Discussion Paper Series

Information aggregation and learning in a dynamic asset pricing model

By

Michele Berardi

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

June 2018

Number 241

Information aggregation and learning in a dynamic asset pricing model*

Michele Berardi
The University of Manchester

June 26, 2018

Abstract

This paper analyses a dynamic framework where an unobservable fundamental can be learned over time through two signals: one exogenous and private and the other, prices, endogenous and public. As information cumulates over time through Bayesian learning, prices become fully revealing and agents disregard their private information, suggesting a possible route through which fundamental values and prices can become misaligned. The analysis is then extended to a setting where agents need to infer the statistical properties of the signals they receive, merging Bayesian with adaptive learning. By introducing uncertainty about the moments of the relevant distributions used for Bayesian learning, adaptive learning can improve the ability of prices to track changes in fundamentals and thus their efficiency.

Key words: uncertainty, information, Bayesian learning, adaptive learning, asset prices.

JEL classification: D83, D84, G12, G14.

* I would like to thank participants at the Royal Economic Society 2016 Conference, the Computations in Economics and Finance 2016 Conference, the Manchester Economic Theory Workshop 2017, the Computations in Economics and Finance 2017 Conference and at various seminars, for useful comments, suggestions and discussions on early drafts of this paper. I claim ownership of all remaining errors.

Information aggregation and learning in a dynamic asset pricing model

1 Introduction

The fundamental value of an asset summarizes the present value of 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 values in a context of imperfect information. To this end, I merge two lines of literature: one on signal extraction and Bayesian learning, the other on adaptive learning. From the first, I 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, I 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. The main results are that: i) under Bayesian learning, public information can harm agents' incentives to acquire private information and lead to a disconnection between prices and fundamentals; and ii) adaptive learning, by introducing uncertainty about the moments of the relevant distributions used for Bayesian learning, can improve the outcomes and price efficiency.

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 is a crucial element for investment strategies of fundamentalist traders, who want to buy assets that are underpriced and sell assets that are overpriced with respect to their fundamental value. To better isolate the link between fundamentals and prices, I assume traders are only concerned about the fundamental value of their portfolio, relative to its price, and they do not try to profit from exploiting short term capital gains. As the fundamental value is not known and never observable, it can only be inferred using observables such as news on a firm's profitability and prices as indirect information.

In particular, I assume that agents receive a noisy exogenous private signal about the fundamental value of an asset: this could be thought of as the subjective interpretation of news about the value and profitability of a firm. Besides this exogenous signal, agents also use prices in their inference, as prices 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 I derive such weight in a static setting and discuss its implications for asset prices. I then extend the model to a dynamic framework where agents repeatedly observe signals and can cumulate information over time. I 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 as prices become fully revealing. This can open up the door for deviations of prices from fundamental values, should such fundamentals change.

In order to investigate this possibility, I extend the framework to allow for the possibility of changes in the fundamental: as agents are aware of this possibility, they discount past observations accordingly. Numerical results show that a key parameter in this setting is the probability that the fundamental value changes at each period: for values of such probability not too high, the weight on private information still decreases to zero over time. Only for relatively high values of such probability, private information is used asymptotically. Such a probability effectively governs a trade-off between the ability to track movements in the fundamental and the volatility of prices.

Bayesian learning relies on the assumption that agents know the precision of the signals they receive, in order to combine them optimally in their inference. This is a strong assumption, especially in contexts where such precision evolves over time endogenously, as it is the case here for prices. I thus extend the analysis by combining Bayesian with adaptive learning, departing from the assumption that agents know a priori the statistical properties of their environment and instead assuming that they need to learn such properties through experience.

I first show that, under a decreasing gain algorithm, learning converges to the Bayesian equilibrium. This is not an obvious result: because prices are endogenous and depend on agents' beliefs, higher beliefs could lead to higher prices and thus to even higher beliefs, in a self-reinforcing destabilizing loop. This doesn't happen, though, because Bayesian weights have a stabilizing effect: a higher variance of prices decreases the relative weight put on prices in the signal extraction problem, thus helping stabilizing the system.

Adaptive learning is particularly suited to investigate the case where fundamentals can change over time: in particular, a constant gain algorithm has been suggested in the literature as an effective way to track changes in estimated parameters. I thus substitute the decreasing gain with a constant gain in the learning algorithm, capturing the idea that agents fear changes in the statistical properties of their environment. This framework allows me to highlight a connection between the constant gain parameter in the adaptive learning algorithm and the probability of changes in fundamentals in the Bayesian learning setting.

While the constant gain allows agents to effectively capture changes in the exogenous signal, it still doesn't fully prevent them from relying only on prices in the long run and discard private information. Still, for sensible values of the gain parameter, which governs how new information gets incorporated into estimates, agents keep using their private information for a long time and prices can thus reflect changes in fundamentals.

Results in this paper show that an over-reliance on public signals can arise when information is accumulated over time: 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 signal. This finding is reminiscent 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 my 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. Adaptive learning, introducing uncertainty over the precision of the signals, can attenuate this effect and effectively allow prices to incorporate private information over possible changes in fundamentals over time.

1.1 Literature review

This work touches upon different streams of literature. The main link is with the body of work on noisy rational expectations equilibria and information aggregation in asset markets. In noisy rational expectations models, prices reveal only partially the information available to agents, operating as noisy aggregators for information. In seminal work, Grossman (1976) and Grossman (1978) show how prices can aggregate information perfectly and substitute for private information on capital markets; Hellwig (1980) and Diamond and Verrecchia (1981) instead show conditions under which prices can only be partially revealing of the private information of agents, resulting in a noisy and partially aggregating equilibrium. My work will show how, under certain conditions, prices can effectively cease to aggregate private information. Admati (1985) extends the framework to a multi-asset market and finds that the presence of many risky assets introduces novel features into the noisy equilibrium. Grossman (1981) provides a general discussion of the informational role played by prices in contexts of asymmetric information, while Admati (1991) considers the same problem from the perspective of the market microstructure and the impact of different trading arrangements on market performance when agents have asymmetric information. Particularly relevant for my paper is the line of work on dynamic noisy rational expectations equilibria. Vives (1993) and

Vives (1995b) study, respectively, the rate at which dispersed information is incorporated into prices and asymmetrically informed agents can learn the return of an asset. Kyle (1985) analyses how quickly new private information about the underlying value of a speculative commodity gets incorporated into market prices and how it affects the liquidity of the market. Vives (1995a) considers the effect of investors' horizons on the information content of prices and shows that such effect depends on the temporal pattern of private information arrival. In particular, long horizons reduce the final price informativeness when the arrival of information is diffuse; nevertheless, as the number of trading periods increases without bound, prices converge to the fundamental value, as it is the case in my work with constant fundamental value. Amador and Weill (2012) consider a dynamic framework where agents receive repeated public and private signals about the state of the world. A key difference with my work is that in their case, while the initial signals realizations are centered around the true state of the world, subsequent realizations are centered around the average population action: that is, after the first period, all new information about the state of the world comes from others and there is no additional arrival of "new" information.

Signal extraction problems and Bayesian learning have been applied in the literature on global games in order to analyze various scenarios where agents face coordination problems with heterogeneous information. While often agents rely on exogenous signals in this line of literature, an analysis of a coordination problem with an endogenous signal is provided by Allen et al. (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 the static setting proposed in this paper. Angeletos, Hellwig and Pavan (2007) propose instead a dynamic global game, where agents can take repeated actions, but contrary to my framework, both public and private signals are there exogenous. The endogeneity of prices as signal is crucial for results in this paper.

My work 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. While

I do not have a proper informational cascade in this work, some of the results I find share the same flavour, as under certain conditions agents will discard their own private information in their decision problem.

Bray and Kreps (1987) highlight the distinction between learning within and learning about an equilibrium: I will consider both types of learning, within and about the equilibrium, in the form of, respectively, Bayesian and adaptive learning, and shed some light on the relationship between the two concepts. Berardi (2015) considers a coordination problem with both Bayesian and adaptive learning, but crucially in that setting both private and public signals used in the signal extraction problem are exogenous and thus their precision is not affected by agents' behavior.

Less directly related to this 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 common feature of all these works is that agents are uncertain (and care) about the future price of the asset, rather than about its fundamental value, as it is the case in this paper.

2 Learning from prices

I follow Allen, Morris and Shin (2006) in the basic set up of the demand function for traders but modify it to allow for information over the fundamental to cumulate over time. I assume there is an asset available for trade on the market, whose fundamental value is denoted by θ. Such value can be thought of as representing the net (per share) value of the firm. The asset is traded at dates $t = 1, ..., T$ and is liquidated at time $T + 1$ at the value θ. I consider both a static setting, as benchmark, where $T = 1$, and a dynamic setting where information is cumulated over time, with $T > 1$. For simplicity I assume no discount and zero risk free rate, so the present discounted value of the fundamental at each time t is equal to θ:

this assumption serves the purpose to make the equilibrium price stationary, and it avoids complications in the derivations. It does not affect the problem of aggregation of information and the link between the fundamental value and prices that are the focus of the analysis.

I assume that agents are only concerned with the long term (expected) return on their portfolio, defined as the difference between the present value of the fundamental and the current price. Agents are mean variance maximizers, so they attach a penalty to the portfolio variance in their maximization problem.[1] At each time, agents consider the difference between the current price and the expected final liquidation value (the fundamental) and trade accordingly. This could be thought of as representing successive generations of long-term investors, who only trade once and never re-enter the market, while information is passed on from one generation to the next. In particular, at each time t agent i inherits the information of agent i in the time $t-1$ population.

I assume that, at each time t, there is a continuum of agents of unit mass, indexed by $i \in [0,1]$. Agents are homogeneous in all aspects except for the private information they receive. Throughout the paper, I 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 onto the real line. Moreover, for a given t, each function z_t^i is assumed to be continuous and bounded in i.

The problem for a trader i at time t is thus to choose the number of shares (k_t^i) such that

$$\max_{k_t^i} E_t^i W_t^i - \frac{\gamma}{2} Var_t^i \left(W_t^i \right)$$

where γ is the coefficient of risk aversion and the return on the portfolio at time t, W_t^i, is defined as the difference between its final value and what one has to pay for it

$$W_t^i = k_t^i \left(\theta - p_t \right).$$

It follows that the optimal demand for trader i is

$$k_t^i = \frac{E_t^i \theta - p_t}{\gamma \left({}^i\sigma_{w,t}^2 \right)}, \tag{1}$$

where ${}^i\sigma_{w,t}^2$ is agent i's conditional variance for the return on the asset, defined as the difference between its price and the fundamental. With θ unknown and prices observable, such

[1]Basak and Chabakauri (2010) show that dynamic mean-variance portfolio choices are not dynamically consistent. While this is an important issue, mean-variance behavior is still commonly assumed as it leads to analytically tractable demand functions.

conditional variance is equal to the conditional variance for θ, the only source of uncertainty for agents. Agents thus need to form expectations about the unobservable θ and its conditional variance in order to implement their strategy. With homogeneous information, agents would all have the same conditional expectations and variance and therefore demand the same quantity.

In order to close the model with an equilibrium condition, I assume an exogenous and stochastic supply of shares $s = \varepsilon_t$, which follows a normal distribution with zero mean and variance σ_ε^2. This noise term will prevent prices from being fully revealing later on when agents have imperfect information, a common result in the literature on noisy rational expectations equilibria.[2]

Equilibrium with homogeneous information therefore implies

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

where $E_t\theta$ and $\sigma_{w,t}^2$ represent, respectively, agents' common conditional expectations and variance for θ. With no uncertainty then, $E_t\theta = \theta$ and $\sigma_{w,t}^2 = 0$ and thus prices would be constant at the fundamental value, i.e., $p_t = \theta$.

2.1 Information structure and equilibrium: static setting

I now introduce uncertainty and heterogeneous information in the model. In particular I assume that agents don't observe the value of the fundamental directly but receive two signals on it: one, exogenous and private, and the other, prices, endogenous and public.

I start presenting the static setting: with $T = 1$, there is only one period available for agents to receive information and to trade. Nature moves first and draws θ from an improper uniform distribution over the real line $\mathbb{R}$.[3] This gives agents a flat uninformative prior for the fundamental, which simplifies the Bayesian updating from the information they receive. 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). This last component can be interpreted as agents receiving different news because accessing different sources of information, or as the subjective interpretation of the same news. News could be about any

[2] Such assumption can be interpreted to capture variations in the availability of publicly tradable shares (asset float). See, e.g., Mele and Sangiorgi (2015) and Branch and Evans (2011).

[3] Note that the fundamental can be negative, and so do prices. Such feature can be justified by assuming that there is no free disposal of the risky asset. This simplifies the set up, as it avoids having a truncation in the support of agents' beliefs, but does not affect the main results of the paper.

element that is perceived to affect the long term value of an asset. This static setting is equivalent to the one proposed in Hellwig (1980), and so are the results. I propose it here as benchmark on which to build the dynamic setting in the next Section.

The exogenous signal (news) is represented as

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

where v_t^i is an i.i.d. random variable, normally distributed with zero mean and variance σ_v^2.

With signals normally distributed and conditionally independent, Bayesian updating gives a posterior $E\left[\theta \mid x_t^i, p_t\right]$ that is linear in the two signals and equal to

$$\tilde{\theta}_t^i \equiv E\left[\theta \mid x_t^i, p_t\right] = \alpha x_t^i + (1-\alpha)\, p_t, \tag{4}$$

where (see Appendix 6.1 for a derivation) the optimal value for α (denoted α^*) is given by the solution to

$$\alpha^* = \frac{\sigma_{p,t}^2}{\sigma_{p,t}^2 + \sigma_v^2}, \tag{5}$$

with $\sigma_{p,t}^2$ denoting the conditional 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

$$k_t^i = \frac{\alpha^*\left(x_t^i - p_t\right)}{\gamma \sigma_{w,t}^2}, \tag{6}$$

where $\sigma_{w,t}^2$ is the variance of the value of the asset conditional on x_t^i and p_t, common for all agents and given by

$$\begin{aligned}
\sigma_{w,t}^2 &= E\left[\left(\theta - \tilde{\theta}_t^i\right)^2 \mid x_t^i, p_t\right] \\
&= \left(\alpha^*\right)^2 \sigma_v^2 + \left(1-\alpha^*\right)^2 \sigma_{p,t}^2. \qquad (7)
\end{aligned}$$

Aggregate demand is then given by

$$K_t \equiv \int \frac{\alpha^*\left(x_t^i - p_t\right)}{\gamma \sigma_{w,t}^2} di = \frac{\alpha^*\left(\theta - p_t\right)}{\gamma \sigma_{w,t}^2}$$

and prices evolve according to

$$p_t = \theta - \frac{\gamma \sigma_{w,t}^2}{\alpha^*} \varepsilon_t, \tag{8}$$

which implies

$$\sigma^2_{p,t} = \left(\frac{\gamma\sigma^2_{w,t}}{\alpha^*}\right)^2 \sigma^2_\varepsilon. \tag{9}$$

Using (5) together with (7) and (8) I then get the price equation (see Appendix 6.2 for details)

$$p_t = \theta - \gamma\sigma^2_v\varepsilon_t, , \tag{10}$$

which then determines

$$\sigma^2_{p,t} = \gamma^2\left(\sigma^2_v\right)^2\sigma^2_\varepsilon, \tag{11}$$

confirming that prices are normally distributed and conditionally independent from the private signals. As already noted by Angeletos and Werning (2006), it can also be seen from (11) 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, say x_t, $x_t = \theta + v_t$, $v_t \sim N\left(0, \sigma^2_v\right)$, (5) would become

$$\alpha^* = \frac{E_t p^2_t - E_t p_t x_t}{E_t p^2_t + E_t x^2_t - 2E_t p_t x_t} = \frac{\sigma^2_{p,t} - \sigma^2_v}{\sigma^2_{p,t} + \sigma^2_v - 2\sigma^2_v} = 1. \tag{12}$$

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: the optimal value for α would thus be equal to one. In other words, in order for prices to have any informational content above and beyond what is provided by the idiosyncratic signal, it must be that the aggregation process that generates prices averages out some noise.

With private signals, instead, the optimal weight on private information is

$$\alpha^* = \frac{\sigma^2_{p,t}}{\sigma^2_{p,t} + \sigma^2_v} = \frac{\gamma^2\left(\sigma^2_v\right)^2\sigma^2_\varepsilon}{\sigma^2_v + \gamma^2\left(\sigma^2_v\right)^2\sigma^2_\varepsilon} = \frac{\gamma^2\sigma^2_v\sigma^2_\varepsilon}{1 + \gamma^2\sigma^2_v\sigma^2_\varepsilon}, \tag{13}$$

which shows that the optimal weight on private information depends positively on the coefficient of risk aversion, the variance of the noise in the private information and the variance of the supply.

It is instructive to compute some limiting cases:

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

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.

Furthermore, if the variance of the idiosyncratic noise in the exogenous signal goes to zero, α^* goes to zero, as can be easily seen from (13): 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:[4] this is due to the fact that the volatility of prices originates from the volatility of the supply and from the volatility of the demand (multiplicatively): the latter arises from uncertainty about the fundamental and in particular is quadratic in the (conditional) variance of the private signal.

It is interesting to consider the condition under which agents put more weight on their private information than on prices, i.e., $\alpha^* > 1 - \alpha^*$, which requires $\gamma^2 \sigma_v^2 \sigma_\varepsilon^2 > 1$. When $\sigma_v^2 \sigma_\varepsilon^2$ is small, thus, agents pay little attention to their private signal: this happens when either the variance of the supply is low (and thus aggregate noise is low and prices are more informative), or when private information is very precise (because, as explained above, this enhances the informativeness of prices).

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 (13). 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.

Remark 1 *In the static setting presented in this Section, the optimal Bayesian weight on private information is given by (13) and prices are defined by (10).*

It is clear, by comparing (10) 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.

[4]The variance of prices is in fact quadratic in σ_v^2: see (11).

The limiting results on the optimal use of information just discussed have implications also for traders' demand. In fact, rewriting (6) as

$$k_t^i = a^* \left(x_t^i - p_t \right) ,$$

with $a^* = \frac{\alpha^*}{\gamma \sigma_{w,t}^2} = \left(\gamma \sigma_v^2 \right)^{-1}$, we can see that $\sigma_v^2 \to 0$ implies $k_t^i \to \infty$: as private information becomes infinitely precise, the demand's response to deviations of prices from the fundamental grows without bound as agents try to exploit any (risk free) arbitrage available.

Note instead that as $\sigma_\varepsilon^2 \to 0$, prices become infinitely precise through the aggregation of private information but individual demands do not depend on the volatility of the supply: agents demand the same quantity, irrespective of σ_ε^2. As the volatility of supply decreases towards zero, in fact, $\alpha \to 0$, which makes agents respond less and less to deviations of their private signal from prices; at the same time, though, the conditional variance of the return on the asset also decreases, which increases the demand of risk averse agents, with the result that the amount demanded is constant with respect to σ_ε^2.

2.2 Dynamic setting

In the previous section I have considered a static framework where information is only received once and trading happens only in one period. The purpose of this work, though, is to analyze the effect of information accumulation and learning on the ability of prices to reveal information about fundamentals. I extend therefore the framework to investigate such issues.

I set $T > 1$, and arbitrarily large. Nature draws the fundamental θ at time $t = 0$ and agents at every period t, $1 \leq t \leq T$ 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). The supply of shares is still exogenous and stochastic, with $\varepsilon_t \sim N\left(0, \sigma_\varepsilon^2\right)$ an i.i.d. process.[5]

Mean and precision of the posterior of θ at each time t, conditional on the history of x_t^i,

[5] This assumption is used, for example, in Allen, Morris and Shin (2006). After providing an interpretation for such modelling choice, they write: "Clearly, the interpretation given above is somewhat contrived, but we advance it merely as a modeling device that serves the purpose of preventing prices being fully revealing, and preserving the independence of the supply shocks over time, so as to aid tractability of the analysis." The same motivation for such modelling choice applies here.

can be written, respectively, as

$$\bar{x}_t^i = \frac{\beta_{t-1}}{\beta_t}\bar{x}_{t-1}^i + \frac{\sigma_v^{-2}}{\beta_t}x_t^i = \frac{1}{t}\sum_{z=1}^{t} x_z^i \tag{14}$$

$$\beta_t = \beta_{t-1} + \sigma_v^{-2} = \sum_{z=1}^{t}\sigma_v^{-2} = \frac{t}{\sigma_v^2}, \tag{15}$$

with $\beta_0 = 0$ and $\bar{x}_0^i = 0$ (that is, $\bar{x}_1^i = x_1^i$).

In terms of the public endogenous signal represented by prices, the mean and precision of the posterior of θ at time t, conditional on the history of p_t, are given, respectively, by

$$\bar{p}_t = \frac{\omega_{t-1}}{\omega_t}\bar{p}_{t-1} + \frac{\sigma_{p,t}^{-2}}{\omega_t}p_t = \frac{\sum_{z=1}^{t}\sigma_{p,z}^{-2}p_z}{\sum_{z=1}^{t}\sigma_{p,z}^{-2}} \tag{16}$$

$$\omega_t = \omega_{t-1} + \sigma_{p,t}^{-2} = \sum_{z=1}^{t}\sigma_{p,z}^{-2}, \tag{17}$$

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 $E\left[\theta \mid \bar{x}_t^i, \bar{p}_t\right]$ is given by

$$E\left[\theta \mid \bar{x}_t^i, \bar{p}_t\right] = \alpha_t \bar{x}_t^i + (1-\alpha_t)\,\bar{p}_t, \tag{18}$$

with α_t determined by the relative precision of the two posteriors at each time t and given by

$$\alpha_t = \frac{\beta_t}{\beta_t + \omega_t}. \tag{19}$$

Using the demand equation

$$k_t^i = \frac{E_t^i\theta - p_t}{\gamma\sigma_{w,t}^2} \tag{20}$$

with

$$E_t^i\theta = E\left[\theta \mid \bar{x}_t^i, \bar{p}_t\right]$$

and

$$\sigma_{w,t}^2 = \alpha_t^2\beta_t^{-1} + (1-\alpha_t)^2\,\omega_t^{-1} = \frac{1}{\beta_t + \omega_t}, \tag{21}$$

aggregating and imposing demand equal supply, leads to

$$p_t = (\alpha_t\theta + (1-\alpha_t)\,\bar{p}_t) - \gamma\sigma_{w,t}^2\varepsilon_t. \tag{22}$$

Solving for p_t, using (16), (17) and (19), gives the price equation

$$p_t = \frac{\beta_t}{\beta_t + \omega_{t-1}}\theta + \frac{\omega_{t-1}}{\beta_t + \omega_{t-1}}\bar{p}_{t-1} - \frac{\gamma\varepsilon_t}{\beta_t + \omega_{t-1}}. \tag{23}$$

Note that, with $\omega_0 = 0$ and $\bar{p}_0 = 0$, at time $t = 1$, p_1 from (23) reduces to (10), as in the static case, since there is no previous information to exploit.

From (23), it then follows that the conditional variance of prices is

$$\sigma^2_{p,t} = \frac{\gamma^2\sigma^2_\varepsilon}{(\beta_t + \omega_{t-1})^2}. \tag{24}$$

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

2.3 Limiting dynamics

The main result of this Section is that, over time, the relative weight put on private information decreases towards zero and agents rely only on prices in their Bayesian signal extraction problem. I state the result formally in the following Proposition

Proposition 2 *In the dynamic setting presented in Section (2.2), where the optimal Bayesian weight on private information is given by (19) and prices evolve according to (22), in the limit, α_t converges to 0 and p_t converges to the fundamental value θ.*

Proof. The proof consists in deriving the limiting outcomes of the system as $t \to \infty$ (which clearly requires that also $T \to \infty$). Starting with the exogenous signal, by the law of large numbers

$$\operatorname*{plim}_{t\to\infty} \bar{x}^i_t = \theta \tag{25}$$

and clearly

$$\lim_{t\to\infty} \beta_t = \frac{t}{\sigma^2_v} = \infty. \tag{26}$$

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 from (21) that

$$\lim_{t\to\infty} \sigma^2_{w,t} = 0 \tag{27}$$

and from (24)

$$\lim_{t\to\infty} \sigma^2_{p,t} = 0. \tag{28}$$

This last result also implies, from (17), that

$$\lim_{t\to\infty} \omega_t = \infty. \tag{29}$$

Finally, since

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

it follows that

$$\alpha_t = \frac{\beta_t}{\beta_t + \omega_t} = \frac{1}{1 + \frac{\omega_{t-1}}{\beta_t} + \left(\gamma^2\sigma_\varepsilon^2\right)^{-1}\left(\beta_t + \frac{\omega_{t-1}^2}{\beta_t} + 2\omega_{t-1}\right)}.$$

Combining then with results (26) and (29),

$$\lim_{t\to\infty} \alpha_t = 0. \tag{30}$$

Looking then at prices, starting from (22) and noting results (27) and (30)

$$\operatorname*{plim}_{t\to\infty} p_t = \operatorname*{plim}_{t\to\infty}\left[\alpha_t\theta + (1-\alpha_t)\,\bar{p}_t - \gamma\sigma_{w,t}^2\varepsilon_t\right] = \operatorname*{plim}_{t\to\infty}\bar{p}_t. \tag{31}$$

Given that $\bar{p}_t$ is a weighted average of all p_z, $1 \le z \le t$, each one centered around θ and with decreasing variance (and that, from (16), the weight on each p_z is inversely proportional to its variance)

$$\operatorname*{plim}_{t\to\infty}\bar{p}_t = \theta$$

and thus p_t converges over time to the fundamental θ. ■

The key element to understand result (30) is to note that while both β_t and ω_t tend to ∞ as $t \to \infty$, β_t grows linearly while ω_t grows 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. In the limit, given enough time to trade, the price converges to the fundamental value, a result consistent with Vives (1995a). At the same time, prices become the only source of information used by agents in predicting the fundamental value.

Because agents try to minimize the variance of the return on 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 the fundamental value changes at some point in time, such change might not get factored properly into prices, thus leading to a divergence between the two. In order to investigate this issue properly, I modify the framework to allow for changes in the fundamental value. As agents are aware of this possibility, they will account for it in their signal extraction problem.

2.4 Allowing for changes in the fundamental

I have considered so far the case where the environment is stationary and agents accordingly use past information to improve their estimates of the fundamental value. 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 address this issue, I consider now a framework where the fundamental is allowed to change over time and agents are aware of this possibility. In particular, given that the fundamental value of an asset is not observables, agents will never be sure whether a change has actually taken place at any specific time, and only entertain subjective probabilities on such events. I assume that such subjective probability is the same as the true (ex-ante) unconditional probability that the fundamental changes at every period. I therefore do not allow agents to update such prior based on ex post evidence from the signals. One could instead think of allowing agents to revise their prior based on how far the new observation for the exogenous signal has fallen from their most recent estimate of the fundamental: the further that is, the more likely it is that such discrepancy comes from a change in the fundamental rather than from noise in the private signal. I abstain here from such complication and assume instead that the subjective probability of a change in fundamental is equal to its ex ante unconditional probability: in other words, there is no learning from agents along such dimension. I will discuss later on, in light also of results from simulations, how results would likely differ if I were to allow agents to learn also along such dimension.

As before, Nature draws the fundamental from an improper uniform distribution over $\mathbb{R}$ at time $t = 0$. Nature, though, can now also re-draw, with some fixed and known (to agents) probability $0 \leq \pi \leq 1$, a new value for the fundamental, from the same improper distribution, at the beginning of each time $t > 1$.

I first define

$$\bar{x}^i_{t,j\le t} = \frac{1}{t-j+1}\sum_{z=j}^{t} x^i_z,$$

the posterior of θ at time t for agent i if a change in the fundamental had occurred at (the beginning of) time $j \le t$ and agents knew it. This is simply the mean of the sample of relevant observations for the exogenous signal, since the change in fundamental took place. Given that agents don't know if and when a change in the fundamental took place, the best predictor for the fundamental at time t, conditional on exogenous signals only, is then given by

$$\tilde{x}^i_t = \sum_{j=1}^{t} a^t_j \bar{x}^i_{t,j\le t}$$

where

$$\begin{aligned} a^t_1 &= (1-\pi)^{t-1} \\ a^t_j &= (1-\pi)^{t-j}\pi, t \ge j > 1. \end{aligned}$$

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

Note that

$$\sum_{j=1}^{t} a^t_j = 1.$$

It is then possible to rewrite $\tilde{x}^i_t$ as a weighted sum of current and past values of x^i_t

$$\tilde{x}^i_t = \sum_{j=1}^{t} h^t_j x^i_j, \tag{32}$$

where

$$h^t_1 = \frac{(1-\pi)^{t-1}}{t} \tag{33}$$

$$h^t_j = h^t_{j-1} + \frac{(1-\pi)^{t-j}\pi}{t-j+1}, t \ge j > 1. \tag{34}$$

Again, $\sum_{j=1}^{t} h^t_j = 1$. Clearly if $\pi = 1$, $h^t_j = 0$ for $j < t$ and $h^t_j = 1$ for $j = t$: only the last observation matters. If instead $\pi = 0$, then all observations receive the same weight $1/t$.

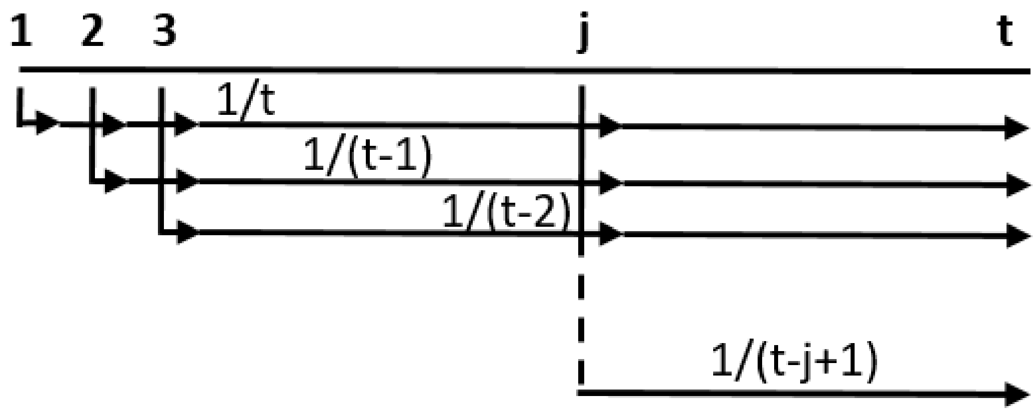

Figure 1: Weighting structure of signals.

To understand better the weighting structure, I propose Fig (1). Observation x_1^i, $\forall i$, is relevant for inference about the fundamental at time t only if Nature never re-drew over the whole sample period from 1 to t, which happened with probability $(1-\pi)^{t-1}$: in such case each observation in that sample should be weighted equally, with weight $1/t$. Observation x_2^i, $\forall i$, is relevant if Nature never re-drew (again, with weight $1/t$), which happened with probability $(1-\pi)^{t-1}$, or if it re-drew at the beginning of period 2 and never after (and in this case, with weight $\frac{1}{t-1}$), which happened with probability $(1-\pi)^{t-2}\pi$. And so on.

I can then define

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

which, since all the x_j^i are i.i.d. over time, reduces to

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

Again, if $\pi = 0$ then $\tilde{\beta}_t = \beta_t = t\sigma_v^{-2}$ as in the dynamic case with constant fundamental. If $\pi = 1$ then $\tilde{\beta}_t = \sigma_v^{-2}$ as in the static case.

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

$$\bar{p}_{t,j\le t} = \frac{\sum_{z=j}^{t}\sigma_{p,z}^{-2}p_z}{\sum_{z=j}^{t}\sigma_{p,z}^{-2}},$$

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

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

or

$$k_j^t = \sigma_{p,j}^{-2}\left[\frac{(1-\pi)^{t-1}}{\sum_{z=1}^{t}\sigma_{p,z}^{-2}} + \pi\sum_{z=2}^{j}\frac{(1-\pi)^{t-z}}{\sum_{m=z}^{t}\sigma_{p,m}^{-2}}\right] \text{ for } t \geq j > 1, \tag{35}$$

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

Note that for $\pi = 0$, $k_j^t = \frac{\sigma_{p,j}^{-2}}{\sum_{z=1}^{t}\sigma_{p,z}^{-2}}$, and for $\pi = 1$, $k_{j\neq t}^t = 0$ and $k_t^t = 1$. That is, for $\pi = 0$ 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. For $\pi = 1$, instead, the framework converges to the static setting seen in Section 2.1, since every period agents discount completely the past.[6] The logic behind the weighting structure is the same as the one seen in Fig (1) for the exogenous signal, with the only difference now that price observations are also weighted by their relative variance, since $\sigma_{p,t}^{-2}$ differs from period to period.

[6]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 $\pi \to 1$.

I can then define

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

which, since prices are conditionally independent, reduces to

$$\tilde{\omega}_t = \left(\sum_{j=1}^{t} \left(k_j^t\right)^2 \sigma_{p,j}^2\right)^{-1}. \tag{36}$$

It can be seen that, for $\pi = 0$, $\tilde{\omega}_t = \omega_t$ from (17) and for $\pi = 1$, $\tilde{\omega}_t = \sigma_{p,t}^{-2}$ as in the static case.

Consider then the pricing equation

$$p_t = E_t\theta - \gamma\sigma_{w,t}^2\varepsilon_t, \tag{37}$$

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

$$\tilde{\alpha}_t = \frac{\tilde{\beta}_t}{\tilde{\beta}_t + \tilde{\omega}_t}. \tag{38}$$

Since $\tilde{p}_t$ includes current prices, solving (37) for p_t gives

$$\begin{aligned} p_t &= \tilde{\alpha}_t \int \tilde{x}_t^i di + (1-\tilde{\alpha}_t)\,\tilde{p}_t - \gamma\sigma_{w,t}^2\varepsilon_t \\ &= \frac{\tilde{\alpha}_t}{1-(1-\tilde{\alpha}_t)\,k_t^t}\int \tilde{x}_t^i di + \frac{(1-\tilde{\alpha}_t)}{1-(1-\tilde{\alpha}_t)\,k_t^t}\sum_{j=1}^{t-1} k_j^t p_j - \frac{\gamma\sigma_{w,t}^2\varepsilon_t}{1-(1-\tilde{\alpha}_t)\,k_t^t}, \end{aligned} \tag{39}$$

where

$$\begin{aligned} \int \tilde{x}_t^i di &= \int \left(\sum_{j=1}^{t} h_j^t x_j^i\right) di = \sum_{j=1}^{t}\left(h_j^t \int x_j^i di\right) \\ &= \sum_{j=1}^{t} h_j^t \theta_j. \end{aligned}$$

Note that, since the fundamental can now change, I have denoted θ_j as the fundamental prevailing at time j.

I thus obtain the price equation

$$p_t = \frac{\tilde{\alpha}_t}{1-(1-\tilde{\alpha}_t)\,k_t^t}\sum_{j=1}^{t} h_j^t \theta_j + \frac{(1-\tilde{\alpha}_t)}{1-(1-\tilde{\alpha}_t)\,k_t^t}\sum_{j=1}^{t-1} k_j^t p_j - \frac{\gamma \sigma_{w,t}^2 \varepsilon_t}{1-(1-\tilde{\alpha}_t)\,k_t^t}. \tag{40}$$

The conditional variance of prices, $\sigma_{p,t}^2$, is endogenous and only implicitly defined by the above equations. From (40), I first derive

$$\sigma_{p,t}^2 = \frac{\gamma^2 \sigma_\varepsilon^2}{\left[\tilde{\beta}_t + \tilde{\omega}_t\left(1-k_t^t\right)\right]^2}, \tag{41}$$

where both $\tilde{\omega}_t$ and k_t^t depend on $\sigma_{p,t}^2$. I then use (36) and (35) (with $j = t$) to substitute k_t^t and $\tilde{\omega}_t$ in (41) and solve numerically for $\sigma_{p,t}^2$. Note that, for $\pi = 0$, $\sigma_{p,t}^2$ from (41) reduces to $\sigma_{p,t}^2$ from (24). As (41) is non-linear and could admit more than one real root, I search locally for a solution for $\sigma_{p,t}^2$ near $\sigma_{p,t-1}^2$, starting with $\sigma_{p,0}^2$ as the equilibrium value for the dynamic case derived in (11), which is indeed the correct value for the first period, where there is no previous information to exploit. The idea is that new information doesn't create big jumps in the system as it gets incorporated into prices.

In order to understand the behavior of this system, I now turn to simulations.

2.4.1 Simulations

I first investigate the long run behavior of $\tilde{\alpha}_t$ for all possible values of π. To this end, I let the system run until convergence of $\tilde{\alpha}_t$ and thus compute numerically the function $\alpha(\pi)$, representing the long run value for $\tilde{\alpha}_t$ for all possible values of $\pi \in [0, 1]$. I repeat this exercise for different values of the risk aversion parameter γ. I report results for two commonly used values of risk aversion, $\gamma = .75$ and $\gamma = 1.5$, respectively in Fig. (2) and Fig. (3).[7] The red circle corresponds to α^*, the optimal value for the static case, where the system converges to for $\pi = 1$, as Nature re-draws the fundamental every period with probability 1.

As it can be seen from the pictures, for relatively low values of π (roughly below .6 for $\gamma = .75$ and below .4 for $\gamma = 1.5$) the optimal weight on private information in the long run approaches zero. For example, with $\gamma = .75$, for $\pi = .6$ I find that $\alpha(.6) = 1.9767e - 016$. This is due to the fact that, while the precision of the private information remains bounded for any value of π, the precision of the public signal instead increases without bound for relatively low values of $\pi > 0$ (see Fig. 4, where I report the convergence values for the

[7] The other two relevant parameters are set to $\sigma_v^2 = \sigma_\varepsilon^2 = 1$.

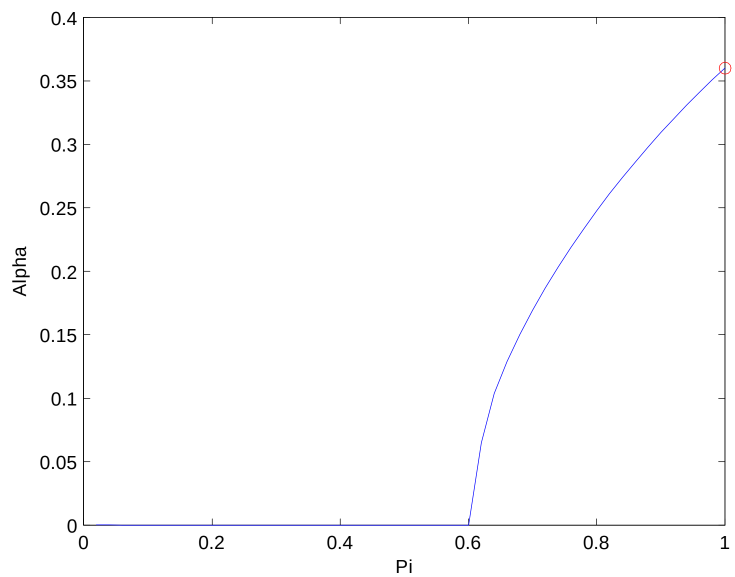

Figure 2: Long run values for $\tilde{\alpha}_t$. $\gamma = .75$.

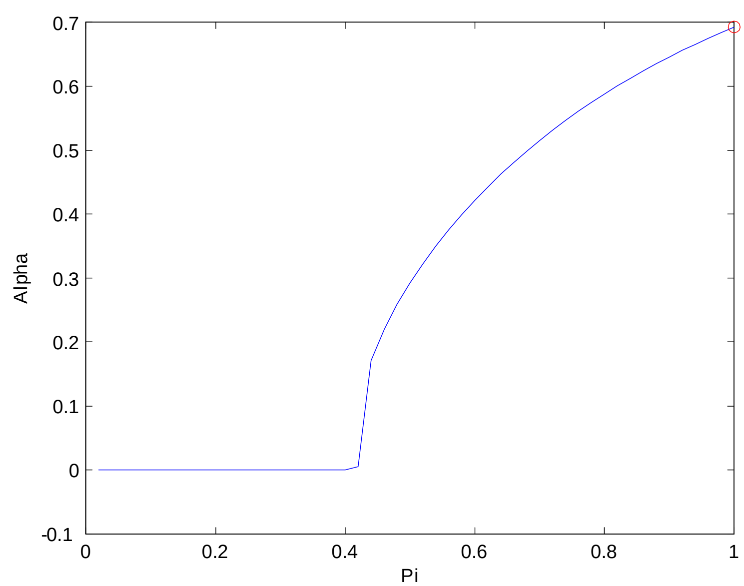

Figure 3: Long run values for $\tilde{\alpha}_t$. $\gamma = 1.5$.

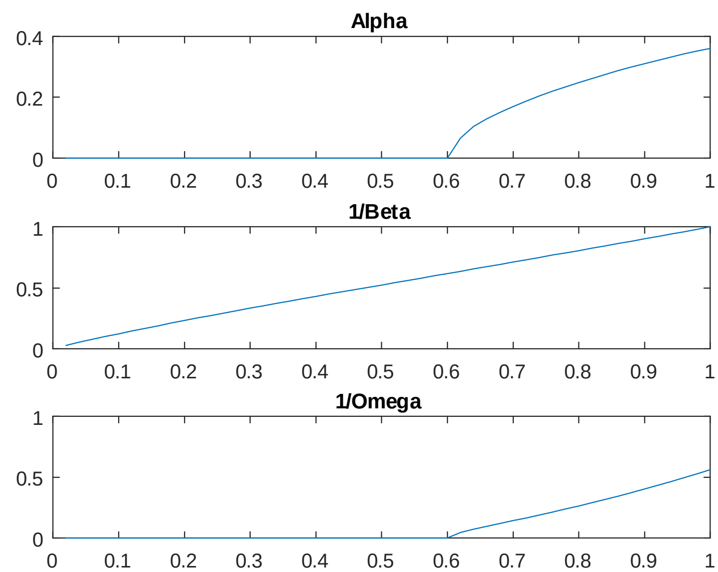

Figure 4: Precision of public and private signal and relative weight on the two signals.

inverse of $\tilde{\beta}_t$ and $\tilde{\omega}_t$ against values of π): as the variance of prices is endogenous, for low values of π a self-reinforcing mechanism is at play by which an increase in the precision of the public signal leads to an increased weight on public information, which in turns leads to a further increase in the precision of prices. If π is large enough, instead, the uncertainty about whether past observations are relevant for current inference is high enough to prevent the precision of the public signal, now heavily skewed towards most recent observations, from increasing without bound.

One important issue is how fast $\tilde{\alpha}_t$ converges to its long run value, in particular in relation to the expected frequency of changes in the fundamental. For example, with $\pi = .1$ it could be expected on average to have a change in the fundamental every 10 periods: would $\tilde{\alpha}_t$ have already converged to its equilibrium value $\alpha\,(.1)$ at that point? To try and answer this question I show the simulated path for $\tilde{\alpha}_t$ for $\pi = \{.1, .4, .7, .9\}$ where in each case I initialize the algorithm at $\tilde{\alpha}_t = \alpha^*$, the optimal level for the static setting. I report in Fig (5) results for $\gamma = .75; \sigma_v^2 = \sigma_\varepsilon^2 = 1$.

It is possible to see that, for $\pi = .1$, $\tilde{\alpha}_t$ converges to zero after about 7 periods: given that $\pi = .1$ implies that Nature redraws on average once every 10 periods, this means that changes in the fundamental are likely to be missed out and not passed on to prices. On the

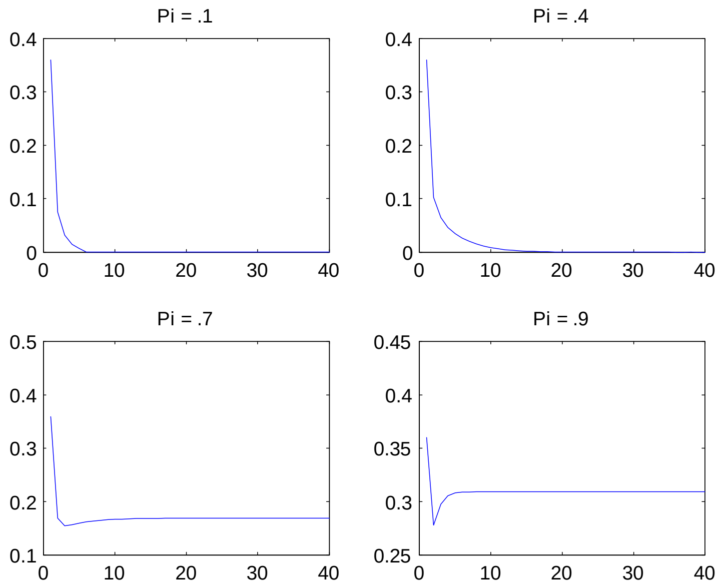

Figure 5: Convergence paths for $\tilde{\alpha}_t$, for $\pi = .1, .4, .7, .9$. $\gamma = .75$.

other end, for $\pi = .9$, $\tilde{\alpha}_t$ converges after about 5 periods to a positive value and never drops down to zero: any change in the fundamental is going to be passed on to some extend to prices.

The main conclusion from this Section is that for relatively low values of π the weigh on the private signal quickly converges to zero, thus preventing any change in the fundamental to be passed on to prices, while if π is high enough, the possibility of a change in the fundamental having occurred in the past prevents the precision of the public signal from increasing without bounds and thus maintains a positive weight on private information in the signal extraction problem.

3 Adaptive learning

Bayesian learning assumes agents know the relative precision of the different signals they receive, and use such information optimally. I 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. I take here the view that agents know their own preferences and know how to solve relatively simple problems of signal extraction. This means that I assume they can work out the way the optimal weight on private versus public information 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 two signals.

Agents will thus need to learn about mean and variance of the two signals they receive. While prices are publicly observable, the exogenous signal is instead idiosyncratic: it follows that, under adaptive learning, agents will hold homogeneous beliefs about prices' moments (assuming common initial beliefs and updating algorithms) but heterogeneous beliefs about the statistical properties of the exogenous signal.

I first provide convergence results under decreasing gain in a stationary environment, with constant fundamental, and then consider the case of constant gain, which will allow me 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 is therefore better suited to analyze adaptive learning in a setting where changes in the fundamental can take place.

In order to implement their optimal demand (20), agents need to learn about $\bar{x}_t^i, \bar{p}_t, \beta_t, \omega_t$ and $\sigma_{w,t}^2$. At each time t, each agent i observes the whole history from 1 to t of his own private signal and of prices, i.e., he observes $\{x_z^i\}_{z=1}^t$ and $\{p_z\}_{z=1}^t$. An agent i at time t is thus defined by the history of signals he receives and his beliefs system $\hat{x}_t^i, \hat{\beta}_t^i, \hat{p}_t, \hat{\omega}_t, s_{w,t}^i$, which are the learning counterparts to $\bar{x}_t^i, \beta_t, \bar{p}_t, \omega_t, \sigma_{w,t}^2$, plus $s_{x,t}^i$ and $s_{p,t}$, which are the estimated second raw moments for the two signals. These last two measures are needed in the adaptive learning environment in order for agents to be able to estimate the second central moments of the signals (see equations (44)-(45) later on). Individual demand is then given by

$$k_t^i = \frac{\hat{\alpha}_t^i \hat{x}_t^i + \left(1 - \hat{\alpha}_t^i\right) \hat{p}_t - p_t}{\gamma s_{w,t}^i}, \tag{42}$$

where

$$\hat{\alpha}_t^i = \frac{\hat{\beta}_t^i}{\hat{\beta}_t^i + \hat{\omega}_t}, \tag{43}$$

with

$$\hat{\beta}_t^i = \hat{\beta}_{t-1}^i + \left(s_{x,t}^i - \left(\hat{x}_t^i\right)^2\right)^{-1} \tag{44}$$

$$\hat{\omega}_t = \hat{\omega}_{t-1} + \left(s_{p,t} - \left(\hat{p}_t\right)^2\right)^{-1} \tag{45}$$

and

$$s_{w,t}^{i} = \left(\hat{\alpha}_{t}^{i}\right)^{2}\left(\hat{\beta}_{t}^{i}\right)^{-1} + \left(1-\hat{\alpha}_{t}^{i}\right)^{2}\hat{\omega}_{t}^{-1} = \frac{1}{\hat{\beta}_{t}^{i}+\hat{\omega}_{t}}. \tag{46}$$

Equations (44)-(45) 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

$$K_{t} = \int \frac{\left(\hat{\alpha}_{t}^{i}\hat{x}_{t}^{i} + \left(1-\hat{\alpha}_{t}^{i}\right)\hat{p}_{t}\right) - p_{t}}{\gamma s_{w,t}^{i}} di. \tag{47}$$

From the market clearing condition, it follows that prices evolve according to

$$p_{t} = \frac{\left(\int \frac{\hat{\alpha}_{t}^{i}}{s_{w,t}^{i}} di\right)\theta + \left(\int \frac{\left(1-\hat{\alpha}_{t}^{i}\right)}{s_{w,t}^{i}} di\right)\hat{p}_{t} - \gamma\varepsilon_{t}}{\int \frac{1}{s_{w,t}^{i}} di}, \tag{48}$$

or, since $\hat{p}_{t-1}$ is common across agents, x_{t}^{i} 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 $\int x_{t}^{i} di = \theta$,

$$p_{t} = \frac{\int \hat{\beta}_{t}^{i} di}{\int \hat{\beta}_{t}^{i} di + \hat{\omega}_{t-1}}\theta + \frac{\hat{\omega}_{t-1}}{\int \hat{\beta}_{t}^{i} di + \hat{\omega}_{t-1}}\hat{p}_{t-1} - \frac{\gamma\varepsilon_{t}}{\int \hat{\beta}_{t}^{i} di + \hat{\omega}_{t-1}}, \tag{49}$$

which is the adaptive learning counterpart of (23).

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

$$s_{x,t}^{i} = s_{x,t-1}^{i} + g_{t}\left[\left(x_{t}^{i}\right)^{2} - s_{x,t-1}^{i}\right] \tag{50}$$

$$s_{p,t} = s_{p,t-1} + g_{t}\left[\left(p_{t}\right)^{2} - s_{p,t-1}\right], \tag{51}$$

and the means of the two signals according to

$$\hat{x}_{t}^{i} = \hat{x}_{t-1}^{i} + g_{t}\left[x_{t}^{i} - \hat{x}_{t-1}^{i}\right] \tag{52}$$

$$\hat{p}_{t} = \hat{p}_{t-1} + g_{t}\left[p_{t} - \hat{p}_{t-1}\right]. \tag{53}$$

Equations (49)-(53), together with (44)-(45), 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

I provide 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), I can represent the evolution, in notional time τ, of estimated variables through ordinary differential equations (ODEs). Fixed points of these ODEs represent possible limiting points of the original stochastic recursive algorithm and thus allow one to analyze the long run behavior of such algorithm. Results from this section require that $T \to \infty$ in the dynamics setting, so that asymptotic results can apply.

The analysis of the learning equation (50) is rather straightforward, 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

$$\frac{ds_x^i}{d\tau} := \lim_{t\to\infty} E\left(\left(x_t^i\right)^2 - s_{x,t-1}^i\right) = \theta^2 + \sigma_v^2 - s_x^i. \tag{54}$$

Clearly this ODE has a unique fixed point $s_x^i = \theta^2 + \sigma_v^2$, $\forall i$. Moreover, this fixed point is stable since the derivative of the ODE w.r.t. s_x^i is equal to -1.

Similarly for (52), which gives rise to the ODE

$$\frac{d\hat{x}^i}{d\tau} := \lim_{t\to\infty} E\left(x_t^i - \hat{x}_{t-1}\right) = \theta - \hat{x}^i, \tag{55}$$

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

The two results above imply that, $\forall i$,

$$\lim_{t\to\infty}\left(s_{x,t}^i - \left(\hat{x}_t^i\right)^2\right) = \sigma_v^2 \tag{56}$$

and thus, from (44),

$$\lim_{t\to\infty}\hat{\beta}_t^i = \infty. \tag{57}$$

In terms of the endogenous variables, starting from $\hat{p}_t$, since ε_t is a zero mean i.i.d. process independent from the other variables, the relevant ODE derived from (53) is

$$\begin{aligned}
\frac{d\hat{p}}{d\tau} \quad : &= \lim_{t\to\infty} E\left(\frac{\int \hat{\beta}_t^i di}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}}\theta + \frac{\hat{\omega}_{t-1}}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}}\hat{p}_{t-1} - \frac{\gamma\varepsilon_t}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}} - \hat{p}_{t-1}\right) \\
&= \lim_{t\to\infty} E\left(\frac{\int \hat{\beta}_t^i di}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}}\theta + \left(\frac{\hat{\omega}_{t-1}}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}} - 1\right)\hat{p}_{t-1}\right) \\
&= \lim_{t\to\infty} \frac{\int \hat{\beta}_t^i di}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}}\left(\theta - \hat{p}\right) \qquad (58)
\end{aligned}$$

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, I will show, 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

$$\begin{aligned}
\frac{ds_p}{d\tau} \quad : &= \lim_{t\to\infty} E\left(\begin{array}{c}\left(\frac{\int \hat{\beta}_t^i di}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}}\theta + \frac{\hat{\omega}_{t-1}}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}}\hat{p}_{t-1} - \frac{\gamma\varepsilon_t}{\left(\int \hat{\beta}_{t-1}^i di + \hat{\omega}_{t-1}\right)}\right)^2 + \ldots \\ \ldots - s_{p,t-1}\end{array}\right) \\
&= \lim_{t\to\infty}\left(\begin{array}{c}\left(\frac{\int \hat{\beta}_t^i di}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}}\right)^2\theta^2 + \left(\frac{\hat{\omega}_{t-1}}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}}\right)^2\left(\hat{p}_{t-1}\right)^2 + \frac{\gamma^2\sigma_\varepsilon^2}{\left(\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}\right)^2} + \ldots \\ \ldots + 2\frac{\int \hat{\beta}_t^i di}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}}\frac{\hat{\omega}_{t-1}}{\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}}\theta\hat{p}_{t-1} - s_{p,t-1}\end{array}\right) \qquad (59)
\end{aligned}$$

and, using result (57) together with the fact that, by definition, $\hat{\omega}_t > 0$,

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

which then implies

$$\begin{aligned}
\frac{ds_p}{d\tau} &= \lim_{t\to\infty}\left(\frac{\left(\int \hat{\beta}_t^i di\right)^2 + \left(\hat{\omega}_{t-1}\right)^2 + 2\left(\int \hat{\beta}_t^i di\right)\hat{\omega}_{t-1}}{\left(\int \hat{\beta}_t^i di + \hat{\omega}_{t-1}\right)^2}\theta^2 - s_{p,t-1}\right) \\
&= \theta^2 - s_p. \qquad (60)
\end{aligned}$$

The fixed point of this ODE is thus $s_p = \theta^2$, which implies that the variance of prices converges to zero since the second raw moment converges to the squared mean.

From (45), using the results that, in equilibrium, $s_p = \theta^2$ and $\hat{p} = \theta$,

$$\lim_{t\to\infty} \hat{\omega}_t = \infty$$

and thus

$$\lim_{t\to\infty} s^i_{w,t} = \lim_{t\to\infty} \frac{1}{\hat{\beta}^i_t + \hat{\omega}_t} = 0.$$

Finally, convergence in all the learning algorithms ensures that α^i_t also converges to its long run equilibrium value, i.e.,

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

I summarize results from this Section in the following Proposition:

Proposition 3 *Consider the system of ODEs (54), (55), (58), (60), 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* $\hat{\alpha}^i_t$ *converges to* 0 *as* t *grows without bound and the system converges asymptotically to the Bayesian equilibrium defined in Sections (2.2)-(2.3).*

The economy under combined Bayesian and adaptive learning thus converges asymptotically to the same equilibrium derived under Bayesian learning only, where agents had full information about the statistical properties of their signals. This means that, as agents learn the relevant moments for the signals, the precision of those signals still increases over time without bound: in particular, prices still become a more precise predictor for the fundamental even when agents's beliefs about the moments of the signals evolve over time in light of new information generated within the system itself.

Establishing convergence of the adaptive learning algorithm under decreasing gain is a necessary step in order to consider the evolution of prices under real time learning with constant gain, which has the property of discounting past information more heavily and has been used in the literature to model adaptive learning in situations where agents might fear structural breaks taking place in the parameters they are estimating. If the learning algorithms were not stable under decreasing gain, in fact, they would not be stable under constant gain either, and would therefore be unsuitable to model the evolution of beliefs in environments where variables are not observed to diverge.

3.2 Constant gain

Having established convergence to the Bayesian equilibrium under adaptive learning with decreasing gain, I now consider the dynamics of the system under constant gain adaptive learning. Constant gain learning discounts past observations exponentially and thus allows new information to play a larger role in the determination of beliefs. It has been proposed as a practical way to allow learning algorithms to incorporate changes in the estimated parameters and seems thus suitable to capture the need to allow for possible time variation in the fundamental value of the asset in this model.

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)). By allowing beliefs to change endogenously according to evidence, these models account for the co-movements of expectations and economic outcomes and are able to capture the important self-referentiality element of this joint determination while at the same time allowing for delays in the adjustments.

In order to understand the impact of constant gain learning in my model, I simulate the system composed by equations (49) to (53) plus (44)-(45) setting $g_t \equiv g$, a small constant. Though there is no direct evidence of the appropriate value for such parameter, Berardi and Galimberti (2017) 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. I will take $g = 0.025$, a fairly common value used in the macroeconomic applied literature, as benchmark in my simulations, and compare results with a lower gain. I also set $\sigma_v^2 = \sigma_\varepsilon^2 = 1; \gamma = .75$ as before.

In Fig. (6) I show $\hat{\omega}_t$ and average (across agents) $\hat{\beta}_t^i$ and $\hat{\alpha}_t^i$ for $g = .025$. It can be seen that while average $\hat{\beta}_t^i$ grows linearly, $\hat{\omega}_t$ grows exponentially and $\hat{\alpha}_t^i$ converges towards zero: the shift in weight towards more recent observations induced by the constant gain is not sufficient to prevent the weight on the public signal to decrease towards zero. Note though that such decrease is rather slow, as average $\hat{\alpha}_t^i$ is still positive at $t = 200$. This means that changes in fundamental would be factored into prices for quite a long time.

To investigate the impact of the learning gain on the evolution of the relative weight on private versus public information, I repeat the simulations with the lower gain $g = .005$, using the same history of shocks ε_t and v_t^i. Results are reported in Fig. (7). It can be

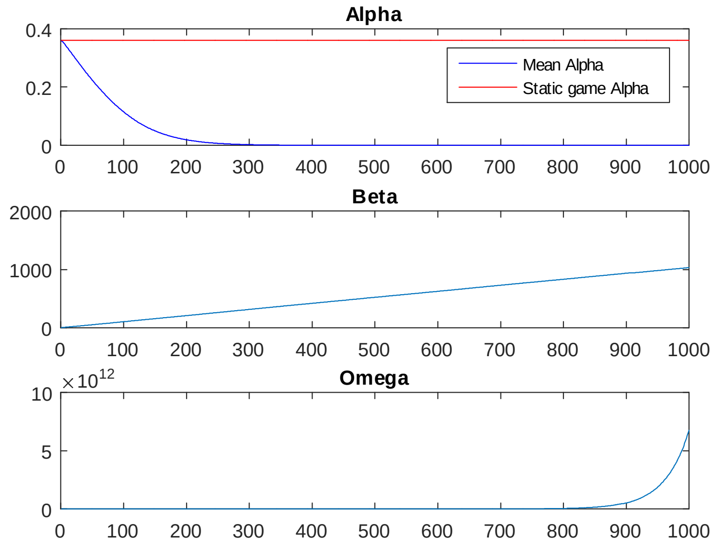

Figure 6: Evolution of $\hat{\omega}_t$ and average $\hat{\beta}_t^i$ and $\hat{\alpha}_t^i$ under constant gain learning. $g = .025$.

seen that a lower gain parameter implies a slower convergence of average $\hat{\alpha}_t^i$ towards zero: this is due to the fact that a lower gain implies a slower increase in the precision of the public signal, as estimates of the variance of prices converge more slowly. The impact on the estimated precision of the private signal, $\hat{\beta}_t^i$, is instead negligible. This is due to the fact that x_t^i is exogenous and its variance is thus not affected by the beliefs of agents: while a lower gain does marginally slow down convergence also in this case, the self-reinforcing mechanism introduced by the endogeneity of the variance of prices with respect to its estimates is not at play here.

One issue with the way I have set up the adaptive learning algorithm is that (44)-(45) are the counterpart of (15)-(17) and thus do not account for possible changes in the fundamental value (that is, they represent the sample variance of the whole series of observations for private and public signals from day one, not discounted for the possibility that a change in the fundamental value has taken place at some point in time). While there is no clear way to amend these equations under a constant gain learning, I try to account for this shortcoming by considering a framework that embeds both extreme cases of no change and sure change in the fundamental as special cases. In particular, I redefine equations (44)-(45) as

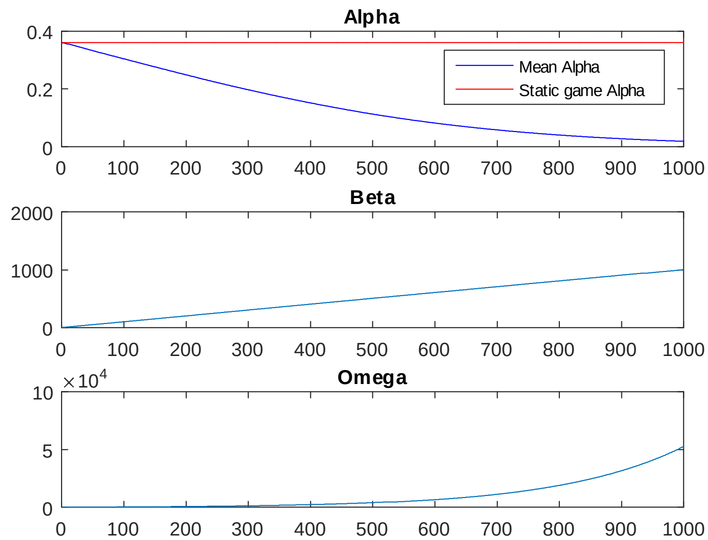

Figure 7: Evolution of $\hat{\omega}_t$ and average $\hat{\beta}_t^i$ and $\hat{\alpha}_t^i$ under constant gain learning. $g = .005$.

$$\hat{\beta}_t^i = (1-g)\,\hat{\beta}_{t-1}^i + \left(s_{x,t}^i - \left(\hat{x}_t^i\right)^2\right)^{-1} \tag{61}$$

$$\hat{\omega}_t = (1-g)\,\hat{\omega}_{t-1} + \left(s_{p,t} - \left(\hat{p}_t\right)^2\right)^{-1}. \tag{62}$$

Though not grounded in any formal optimization problem, this framework captures the fact that, when Nature can redraw the fundamental value, there is uncertainty about the cumulative nature of the signals. In particular, for $\pi = 1$ theory would dictate that $g = 1$ and equations (61)-(62) would reduce to a static framework where the optimal signals (both endogenous and exogenous) are represented only by current realizations; for $\pi = 0$, instead, theory would dictate that $g = 1/t$, which converges to zero at t increases without bound and leads to a dynamic framework where no changes in the fundamental are allowed and the optimal signal is a weighted cumulative aggregation of all past realizations. Simulating the system with these new specifications for $\hat{\beta}_t^i$ and $\hat{\omega}_t$, I find that for values of g usually considered in the literature (between .005 and .05), $\hat{\alpha}_t^i$ still converges to zero: the reason is that, while $\hat{\beta}_t^i$ remains bounded, $\hat{\omega}_t$ increases without limit. This result is similar to the one found in Section 2.4.1, where a low π led to $\tilde{\alpha}_t$ converging to zero as the precision of the public signal grew without bounds. Values of the constant gain usually adopted in the literature, thus, are not high enough to compensate for the tendency of the precision of the public signal to increase without bound due to the endogeneity of prices. I will discuss in more detail the relationship between π and g in the next Section.

The main result from this Section, thus, is that while constant gain adaptive learning is effective in tracking changes in the exogenous variables (simulations show that estimates of x_t^i and $s_{x,t}^i$ adjust quickly after a change in the fundamental value), it does not prevent the weight on private information from decreasing towards zero over time as the estimated variance of prices decreases towards zero and the precision of the public signal grows without bound. At the same time, though, because such convergence is rather slow at conventional values for the gain parameter, changes in the fundamental can be passed onto prices under such a learning scheme for a long time.

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 the 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 probability of a change in the fundamental happening at each time t. 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), the updating rule for the adaptive learning scheme, equation (52), with constant gain leads to

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

assuming $\hat{x}_0^i = 0$. I 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, $\tilde{x}_t^i$ is given by (32), with weights from (33)-(34), which, in a non-recursive way, can be rewritten as

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

While the weighting structure in the Bayesian framework is more convoluted, it can be seen 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, I show b_j^t and h_j^t in Figure 8. Curves are computed for $g = 0.025$ and $\pi = 0.01$, with $t = 100$. It can be seen that, despite being derived in different

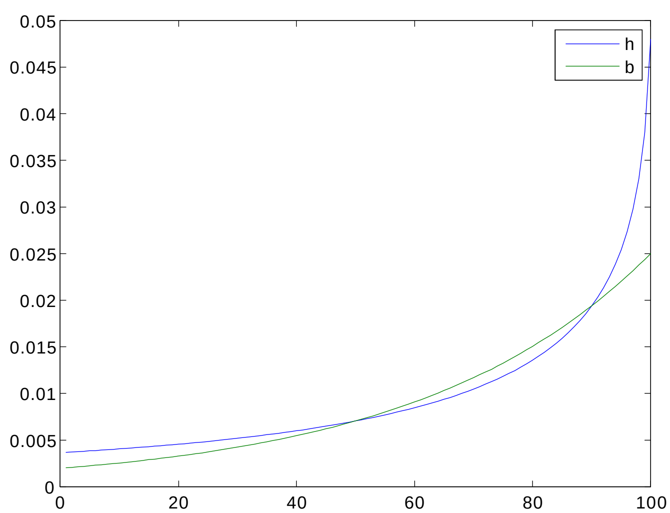

Figure 8: Weights on past observation under constant gain adaptive learning (b) and Bayesian learning (h).

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.

5 Discussion and conclusions

I 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 drives the volatility of signals to zero, but it also implies that all weight in the limit is put on the public endogenous signal (prices) and none on the private exogenous one. While prices aggregate information perfectly in the limit, this opens up the possibility of a misalignment of prices from the fundamental should such fundamental change.

To investigate this issue, I have extended the framework to allow for the fundamental to change with some fixed, exogenous probability. While past information gets discounted according to the probability of it being still relevant for current inference, even for a relatively high probability of changes in the fundamental happening every period agents still end up relying only on prices as signals in the long run. Private information is disregarded in the limit and only the public signal is used: a form of rational herding with informational cascade emerges.

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 derived from previous agents' behavior. Though in the present model all agents act simultaneously, in the dynamic framework where agents trade repeatedly over time I obtain a similar outcome: 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 my framework, similarly, as the weight on private information 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 of information effectively fails. If a change in the fundamental happens at this point, it does not get factored into prices.

Two assumptions are important for the results in this paper and are worth discussing in detail here. First, the fact that changes in the fundamental are statistically uncorrelated. That is, when the fundamental changes, the new value and the old one are independent and drawn from the same improper distribution. Moreover, there is an exogenous probability π that such changes take place at each time. This implies that, while past signals never become completely uninformative, as there is always a chance that a change in the fundamental never happened, such probability decreases very fast even for relatively large values of π. For example, for $\pi = .5$, the probability that a signal received at time $t = 1$ will is still relevant at time $t = 20$ is .000095%. If changes in the fundamental were smooth and correlated, such decline in probability would be much more gradual, thus quantitatively limiting the effect of agents discarding private signals when forming beliefs about the fundamental.

The second important assumption is that agents don't update the probability of changes in the fundamental, their subjective probability always being equal to the unconditional probability π. One could instead allow agents to update their (prior) subjective probability based on the observed signal: signals close to the current estimate of the fundamental would be more likely to come from a normal distribution centered around the old fundamental rather than from a new draw from the improper distribution of fundamentals. If such updates were allowed, it is possible to conjecture that, for any given π (and thus prior), the subjective posterior that the fundamental has changed would be higher the farther away the new observation falls from the most recent estimate of the fundamental. Results from simulations in Section (2.4.1) thus suggest that in this case small changes in the fundamental would be more likely to be overlooked, as they would more likely be misinterpreted as noise in the signal rather than as a change in the fundamental. That is, small changes in the fundamental would lead to a relatively smaller posterior (for any given π) for the subjective probability of changes, which would have the effect of making $\tilde{\alpha}_t$ more likely to converge to zero and at a faster rate.

In the second part of the paper I then relaxed a strong assumption in Bayesian learning, namely that agents know the statistical properties of the signals they are receiving. This assumption is particularly troublesome in cases where, like in this paper, the precision of a signal is endogenous and changes over time. For this reason I model agents as joint Bayesian and adaptive learners: they use adaptive learning to infer the moments of the distributions of the two signals from past data, and then use such moments to update their priors through Bayesian learning. In this setting, adaptive learning with a decreasing gain allows the investigation of long run convergence in case of a stationary environment; a constant gain algorithm instead provides the suitable framework to analyze outcomes in case changes to the underlying fundamental are possible. Overall, joint Bayesian and adaptive learning lead to results that confirm the possibility of prices and fundamental values to become misaligned in the long run, though over long periods of time agents keep using their private information and prices can reflect changes in fundamentals.

6 Appendix

6.1 Derivation of α^* with private and public signals

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

$$\alpha^* = \arg\min_{\alpha} E_t \left(\theta - \tilde{\theta}_t^i\right)^2 \tag{63}$$

with

$$\tilde{\theta}_t^i \equiv \alpha x_t^i + (1-\alpha)\, p_t. \tag{64}$$

Minimizing (63) subject to (64) leads to the first order condition

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

whose solution implies

$$\alpha^* = \frac{E_t p_t^2 - E_t p_t x_t^i}{E_t p_t^2 + E_t \left(x_t^i\right)^2 - 2E_t p_t x_t^i}. \tag{65}$$

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

$$\alpha^* = \frac{\sigma_p^2}{\sigma_p^2 + \sigma_v^2}. \tag{66}$$

6.2 Derivation of the price equation

Combining (5) and (7) gives

$$\sigma_{w,t}^2 = \left(\frac{\sigma_{p,t}^2}{\sigma_{p,t}^2+\sigma_v^2}\right)^2 \sigma_v^2 + \left(\frac{\sigma_v^2}{\sigma_{p,t}^2+\sigma_v^2}\right)^2 \sigma_{p,t}^2 \tag{67}$$

$$= \frac{\left(\sigma_{p,t}^2\right)^2 \sigma_v^2 + \sigma_{p,t}^2\left(\sigma_v^2\right)^2}{\left(\sigma_v^2+\sigma_{p,t}^2\right)^2} = \frac{\left(\sigma_v^2+\sigma_{p,t}^2\right)\sigma_{p,t}^2\sigma_v^2}{\left(\sigma_v^2+\sigma_{p,t}^2\right)^2} = \frac{\sigma_{p,t}^2\sigma_v^2}{\sigma_v^2+\sigma_{p,t}^2} = \alpha^*\sigma_v^2. \tag{68}$$

Substituting then (68) into (8) leads to

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

and therefore

$$\sigma_{p,t}^2 = \gamma^2\left(\sigma_v^2\right)^2\sigma_\varepsilon^2.$$

References

[1] Adam, K., Marcet, A., Nicolini, J.P., 2016. Stock market volatility and learning. Journal of Finance 71, 33-82.

[2] Admati, A.R., 1985. A noisy rational expectations equilibrium for multi-asset securities markets. Econometrica 53, 629-658.

[3] Admati, A.R., 1991. The informational role of prices. Journal of Monetary Economics 28, 347-361.

[4] Allen, F., Morris, S., Shin, H.S., 2006. Beauty Contests and Iterated Expectations in Asset Markets. Review of Financial Studies 19, 719-752.

[5] Amador, M., Weill, P-O., 2012. Learning from private and public observations of others' actions. Journal of Economic Theory 147, 910-940.

[6] 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.

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

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

[9] Basak, S., Chabakauri, G., 2010. Dynamic mean-variance asset allocation. Review of Financial Studies 23, 2970-3016.

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

[11] Berardi, M., Galimberti, J.K., 2017. Empirical calibration of adaptive learning. Journal of Economic Behavior and Organization 144, 219-237.

[12] 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

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

[14] 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.

[15] Bray M., Kreps, D.M., 1987. Rational learning and rational expectations. In G.R. Feiwel (ed.), Arrow and the Ascent of Modern Economic Theory, 597-625. New York: New York University Press.

[16] 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.

[17] 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.

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

[19] Diamond D.W., Verrecchia, R.E., 1981. Information aggregation in a noisy rational expectations economy. Journal of Financial Economics 9, 221-235.

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

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

[22] Grossman, S., 1976. On the efficiency of competitive stock markets where traders have diverse information. Journal of Finance 31, 573-585.

[23] Grossman, S., 1978. Further results on the informational efficiency of competitive stock markets. Journal of Economic Theory 18, 81-101.

[24] Grossman, S., 1981. An introduction to the theory of rational expectations under asymmetric information. The Review of Economic Studies 48, 541-559.

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

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

[27] Kyle, A. S., 1985. Continuous auctions and insider trading. Econometrica 53, 1315-1335.

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

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

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

[31] 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.

[32] Vives, X., 1993. How fast do rational agents learn? Review of Economic Studies 60, 329-347.

[33] Vives, X., 1995a. Short-term investment and the information efficiency of the market. The Review of Financial Studies 8, 125-160.

[34] Vives, X., 1995b. The speed of information revelation in a financial market mechanism. Journal of Economic Theory 67, 178-204.

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