 2016. Stock market volatility and learning. Journal of Finance 71, 33-82.

[2] Angeletos, G.M., Hellwig, C., Pavan, A., 2007. Dynamic Global Games of Regime Change: Learning, Multiplicity, and the Timing of Attacks. Econometrica 75, 711-756.

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

[4] Banerjee, A., 1992. A Simple Model of Herd Behavior. Quarterly Journal of Economics 107, 797-817.

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

[6] Berardi, M., Galimberti, J.K., 2015. Empirical calibration of adaptive learning. KOF Working Paper n. 392.

[7] Bikhchandani, Sushil, David Hirshleifer and Ivo Welch, 1992. A theory of fads, fashion, custom, and cultural change as informational cascades, Journal of Political Economy 100, 992-1026

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

[9] 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.

[10] 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.

[11] 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.

[12] Devenow, Andrea and Ivo Welch, 1996. Rational herding in financial economics. European Economic Review 40, 603-615.

[13] Evans, G.W., Honkapohja, S., 2001. Learning and Expectations in Macroeconomics. Princeton University Press.

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

[15] Hommes, C., Zhu, M., 2014. Behavioral learning equilibria. Journal of Economic Theory 150, 778-814.

[16] Mele, A., Sangiorgi, F., 2015. Uncertainty, information acquisition, and price swings in asset markets. Review of Economic Studies 82, 1533-1567.

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

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

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

[20] 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.

[21] Welch, Ivo, 1992. Sequential sales, learning, and cascades, Journal of Finance 47, 695-732.