Informational differences and learning in an asset market with boundedly rational agents

Cees Diks & Pietro Dindo

Center for Nonlinear Dynamics in Economics and Finance, Department of Quantitative Economics,

University of Amsterdam, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands

July 25, 2006

Abstract

The price formation mechanism in an asset market with boundedly rational agents can be viewed as a filter acting on incoming news about economic fundamentals such as future dividends. Here we study the properties of an asset pricing market filter obtained under some simple behavioral assumptions, and examine the resulting dynamical structure of the fluctuations of the market price around the time-varying underlying fundamental reference price. The starting point is an asset pricing model in which agents can choose among two different degrees of information on fundamentals. At the same time agents are also learning the parameter of the dividend generating process. This leads to prices that deviate substantially and persistently from the fundamental value in the short run but stay close to it in the long run. In particular, prices have a time-varying nonlinear mean reverting dynamics which we show to be related to agents' interaction triggered by informational differences.

Keywords: Asset price dynamics; Bounded rationality; Informational efficiency; Nonlinear mean reversion

JEL Classification: G12, D8, E3

1 Introduction

Since the beginning of the eighties, the validity of the efficient market hypothesis has been questioned on the basis of empirical evaluation of so-called financial anomalies. Well-known examples are excess volatility, as described by Shiller (1981) and LeRoy and Porter (1981), reversion of the asset prices to its mean, as documented by Poterba and Summers (1988) and Fama and French (1988b), and correlation between returns and lagged returns or lagged dividend yields, as shown by Shiller (1984) and Fama and French (1988a).

Stimulated by these findings, a part of the scientific community has investigated whether such anomalies can be explained by assuming that the agents operating in the market are boundedly rational. Although the exact implication of bounded rationality varies among the different models, a common characteristic is that boundedly rational agents act in an economic setting which they do not know in full detail. Furthermore, agents are often assumed to be able to optimize an objective function under certain constraints but unable to optimally anticipate the effect of their and other agents' actions. In particular boundedly rational agents are not assumed to be able to coordinate their actions such that their beliefs are self fulfilling. In other words, expectations of boundedly rational agents need not to be rational.

In order to explain fluctuations in prices that are not due to fluctuations in economic fundamentals, one class of models with boundedly rational agents concentrates on the interaction of agents choosing different expectation schemes, as first analyzed in Brock and Hommes (1997). In a study of a stylized financial market, Brock and Hommes (1998) assume that agents do not know whether it is more profitable to use a strategy that predicts prices by relying on fundamental information, or a strategy that predicts prices by extrapolating trends. Agents use realized profits or a similar performance measure to decide which strategy to use. The result of this choice for the best strategy, or best predictor, leads to complex endogenous price fluctuations. Since such endogenous fluctuations can already arise in the absence of exogenous influences such as time-varying fundamentals, explicit exogenous noise sources are often avoided, for instance by assuming constant fundamentals and reformulating the price dynamics in terms of deviations from the fundamental value. An advantage of this assumption is that the price dynamics can be specified in terms of (often nonlinear) difference or differential equations which can be analyzed analytically.

A limitation of models in this class is that they typically do not take into account the direct effect of news about the economic fundamentals on agents' behavior thus excluding one of the most trivial behavioral scenarios one might deem important in asset price formation — the over- or under-reaction of agents, and hence the market, to new information. Generally speaking, market models that tend to a stable equilibrium state in the absence of news can still show fluctuations triggered by the arrival of new information. Because *a priori* we do not know if market fluctuations are necessarily self-perpetuating as in chaotic dynamics, we explicitly wish to examine the role of exogenous noise on the price dynamics, thus keeping open the possibility of scenarios where ongoing market fluctuations require repeated triggering by a sequence of exogenous shocks.

In view of this critique, there is another class of models in the literature of asset markets with boundedly rational agents, which explicitly takes into account the role of news on fundamentals on the price dynamics. Early examples are Bulkley and Tonks (1989) and Barsky and De Long (1993) who investigate the effect of agents' learning of the growth rate of dividends from movements in the stock price. More recent examples are Timmermann (1993), Timmermann (1996) and Barucci *et al.* (2004), who assume that agents estimate the parameters that define the relationship between prices and dividends. In all these cases, agents use the rational expectations relationship that would hold between endogenous variables (prices) and exogenous variables (dividends) if the underlying parameters were known. That is, agents do not take into account that their learning effort is modifying the way dividends feed back into prices. When new information about dividends becomes available, it influences returns not only directly but also indirectly as it affects the estimates of the parameters that the agents use to forecast future prices and/or future dividends. These models converge to rational expectations when the agents learn the parameters of the data generating process. A limitation of these models is that they all assume the presence of a representative agent, so that agents' interaction triggered by informational differences or by different expectations do not play a role. Moreover, due to the stochastic components associated with the incoming news about the fundamentals, results are practically always obtained by means of simulations.

Motivated by these arguments, the aim of this paper is to construct a framework for examining markets with boundedly rational agents where both parameter estimation and agents' interaction play a role. When we use boundedly rational agents, we do not question that rational behavior, and especially rational expectations, can be a good approximation of the equilibrium of agents' repeated interaction, rather. We rather argue that the convergence to such an equilibrium is worth investigating as it might explain part of the economic variables' fluctuations we observe in reality. Our objective is to characterize how both parameter estimation and agents' interaction transform incoming information into realized market prices. Because it is impossible to carry out this exercise under all conceivable behavioral assumptions, we limit ourselves to a simple class of agent models, where all agents act upon the information available to them regarding fundamentals (including that revealed by prices). Agents' interaction is triggered by different expectations and different expectations can be explained by different degrees of information regarding the future value of dividends. This means that agents neither extrapolate price trends or use other chartists' rules per se, nor expect other agents to do so, so that second – or higher – order expectations play no role.

A convenient characteristic of our model is that it contains two important benchmarks as restrictions. The first benchmark is given by the classical asset pricing model: the equilibrium price we derive coincides with the correct present value price when one discards both the role of informational differences and of agents' learning the growth rate of dividends. The second benchmark is given by the model developed in Barsky and De Long (1993): our equilibrium price coincides with the price derived in the model of Barsky and De Long if does take in to account that agents' are learning the growth rate of dividends and discards the role of informational differences. Our model can thus be seen as an extension of Barsky and De Long model when agents with different information sets are active in the market.

We investigate the extent to which our model is able to explain empirical properties of asset prices. As it turns out, the agent-based dynamics driven by exogenous noise leads to a simple econometric model for prices that can account for several well-documented anomalies such as autocorrelation of returns and large persistent deviations of the market price from the fundamental price in the short run but convergence to it in the long run. In fact, in line with the econometric model proposed by Summers (1986), our model leads to a (log) price which is the sum of a persistent component, proportional to the (log) dividend, and of a transitory component, proportional to the (log) dividend yield, which we derive to follow a stationary AR(1) process with a time-varying AR(1) coefficient. Our analysis shows that whereas the fact that the transitory component follows an AR(1) process is a direct consequence of agents' learning the model parameters, the fact that the AR(1) coefficient is time-varying is due to agents' interaction triggered by informational differences. This offers theoretical support to the empirical evidence that the temporary component in the mean reversion is nonlinear and is switching between regimes with different rates of convergence, as documented both by Gallagher and Taylor (2001) and by Manzan (2003).

As we consider an asset market where agents have different degrees of information, our framework is closely related to that of Grossman and Stiglitz (1980) (henceforth GS). They investigate whether the price is informationally efficient in a repeated market for a one period living asset where agents can decide between two different degrees of information about the value of the asset return at the end of the period. As GS, we also assume that agents operating in the asset market can decide whether or not to be informed about next period's dividend. In contrast to GS, and similar to Bray (1982), Hellwig (1982) or Routledge (1999), we consider a dynamic model rather than a static one. By this we mean that we do not start off by assuming that agents have rational expectations but rather see rational expectations as a possible long run outcome of a learning process of boundedly rational agents which are using simpler rules. Failure of the uninformed agents to learn the relationship between prices and dividends implies deviations of the price from its fundamental value. Moreover, we assume that the fractions of informed and uninformed agents are not constant but change over time based on past performances of both strategies. The fraction of each type of agents is thus an endogenously determined variable. Another difference with the framework of GS and followers is that we model a market for an infinitely living asset rather than of a sequence of identical markets for a single period asset. This implies that agents need to form expectations not only on the future values of the dividend but also on the remaining value of the asset. To our knowledge, Goldbaum (2005) is the first to consider a dynamic multi-period market in a setting where agents have different degrees of information. Whereas Goldbaum assumes the asset's dividend to be stationary in differences, we assume, in order to stay closer to real data, that the asset's dividend is stationary in log-differences. Accordingly, we choose to derive our agents' demand from mean variance maximization of a constant relative risk aversion (CRRA) utility function rather than from a constant absolute risk aversion (CARA) utility function as in Goldbaum (2005). Choosing a CRRA framework also allows us to model the evolution of fractions of informed and uninformed agents by a replicator dynamics type of mechanism (see Weibull 1995 for an introduction to this kind of evolutionary dynamics). In fact, we show that, the evolution of informed and uninformed agents' relative wealth, or market power, leads naturally to the replicator dynamics.

Our model shows that financial markets populated by agents with different degrees of information can be seen as economic systems with negative feedback. This establishes a precise correspondence with the famous cobweb model, see Ezekiel (1938), and with the literature that originated from it, such as Muth (1961) and Brock and Hommes (1997). Our equation for the evolution of the dividend yield as a function of the uninformed agents' expectations has a perfect correspondence with the equilibrium price equation in a cobweb model in the case of linear supply and linear demand. We refer to this literature to justify the expectation formation of boundedly rational agents. In particular Brock and Hommes (1997) show that if rational expectations come at a cost, boundedly rational agents keep switching between an expensive sophisticated and a cheap simple expectation scheme, thus generating complicated price fluctuations. Because we want to keep our model as simple as possible, we do not explicitly model agents' choice between cheap simple expectations and expensive sophisticated expectations. We concentrate on informational difference and model agents' expectations as adaptive. An analysis where both informational differences and the role of expectations scheme choice play a role is performed by De Fontnouvelle (2000). He shows that if agents are allowed to switch among different types of expectations schemes and if rational expectations come at a cost, an asset market of the type proposed by GS leads to similar price fluctuations as Brock and Hommes (1997) found for the cobweb model. By considering informational differences as well as the choice of an expectations scheme, even in the simpler case of a one period living asset, De Fontnouvelle arrives at a rather complicated model that is analyzed mostly by means of simulation, rather than analytically as we do here.

The paper unfolds as follows. Section 2 introduces the model in terms of its three founding elements: the asset market (2.1), the expectation formation (2.2) and the evolution of the fractions of informed and uninformed agents (2.3). Section 3 analyzes the co-evolution of the market price and of the fractions of informed and uninformed agents in a world without uncertainty about future growth rates of dividends. Technically in this section we analyze the deterministic skeleton of the system of difference equations developed in Section 2. Section 4 analyzes the full model, i.e. the evolution of the market price and of fractions of agents when uncertainty about future growth rates of dividends plays a role. Here we also relate the price dynamics generated by our model, by the classical asset pricing model and by the model developed in Barsky and De Long (1993), with respect to some well-know "financial anomalies". Section 5 concludes.

2 The model

2.1 The asset market

We consider a market where shares of a financial asset are traded repeatedly in discrete time periods. The market is populated by agents who believe that the discounted sum of expected future dividends constitutes a "fair" price. As in GS, every agent can decide whether or not to buy information about next period's dividend. As a result, in every period the population of agents is divided in two groups with different degrees of information concerning fundamental variables. The current setting differs from GS in that, in our model, the asset represents a claim on an infinite sequence of future dividends rather than on a single dividend, that is, the asset is infinitely lived and does not perish at the end of the period. As a consequence agents, besides forming expectations on dividends, must also form expectations on future asset prices. Another important difference with respect to the GS framework is that in our model agents are boundedly rational. By this we mean that agents are not able to compute the equilibrium relationship between price and dividends that should arise in the market where informed and uninformed agents operate. The aim of this Section is to characterize how, in this setting, the market price of an asset/share, p_t, and the fraction of informed agents, λ_t, co-evolve given agents' expectations and the dividend process $\{d_t\}$. In order to arrive at such a relationship we first specify the underlying assumptions of our model.

Assumption *(i)* The dividend process, $\{d_t\}$, is stochastic. In the benchmark case $\{d_t\}$ is given by a geometric random walk. At time t, d_t, is given by:

$$d_t = d_{t-1}(1+g)\eta_t,$$

where $\{\eta_t\}$ is a sequence of independent, identically distributed (i.i.d.) random variables with mean 1 and variance σ_η^2. This implies that the constant $(1+g)$ is the long run growth rate of dividends.

Assumption *(ii)* Agents know that the dividend is growing over time at a certain rate which they estimate using past dividend realizations. We let g^e denote their belief, or estimate, of the long run growth rate of dividends. We assume that this belief is the same across agents and that agents use it for predictions "as if" it is the true value in the dividend generating process. For the moment we consider g^e as given and not time dependent. In Subsection 2.2 we will discuss how agents actually form their beliefs, g^e, regarding the long run growth rate g.

Assumption *(iii)* All agents are "fundamentalists" in the sense that they follow the present value model, i.e. the discounted sum of all future dividends is their "fair" value of the asset. The exact relationship between today's price and tomorrow's expected dividend depends on the agent's information about future dividends. The general information set contains past dividend and price realizations and its exact expression is different for different agents' groups. We denote the information set at time t for group H as $\mathcal{F}_t^H$. The fair value, i.e. the expected discounted sum of future dividends, conditional on $\mathcal{F}_t^H$ is denoted by v_t^H:

$$v_t^H = E\left[\sum_{i=1}^{\infty} \frac{d_{t+i}}{(1+r)^i} \,\middle|\, \mathcal{F}_t^H\right]. \tag{2.1}$$

The coefficient $(1+r)$ is the discount rate or required rate of return. We assume that the discount rate is the same across agents and that agents use the same discount rate for all future periods. The latter assumption is made because we want to concentrate on sources of price fluctuations given by agents' interaction and agents' learning rather than agents' changing of their discount rate. In general, the discount rate can be characterized as the sum of the risk free rate and the risk premium, which depends on the risk preferences of agents. In this case, to state that agents use the same constant discount rate means that agents have the same constant risk preferences. See

Appendix A for a derivation of the risk premium in a context where preferences of the agents are explicitly taken into account. We also assume that the discount rate is always bigger than the agent's estimate of the long run growth rate of dividends. Notice that together with assumption *(ii)*, assumption *(iii)* implies that agents use a static Gordon model to evaluate the asset.

Assumption *(iv)* At time t, each agent can decide whether to buy information about the value of d_{t+1} or not to buy it and thus remain uninformed. As a result, in every period there are two groups of agents having a different degree of information regarding the next realization of the dividend process. At time t, the informed agents, group I, are fully informed regarding d_{t+1}. This implies that their current expectation of the $t+1$ dividend is

$$d^{e,I}_{t,t+1} = d_{t+1}, \tag{2.2}$$

where the superscript e, I stands for expectations of the informed agents. We assume that they pay a fixed per period cost $c > 0$ for this information. The uninformed agents, group U, do not know d_{t+1} but can use public information, available in the form of realized dividends d_s and realized prices p_s, $s \leq t$, to form their expectations, $d^{e,U}_{t,t+1}$, about d_{t+1}. The superscript e, U stands for expectations of the uninformed agents. For example, if uninformed agents would rely solely on the public belief of the dividend growth rate, they could use $d^{e,U}_{t,t+1} = (1 + g^e)d_t$. The alternative which we consider here, is that they try to get additional information revealed by the demands of informed agents through the current market price p_t. Uninformed agents consider a relationship between the dividend and the price of the form:

$$d^{e,U}_{t,t+1} = y^e p_t, \tag{2.3}$$

where y^e is agents belief, or estimate, of the market dividend yield, that is, the ratio between future dividend and current price. As we have done for g^e, we start with considering y^e as given and fixed. In Subsection 2.2 we will discuss how agents actually form their beliefs of the market dividend yield. Finally, we let λ denote the fraction of informed agents, so that $1 - \lambda$ is the fraction of uninformed agents. A subscript t is added when we consider a time dependent λ. We use time varying fractions only from Subsection 2.3 where we describe how λ_t evolves endogenously.

Assumption *(v)* At each time t, the ex-dividend market price of one share, p_t, is given by the following market equilibrium pricing equation:

$$p_t = \lambda v^I_t + (1 - \lambda)v^U_t, \tag{2.4}$$

where v^I_t and v^U_t are the "fair" value of the asset conditional on the information of the informed and uninformed respectively, as derived below. Under this assumption the realized price today is a weighted average, with weights being equal to the fraction of each agent type, of the agent's estimate of the share fair value. Although this market equilibrium pricing equation is admittedly stylized, it can be derived by assuming that agents can choose to invest in a risky asset and in a risk free bond and use a mean variance utility to decide how much of their wealth to allocate in each investment. If one starts from such a micro-foundation of agents' demands, assumption (v) translates into the assumption that markets clear as in a Walrasian framework. See Appendix A for a derivation of (2.4).

Given assumptions $(i) - (v)$ the next step is to derive their implication on the price dynamics in (2.4). First, we compute the fair value for the informed and for the uninformed agent.

Informed agents Assumption (ii) implies that informed agents, like all other agents, know that the dividend is growing over time at a certain rate, which they assume to be equal to g^e and which they use for predictions "as if" it is the true value of the growth rate of the dividend process. Assumption (iii) implies that expectations of future share values are directly linked to expected future dividends through equation (2.1). Assumption (iv) implies that at time t the informed agents know the value of d_{t+1} so that their information at time t is given by $\mathcal{F}_t^I = \{d_{t+1}, d_t, \dots, p_t, p_{t-1}, \dots\}$. Hence their expectations of future dividends are:

$$d_{t,t+i}^{e,I} = d_{t+1}(1+g^e)^{i-1}, \quad \text{for} \quad i \geq 1. \tag{2.5}$$

Notice that since agents treat their estimate g^e as if it is the true value of the growth rate of dividends, they do not take into account possible estimation errors in their dividend predictions. Plugging expectations (2.5) into (2.1) we arrive at the informed agents' estimate of the value of the stock,

$$v_t^I = \frac{d_{t+1}}{(r - g^e)}. \tag{2.6}$$

The informed agents' stock evaluation is thus in accordance with the Gordon model. Equation (2.6) implies that informed agents are behaving similarly to so-called fundamentalists in the interacting agents literature, see Hommes (2006) and Le Baron (2006) for recent surveys. In fact, the informed agents' fair value of the asset is proportional to the dividend payed at time $t+1$. For this reason, we shall refer to v_t^I, the fair price of the informed agents, as the fundamental price p_t^*, that is, we define:

$$p_t^* = \frac{d_{t+1}}{(r - g^e)}. \tag{2.7}$$

Notice that the value of the fundamental price depends on g^e, the common agents' belief of the growth rate of dividends.

Uninformed agents We have assumed in (iv) that boundedly rational uninformed agents try to infer the value of d_{t+1} from the market clearing price p_t. In doing that, they use their model (2.3) to explicit the relationship between the current realized market price p_t and expected future dividends d_{t+1}. Combining this with assumptions (ii) and (iii) as we did for the informed agents we obtain

$$d_{t,t+i}^{e,U} = y^e p_t (1+g^e)^{i-1}, \quad \text{for} \quad i \geq 1, \tag{2.8}$$

which, using (2.1), and the uninformed information set $\mathcal{F}_t^U = \{d_t, d_{t-1}, \ldots, p_t, p_{t-1}, \ldots\}$, gives the uninformed agents' estimate of the value of the stock:

$$v_t^U = \frac{y^e p_t}{(r - g^e)}. \tag{2.9}$$

Notice that also for uninformed agents, there is a correspondence with the literature on interacting agents. In fact, our ("fundamentalists") uninformed agents are behaving "as if" they are chartists, that is, they use current prices to estimate the value they attach to the asset. This is an important characteristic of our model and we anticipate its consequences here. Consider g^e and y^e as given. If at time t, $y^e/(r - g^e)$ is bigger than one, uninformed agents behave "as if" they are trend followers, and they drive prices well above the fundamental levels. In this case the realized dividend yield, $y_{t+1} = d_{t+1}/p_t$, will become lower than $(r - g^e)$. The converse happens when $y^e/(r - g^e) < 1$. In this case the uninformed agents behave "as if" they are contrarians. Summarizing, the uninformed agents behave "as if" they are chartists but with a different trend coefficient for different values of g^e and y^e.

Given the expectations of both types of agents, specified in (2.6) and (2.9), and the market equilibrium pricing equation (2.4) we get,

$$p_t = \frac{d_{t+1}}{(r - g^e)} \frac{\lambda(r - g^e)}{(r - g^e - (1 - \lambda)y^e)}. \tag{2.10}$$

This equation shows that the realized price is proportional to the fundamental price p_t^* in (2.7) – the same result one gets through the Gordon model – but that there is an additional behavioral factor due to the presence of the uninformed agents trying to extract information from the market price. In fact, one can interpret this as a generalization of the Gordon equation to a simple setting where agents have different degrees of information.

Equation (2.10) can be written as a relationship between the realized dividend yield $y_{t+1} = d_{t+1}/p_t$, agents' beliefs g^e and y^e, and the fraction of informed agents' λ:

$$y_{t+1} = \frac{(r - g^e)}{\lambda} - \frac{(1 - \lambda)}{\lambda} y^e := f(y^e, g^e, \lambda). \tag{2.11}$$

We call the map f the feedback map because, given a fraction of informed agents λ and common beliefs about the growth rate of dividends g^e, it establishes a feedback between expectations of uninformed agents of the dividend yield and dividend yield realizations. Using the feedback map defined in (2.11), it can be easily derived that, for any $\lambda \in (0, 1]$, if $y^e = r - g^e$ then $y_{t+1} = r - g^e$. When this is the case, the agents' belief regarding the dividend yield is self-fulfilling and $r - g^e$ is thus the rational expectation dividend yield. Notice that when $y^e = r - g^e$, using equation (2.10), one gets that the market price equals the fundamental price p_t^* which we have defined in (2.7) as the fair value of the informed agents. For this reason we denote

$$y^* = r - g^e \tag{2.12}$$

as the fundamental dividend yield. The rational expectations dividend yield is thus equal to the fundamental dividend yield and, as we will specify later, it gives approximately the same price process as derived by Timmermann (1993) or by Barsky and De Long (1993) who also consider a model where agents are updating their estimate of the dividend growth rate g. The novelty here is that, due to the presence of informed and uninformed agents, y_{t+1} may fail to be equal to y^*. In particular the presence of uninformed agents generates a price that differs from the fundamental price. In fact, equation (2.11) establishes a negative feedback system between the dividend yield and its belief or estimate, since $\partial f/\partial y^e = -(1-\lambda)/\lambda < 0$. This implies that positive (negative) deviations of y^e from y^*, lead to negative (positive) realized deviations of y_{t+1} from y^*. This observation links our model to the classical cobweb model (see e.g. Ezekiel (1938) for an early treatment). In fact equation (2.11) for the price dividend ratio is the same as the equilibrium price equation in a cobweb model with linear supply and linear demand. The general asset price dynamics will be more complicated however, since, according to (2.11), asset prices are also driven by the learning of the growth rate of exogenous dividend process and by the evolution of agents' fractions.

2.2 Expectation formation

In this subsection we specify how agents form expectations on the long run growth rate of dividends and on the dividend yield. As far as the growth rate of dividends is concerned we have assumed in *(ii)* that both informed and uninformed agents have homogeneous expectations on the dividend growth rate. We follow Barsky and De Long (1993) and assume that agents use adaptive expectations to estimate its long run value. Adaptive expectations are characterized by

$$g^e_{t,t+1} = \gamma g^e_{t-1,t} + (1-\gamma)\left(\frac{d_t}{d_{t-1}} - 1\right). \tag{2.13}$$

where $g^e_{t,t+1}$ denotes the time t belief or estimate of the long run growth rate for period $t+1$, and where we refer to γ as the memory coefficient specifying the rate at which agents discount past information. Naive expectations are obtained in the special case $\gamma = 0$ and that $g^e_{t,t+1}$ is the mean of all past observations of y when $\gamma = (t-1)/t$. Notice that when time-varying growth rate beliefs, one has to update the definition of the fundamental price and of the fundamental dividend yield from expression (2.7) and (2.12) to, respectively:

$$p^*_t = \frac{d_{t+1}}{r - g^e_{t,t+1}} \tag{2.14}$$

$$y^*_{t+1} = r - g^e_{t,t+1} \tag{2.15}$$

The same expectation technology used for the estimation of the future growth rate of dividends is assumed to be used for the estimation of the value of the future dividend yield y^e. Notice that only uninformed agents need to form expectations about tomorrow's dividend yield, as informed agents know already the value of d_{t+1}. Adaptive expectations for the dividend yield are specified by:

$$y^e_{t,t+1} = \alpha y^e_{t-1,t} + (1-\alpha) y_t, \tag{2.16}$$

where, as before, $y^e_{t,t+1}$ denotes the belief or estimate of the dividend yield of time $t+1$ based on the information up to time t, and the parameter $\alpha \in [0,1]$ is, as γ, the memory parameter, which specifies the rate at which agents discount past information.

Substituting the time varying expectations of the growth rate of dividends and of teh dividend yield into the feedback map (2.11) one obtains:

$$y_{t+1} = f(y^e_{t,t+1}, g^e_{t,t+1}, \lambda) = \frac{(r - g^e_{t,t+1})}{\lambda} - \frac{(1-\lambda)}{\lambda} y^e_{t,t+1}. \tag{2.17}$$

The previous equation evaluated at time t together with Eq. (2.16) gives the explicit dynamics for the expectations of the dividend yield as a function of its lagged values, estimates for the growth rate of dividends, and fractions of informed agents:

$$y^e_{t,t+1} = \frac{(1-\alpha)(r - g^e_{t-1,t})}{\lambda} + \frac{\alpha + \lambda - 1}{\lambda} y^e_{t-1,t}.$$

Using the feedback map (2.11) again one can rewrite this last dynamic expression in terms of realization of the dividend yield instead of in terms of expectations of the dividend yield. expectations. Changing variables from $(y^e_{t-1,t}, y^e_{t,t+1})$ to (y_t, y_{t+1}) gives:

$$y_{t+1} = \frac{r - g^e_{t,t+1}}{\lambda} - \frac{\alpha(r - g^e_{t-1,t})}{\lambda} + \frac{\alpha + \lambda - 1}{\lambda} y_t := h(y_t, g^e_{t,t+1}, g^e_{t-1,t}\lambda) \tag{2.18}$$

The updating map h establishes the dependence of the dividend yield on its lagged value, on the agents' beliefs of the long run growth rate of dividends at two subsequent dates and on the fractions of informed agents λ. Notice that the map h is linear in y_t. We study this map in the Section 3.

Adaptive expectations As shown by Muth (1960) adaptive expectations are optimal when the data generating process is a random walk plus i.i.d. noise. That is, adaptive expectations coincide with rational expectations when the data generating process for the variable to be estimated, say x_t, is given by

$$x_t = \varepsilon_t + (1-\mu) \sum_{i=1}^{t-1} \varepsilon_{t-i} \tag{2.19}$$

where $\{\varepsilon\}$ is an i.i.d. noise process and $\mu \in (0,1)$ measures which fraction of the shock has a persistent component. If one considers (2.19) as a data generating process, then adaptive expectations follow from

$$x^e_{t,t+1} = E_t(x_{t+1}) = \mu x^e_{t-1,t} + (1-\mu) x_t, \tag{2.20}$$

where E_t is the mathematical expectation given the information at time t and where the parameter μ is what we have called memory.

Adaptive expectations for the growth rate of dividend are not consistent with the process we assumed for the dividend. In fact, in *(ii)* the growth rate of dividends has only a temporary component. But the inconsistency is rather small when the memory μ is high. In fact, the effect of the persistent component is very small with μ close to 1.

Adaptive expectations for the dividend yield are also not consistent. or not rational in the sense of Muth (1961), but we assume that agents use them for two reasons. First, empirical support in favor of the use of adaptive expectations for the dividend yield in present values mode has been given by Chow (1989). Second, Brock and Hommes (1997) have shown that with endogenously determined variables as price or dividend yield, if rational expectations come at a cost, agents may switch continuously between costly rational expectations and simpler expectations at no cost. As a result, when one models this expectation choice, convergence to a rational expectations equilibrium is not ensured. In order to keep the structural form of our model at its simplest, we do not model agent's choice of its expectation framework and assume that boundedly rational agents use adaptive expectations. In fact adaptive expectation gives a reasonable trade off between simplicity of use and implementation and consistency with the outcomes of the models.

2.3 Evolution of the fraction of informed agents

So far we have assumed that the fraction of informed and uninformed agents are fixed. In this subsection we model how these fractions change over time. The driving force behind these changes is the trade off between the cost of information and the precision of the dividend yield estimator based on public information. Intuitively, given the cost of information, the more precise the estimate of the dividend yield, the bigger the fraction of uninformed agents. Or, given the precision of the dividend yield estimate, the higher the cost of being informed, and the smaller the fraction of informed agents. The next assumption specifies how exactly we model this process.

Assumption (*vi*) The evolution of the fraction of informed agents λ is modeled by the replicator dynamics mechanism. The replicator dynamics can be motivated in the context of boundedly rational agents who are learning and imitating which strategy to play in a strategic environment (see e.g. Weibull 1995 and Binmore and Samuelson 1997). Furthermore, the replicator dynamics arises naturally in a framework where the equilibrium pricing equation (2.4) is derived from the maximization of a mean variance utility function. In fact, in this framework, as outlined in Appendix 2, λ is related to the fraction of the total wealth possessed by the informed agents. Although it is beyond the scope of this paper to consider other specifications of the dynamics of λ, we believe the result to be valid more generally than just for the replicator dynamics discussed here. Since the objective of the agents is to gather information about future dividends, we assume that the success of a strategy is given by the squared forecast error of the dividend predictor. Informed agents have zero forecast error. Uninformed agents try to forecast the future dividends by estimating the dividend yield, so their squared forecast error for the realization y_t is $(y^e_{t-1,t} - y_t)^2$. The costs of information are $c > 0$ per time step for the informed agents, and 0 for the uninformed. As a result, we can define the fitness π_t of the strategies at time t. The fitness of the strategy to buy information is:

$$\pi^I_t = -c, \tag{2.21}$$

while the fitness of the strategy of remaining uninformed is

$$\pi^U_t = -\left(\frac{y^e_{t-1,t} - y_t}{y^*_t}\right)^2. \tag{2.22}$$

The time-varying fundamental dividend yield y_t^* defined in (2.15) is introduced in the denominator of π_t^U for normalization. Given that dividend yield y_t has y_t^* as reference value, this choice is convenient because it implies that the two fitness measures coincide when, given c, a forecasting error of $\sqrt{10c}\%$ is made. For example, if $c = 0.1$, the two fitness measures are equal with errors of 1% around y_t^*,

In the presence of a time varying fraction λ_t, Eq. (2.17) becomes:

$$y_{t+1} = \frac{r - g^e_{t,t+1}}{\lambda_t} - \frac{(1-\lambda_t)}{\lambda_t} y^e_{t,t+1}.$$

We can use this relation between y_{t+1} and $y^e_{t,t+1}$ and rewrite the fitness of the uninformed agents (2.22) as:

$$\pi_t^U = -\left(\frac{(y_t^* - y^e_{t-1,t})}{\lambda_{t-1} y_t^*}\right)^2, \tag{2.23}$$

Given the fitness measure of both strategies we can now specify the dynamics for λ. Under replicator dynamics the fraction λ_t of informed agents evolves according to

$$\lambda_t = (1-\delta)\frac{\lambda_{t-1}(\pi_t^I + \rho)}{\lambda_{t-1}\pi_t^I + (1-\lambda_{t-1})\pi_t^U + \rho} + \frac{\delta}{2}, \tag{2.24}$$

where the parameter ρ defines the speed of adjustment and δ is to be interpreted as a mutation or experimentation parameter (see, e.g., Young and Foster (1991); Droste *et al.* (2002)). The parameter ρ has to be taken larger than the cost c, $\rho > c$, to assure that the fitness measure of the informed agent is positive which ensures that the fractions are always in the interval $[0, 1]$. The parameter δ is related to what we call evolutionary (or selection) pressure in the following way: as $\delta \to 0$ the updating of the fractions is determined more and more by the selection mechanism. On the other hand when $\delta = 1$ the evolutionary pressure reaches its minimum and both fractions are $1/2$, independently on the fitness of the two strategies. Based on the expressions (2.21) and (2.23) the replicator dynamics (2.24) gives,

$$\lambda_t = (1-\delta)\frac{\lambda_{t-1}(-c+\rho)}{\lambda_{t-1}(-c) - (1-\lambda_{t-1})\left(\frac{y_t^* - y^e_{t-1,t}}{\lambda_{t-1} y_t^*}\right)^2 + \rho} + \frac{\delta}{2}. \tag{2.25}$$

In Appendix B we present an interpretation of this formula starting from wealth evolution of agents' maximizing a mean variance CRRA utility function. We mention here that, in the framework of Appendix B, the adjustment parameter ρ can be interpreted as a fixed profit, e.g. a profit coming from investment in a risk free asset, that washes out differences in profits given by the investment in the risky asset.

2.4 Market returns

To summarize, the full model developed so far is given by the following four equations

$$g^e_{t,t+1} = \gamma g^e_{t-1,t} + (1-\gamma)\left((1+g)\eta_t - 1\right), \tag{2.26}$$

$$y^e_{t,t+1} = \alpha y^e_{t-1,t} + (1-\alpha) y_t, \tag{2.27}$$

$$y_{t+1} = \frac{r - g^e_{t,t+1}}{\lambda_t} - \frac{(1-\lambda_t)}{\lambda_t} y^e_{t,t+1}, \tag{2.28}$$

$$\lambda_t = (1-\delta)\frac{\lambda_{t-1}(-c+\rho)}{\lambda_{t-1}(-c) - (1-\lambda_{t-1})\left(\frac{(y^*_t - y^e_{t-1,t})}{\lambda_{t-1} y^*_t}\right)^2 + \rho} + \frac{\delta}{2}. \tag{2.29}$$

Equation (2.26) gives the common expectation formation regarding the growth rate of dividends g^e, Eq. (2.27) gives the expectation formation of the dividend yield y^e by the uninformed traders, Eq. (2.28) the market equilibrium pricing condition which fixes the dividend yield y, and Eq, (2.29) the dynamics of the fraction of informed agents λ. The shocks and parameters are: η_t, the noise process driving the dividend growth; g, the long run dividend growth; γ, the memory agents use to estimate the future dividend growth; α, the memory agents use to estimate the future dividend yield; r, the required rate of return; δ, the experimentation or mutation rate; c, the cost of information per time step; ρ, the parameter which regulates the speed of adjustment of the replicator dynamics.

From the dynamics of the dividend price ratio and from the dividend process $\{d_t\}$, one can derive the asset price return. Since $y_{t+1} = d_{t+1}/p_t$, we have

$$\log(p_t) = \log(d_{t+1}) - \log(y_{t+1}). \tag{2.30}$$

Now, if we call q_t, the log price return at time t, the previous equation gives:

$$q_t = \log(p_t) - \log(p_{t-1}) = \log(d_t) - \log(d_{t-1}) - \log(y_{t+1}) + \log(y_t).$$

In (2.30), whenever y_t converges to its steady state value, that is, whenever $y^e_{t,t+1}$ converges to the rational expectations value $r - g^e_{t,t+1}$, the price follows $p^*_t = d_{t+1}/r - g^e_{t,t+1}$, i.e. the time varying correspondent of the fundamental price defined in (2.7). Notice that the fundamental price depends on the changing estimates of the growth rate of dividends and that it is the same price which has been derived by Barsky and De Long (1993). If, moreover, $g^e_{t,t+1} \to g$, the fundamental price converges to the "correct" present value price, in fact $y^*_t \to r - g$ and $p^*_t \to d_{t+1}/(r-g)$. If y_t fails to converge to $r - g$, deviations of the price from fundamental price can have two origins. The first is the failure of the deterministic skeleton of the system specified in (2.26-2.29) to converge to its fixed point, stated differently the fact that the adaptive expectations do not converge to rational expectations. This is related to the work of GS and to the fact that prices are not fully informative. The second possible reason is that, even if the system converges to the fixed point, it could approach an equilibrium where $g^e \neq g$. This is a situation where the fundamental price p^* is not equal to the "correct" present value price $d_{t+1}/(r-g)$. This is specifically relevant when the estimate g^e is time varying so that the system in (2.26-2.29) is stochastic. In what follows, we analyze these effects separately as well as their interplay. First, in Section 3, we analyze the conditions of convergence of the deterministic system dynamics of y and λ to their equilibrium values. Then, in Section 4, we complement this analysis by investigating the effect of a time varying stochastic g^e and how the two sources interact.

3 Informational differences

In this section we analyze the impact of informational differences alone on the dynamics of asset prices assuming that $g^e_{t,t+1} \equiv g^e$, i.e. there is no learning of the divided growth rate. Technically, we analyze the system of equations (2.26-2.29) when the memory parameter $\gamma = 1$. Without loss of generality we consider only the case $g^e = g$. The generalization to $g^e \neq g$ is straightforward and implies only a shift of the level of the steady state dividend yield from $r - g$ to $r - g^e$. To simplify the notation, throughout the rest of the paper we write $y^e_{t,t+1} \equiv y^e_{t+1}$ for the forecast of y_{t+1} made at time t.

When $g^e_{t,t+1} = g$, we can reduce the system (2.26-2.29) to a two dimensional (2-D) system in the variables y^e_{t+1} and λ_t. The result is:

$$y^e_{t+1} = \frac{(1-\alpha)(r-g)}{\lambda_{t-1}} + \frac{\alpha + \lambda_{t-1} - 1}{\lambda_{t-1}} y^e_t, \tag{3.1}$$

$$\lambda_t = (1-\delta) \frac{\lambda_{t-1}(-c+\rho)}{\lambda_{t-1}(-c) - (1-\lambda_{t-1}) \left(\frac{(r-g) - y^e_t}{\lambda_{t-1}(r-g)} \right)^2 + \rho} + \frac{\delta}{2}. \tag{3.2}$$

The parameters of the model are α, the memory of uninformed agents for their estimation of the future dividend yield; r, the required rate of return; g, the growth rate of dividends; δ, the experimentation level; c, the cost of information per time step; ρ, the speed of adjustment of the replicator dynamics. Given the dynamics of (y^e_{t+1}, λ_t) specified by (3.1-3.2), the dynamics of the dividend yield y_{t+1} can be derived by using the feedback map f defined in (2.11). Before investigating the full dynamics of (3.1-3.2) it is instructive to consider the 1-D system obtained when the fraction λ_t of informed agents is fixed to a constant value λ. Proofs of all the propositions can be find in the Appendix C at the end of the paper.

3.1 Dividend yield dynamics

Taking $\lambda_t \equiv \lambda$, Eq. (3.1) becomes:

$$y^e_{t+1} = \frac{(1-\alpha)(r-g)}{\lambda} + \frac{\alpha + \lambda - 1}{\lambda} y^e_t \tag{3.3}$$

Given the linearity of (3.3), the analysis of the dynamics is straightforward and it is possible to compute the general solution of the difference equation. That is, given y^e_0 one can compute the value of y^e_t, for all t. The following proposition summarizes the results.

Proposition 1. *Given the memory parameter $\alpha \in (0, 1)$, the fraction of informed agents $\lambda \in (0, 1]$, and the required rate of return $r > g$, we have:*
(i) *The solution of the difference equation (3.3) with initial condition y_0 is given by:*

$$y_t^e = (y_0^e - y^*) \left(\frac{\alpha + \lambda - 1}{\lambda} \right)^t + y^*,$$

where

$$y^* = r - g.$$

(ii) *If*

$$\lambda > \bar{\lambda} \equiv \frac{1 - \alpha}{2}, \tag{3.4}$$

y_t^e converges to the steady state y^* otherwise y_t^e diverges to $\pm\infty$.

Notice that whenever the system converges to its steady state y^*, also the dividend yield y^t converges to y^* through the feedback map (2.11). At the steady state y^* the price equals the fundamental price p_t^* defined in (2.7) and thus fully reveals the information concerning the future dividend. The shaded area in Fig. 1 shows the stability region of (3.3) in the parameter space (α, λ) whereas the white area shows the unstable region. The shaded area is divided in two different regions with different gray scales. In the lighter region, the convergence of the expected dividend yield to its the steady state y^* is oscillatory, whereas in the darker, the convergence is monotone. Notice that the border between the stability and the instability region is characterized by oscillatory behavior of the expected dividend yield y_t^e, and thus of the realized dividend yield y_t too. This implies that failure of the price to fully reveal the fundamental information should be characterized by price fluctuations with negative autocorrelation. This statement is made more precise in the analysis of the 2-D model that follows.

3.2 Dividend yield and fractions dynamics

In general, the fraction of informed agents λ_t is time dependent and the dynamics of the dividend yield and of the fraction of informed agents is nonlinear. The overall system is given by (3.1-3.2). We use local stability analysis to characterize the behavior of the state variables (y_t^e, λ_t) near the steady state of (3.1-3.2). The following proposition characterizes the steady state of the system and its local stability.

Proposition 2. *Given the memory parameter $\alpha \in (0, 1)$, the experimentation level $\delta \in (0, 1)$, the ratio between the speed of adjustment and the cost of being informed, k, such that $k = \rho/c > 1$ and the required rate of return $r > g$, we have:*
(i) *The point (y^*, λ^*) where*

$$y^* = r - g$$

and

$$\lambda^* = \frac{2 - \delta + 2k\delta - \sqrt{-8k\delta + (2 - \delta + 2k\delta)^2}}{4}, \tag{3.5}$$

is the unique steady state of the system (3.1-3.2). Moreover, $\lambda^* \in (0, 1/2)$.
(ii) *The Jacobian of (3.1-3.2) at the steady state is diagonal and given by*

$$\mathbf{J}|_{(y^*,\lambda^*)} = \begin{pmatrix} \dfrac{\alpha+\lambda^*-1}{\lambda^*} & 0 \\ 0 & (1-\delta)\dfrac{k(k-1)}{(k-\lambda^*)^2} \end{pmatrix}. \tag{3.6}$$

If

$$\delta > \bar{\delta} \equiv \frac{(1+\alpha)}{1+\frac{2\alpha k}{(1-\alpha)}} = \frac{(1+\alpha)}{1+\frac{2\alpha\rho}{(1-\alpha)c}}, \tag{3.7}$$

the steady state (y^*, λ^*) *is locally stable. The previous condition corresponds to the stability condition (3.4) of the* 1*-D dynamical system (3.3). That is whenever* $\delta > \bar{\delta}$ *then*

$$\lambda^* > \bar{\lambda}.$$

and vice-versa.

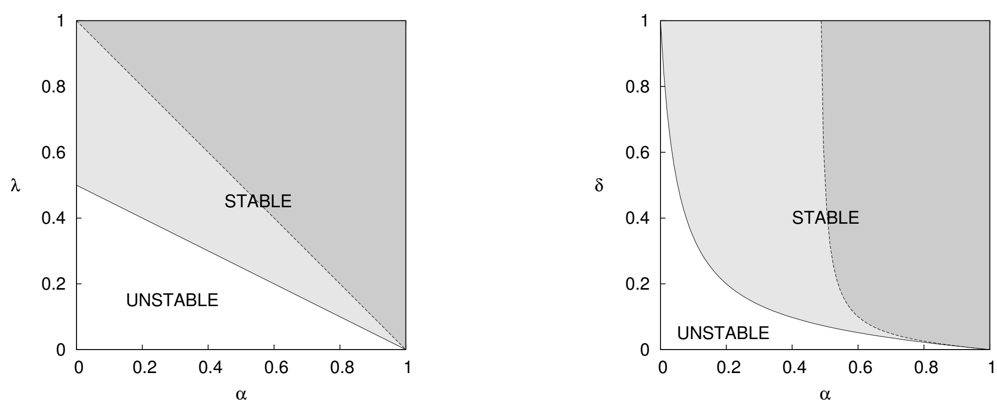

Figure 1: Left panel: Stability region for the 1-D dynamical system in (3.3). The expected dividend yield converges to the steady state $y^* = r - g$ only for values of (α, λ) in the shaded area. In the darker region the convergence to y^* is monotone, whereas in the lighter the convergence is oscillatory. Right panel: Stability region in the 2-D dynamical system in (3.1-3.2) as a function of the mutation rate δ and memory α, when $k = \rho/c = 10$. As in the left panel, the continuous line marks the border of the stability region, and the dotted line marks the border of the region where the convergence of the expected dividend yield to y^* is monotone (darker region) or oscillatory (lighter region).

The local stability condition (3.7) is represented in terms of the parameters (α, δ) in the right panel of Fig. 1 for $k = \rho/c = 10$. In the white area the steady state (y^*, λ^*) is unstable, while in the shaded area the steady state is stable. Since when the expected dividend yield y^e_t converges to y^* also the realized dividend yield y_t converges to y^*, stability of the steady state (y^*, λ^*) implies convergence of teh price to the fundamental price p_t^* and thus to a fully informative price. This result is not contradicting that obtained by GS, which is that markets cannot be informationally efficient. In fact it has been obtained by distorting the dynamics of the fraction of informed agents, i.e. by assuming a sufficiently large experimentation rate δ. Under this condition there always is a fraction of agents that are prepared to buy fundamental information. If, on the other hand, $\delta = 0$, then $\lambda^* = 0$ and y^* is not defined. In general, for any $\alpha \in (0, 1)$ there exists a sufficiently small mutation rate such that the prices are not fully revealing and the system is unstable. The definition of $\bar{\delta}$ in eq. (3.7) of Proposition 2 shows that, for a given α, the stability region of the system (3.1-3.2) shrinks, when the cost of information c increases, and expands when the speed of adjustment ρ increases.

What happens to the dynamics of the expected dividend yield and of the fraction of informed agents when the steady state is unstable? In order to answer this question we analyze the global dynamics of the system (3.1-3.2) for small experimentation levels, $\delta < \bar{\delta}$. When the stability conditions (3.4) and (3.7) do not hold, whereas in the 1-D system the expected and realized dividend yield diverge unboundedly and unrealistically, in the 2-D system our simulations show the emergence of bounded aperiodic cycles. The top left and top right panels of Fig. 2 show the typical evolution of the uninformed agents' expected dividend yield, y^e, and of the fraction of informed agents, λ, respectively. At time $t = 0$, the fraction of informed agents is above the dotted line, which marks the value $\bar{\lambda}$ in (3.4) and gives the stability condition for the steady state of the y_t^e dynamics. As a result, at $t = 0$ both the value of y_t^e and, through (2.11), the value of y_t, are close to their steady state value, y^*. This implies that the price is close to being fully informative, there is no advantage in buying information so that the fraction of informed agents decreases. This process continues until the fraction of informed agents is smaller than the value $\bar{\lambda}$. At this moment there are so few informed agents that the asset price starts to diverge from the fundamental. The dynamics of the expected dividend yield y_t^e is unstable and y_t^e starts to diverge from the steady state y^*. The fraction of informed agents continues to decrease until the price carries so little information about p^* that informed agents are better off. Eventually, it gives a higher fitness to pay the cost of being informed than to use a freely available estimate with a large error. As a result, the fraction of informed agents grows sharply, see e.g. the top right plot around period $t = 50$. The fraction of informed agents reverts to a region where the price is sufficiently informative so that y_t^e returns to values close to y^*. As time flows the process repeats, with λ decreasing again, and so on and so forth. The left and right bottom panels of Fig. 2 show, respectively, the dynamics we have just illustrated in the (y^e, λ) space and the corresponding dynamics of the log price as compared to the log fundamental price.

The numerically obtained phase plots shown in Fig. 3 suggest that the fluctuations of y^e and λ just described, are associated with a so-called homoclinic bifurcation. Similar phenomena are encountered in other multidimensional nonlinear systems and emerge from the interplay between local instability and global stability of the dynamics. Brock and Hommes (1997) and Droste *et al.* (2002) present other economic frameworks where homoclinic bifurcation arise. They also offer detailed discussions of the mathematical aspects behind these interesting phenomena.

Before turning to the economic interpretation of the fluctuations and to the comparison of our results with those of GS, it is instructive to characterize the convergence of the fraction of informed agents in the stability region. Close to the equilibrium λ turns out to change very

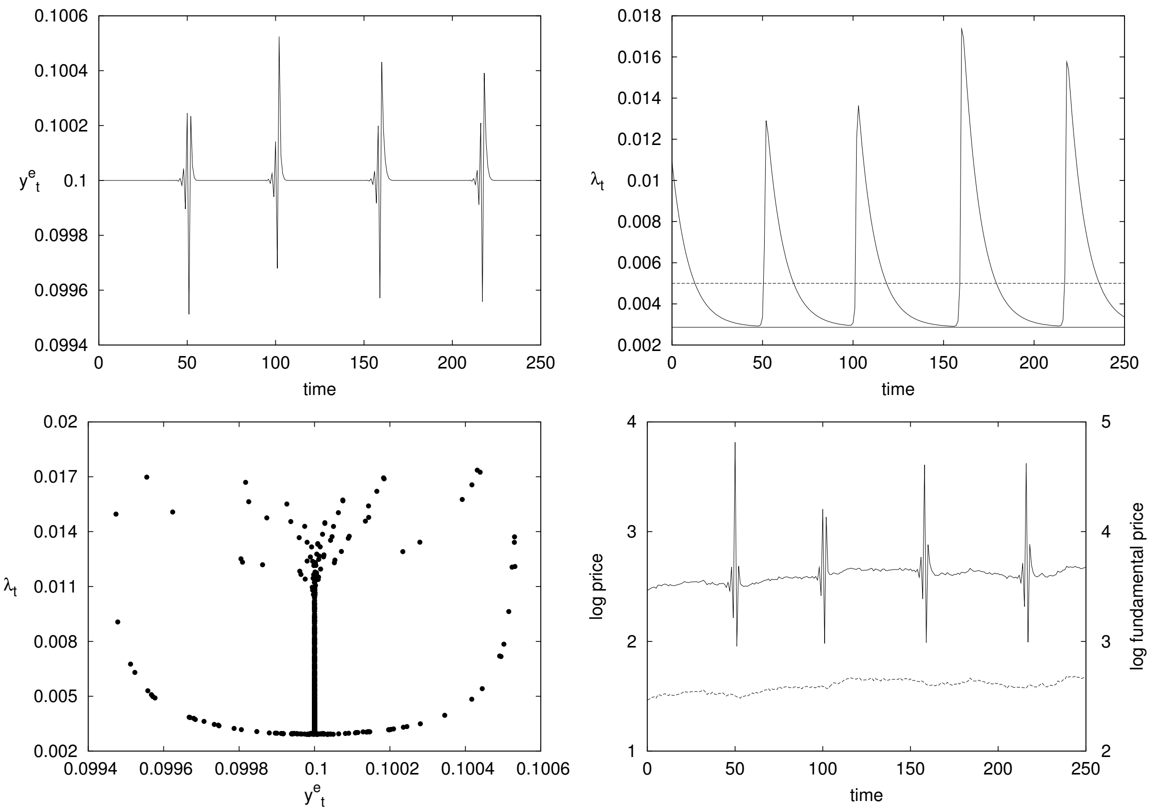

Figure 2: Top panels: time series for the expected dividend yield y^e (left panel) and the fraction of informed agents λ (right panel) produced by (3.1-3.2). In the top right panel the dotted line corresponds to $\bar{\lambda}$ in (3.4), that is to the stability value of λ, whereas the continuous line corresponds to the steady state λ^* given in (3.5). In the bottom left panel, the state space representation of the previous two times series. In the bottom right panel, the log price dynamics derived form the dynamics of y^e and λ . Fundamental prices are given by the dotted line and realized prices by the continuous line. Notice that the two price series levels should be read using two different scales. The left scale gives the value of log prices whereas the right scale gives the value of fundamental log prices. The dividend process is characterized by $\sigma_\eta = 0.02$. Parameter values are $\alpha = 0.99$, $c = 0.1$, $\rho = 1$ (so that $k = \rho/c = 10$), $r = 0.1$, $g = 0$ and $\delta = 0.000575$.

slowly.

Proposition 3. *Given a memory parameter $\alpha \in (0, 1)$, an experimentation rate $\delta \in (0, 1)$ and a speed of adjustment ρ larger than the cost of information, $\rho > c$, if we call ν_2 the eigenvalue characterizing the dynamics of λ in a neighborhood of (y^*, λ^*), we have*

$$1 > \nu_2 > (1 - \delta)\left(1 - \frac{1}{k}\right).$$

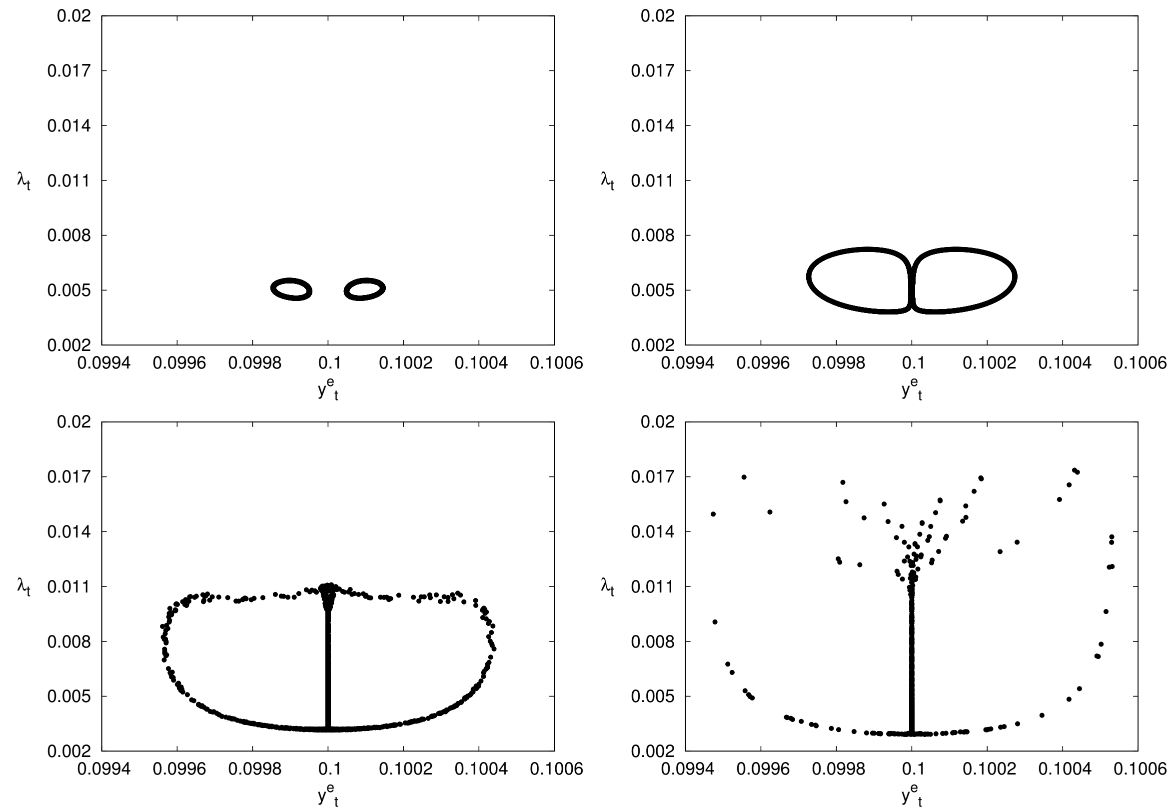

Figure 3: Phase plot of the expected dividend yield y^e and the fraction of informed agents λ produced by the system of difference equations (3.1-3.2). Parameter values are $\alpha = 0.99$, $c = 0.1$, $\rho = 1$ (so that $k = \rho/c = 10$), $r = 0.1$ and $g = 0$. Top left panel: $\delta = 0.000675$, top right panel: $\delta = 0.00066$. Bottom left panel: $\delta = 0.000625$, bottom right panel: $\delta = 0.000575$. At an intermediate value of δ a homoclinic bifurcation occurs.

This proposition shows that when the experimentation rate δ is small and the ratio between the speed of adjustment and the cost for information, k, is big, the value of ν_2 is very close to one. As a result, when the system is stable, changes in the value of the fraction of informed agents λ are very slow, and hence λ is very persistent. We will find confirmation of this statement in the next Section, when we feed our system with exogenous shocks on the dividend growth rate.

3.3 Economic interpretation

The use of rational expectations in a negative feedback framework poses puzzling consequences when one endogenizes the dynamics of the fraction of informed agents λ. In fact when the information cost is positive, if agents had rational expectations the price would fully reveal the available information about future dividends and nobody would pay for information. This implies that the fraction of informed agents would tend to zero. In the limit the price would not contain information about the dividend anymore. This is the same paradox as found by GS, who also consider an asset pricing model where agents can either buy information on fundamentals or try to extract such information from the asset price. In a repeated single period model with rational agents, they show that there cannot exist an equilibrium value of the fraction of informed agents for which the price fully reveals the information about the future dividend. If an equilibrium existed then nobody would pay for the information and prices could not possibly reveal any information. They also show that there exists a rational expectations "equilibrium degree of disequilibrium" where prices fail to be fully informative. In order to get this result they need two key assumptions: (1) the supply of shocks is noisy (this is equivalent to saying that there are noisy traders in the market) and (2) the informational content of the dividend signal is not perfect. Without these two assumptions, an equilibrium degree of disequilibrium would not exist.

Our model is inspired by that of GS but with three important differences. First, our agents are not rational but boundedly rational, that is, they do not use rational but adaptive expectations. Second, we consider a multi period model where future returns are determined by capital gains in addition to dividends, and agents form expectations about both future prices and dividends. Third, the fraction of informed and uninformed agents are dynamic variables of our model. By assuming that agents have boundedly rational expectations and that fractions are endogenously determined, we offer another source for balancing the cost of information and the informational content of the price namely the learning process of the uninformed agents. In our framework we obtain a "dynamic equilibrium degree of disequilibrium" due to endogenous price fluctuations produced by the interaction of boundedly rational agents.

De Fontnouvelle (2000) and Goldbaum (2005) are earlier contributions where bounded rationality and learning offer a different explanation for the existence of an equilibrium degree of disequilibrium. Their framework differs from ours in many ways, most importantly in that they consider a dividend process which follows a random walk rather than a geometric random walk as we do here. Furthermore their resulting systems of the joint evolution of the asset price and the fractions of agents is fairly complicated so that their analysis is performed only via simulations.

An important characteristic of our "dynamic equilibrium degree of disequilibrium" is that the system adjusts itself to a pattern of a-periodic oscillations where the prices are not fully informative. In our framework the price informational content is time varying and switches continuously between being nearly fully informative and hardly informative. The economic intuition behind this is quite clear: when the fraction of informed agents is high enough, we are in a region where the adaptive expectations of boundedly rational uninformed agents are converging to rational expectations and the price is close to being fully informative. This pushes down the fraction of informed agents. As a result, with only few agents being informed, the uninformed agents using adaptive expectations do not converge to rational expectations anymore and the price starts to diverge from its fundamental value and, as a consequence, carries little information. This creates incentives to buy information and pushes the fraction of informed agents up again, and the story repeats. This trade-off between local instability (when too few agents are informed) and global stability (when many agents are informed) leads to complicated dynamic behavior. In our model, the "dynamic equilibrium degree of disequilibrium" is therefore a time-varying learning equilibrium where prices fluctuate between being close to fully revealing and being uninformative, and agents switch between costly information gathering and free riding.

4 Informational differences and parameter learning

In the previous section we have assumed that the agents' estimate of the dividend growth rate is constant. As a result, the equation for the dividend yield is fully deterministic. In this section we analyze the simultaneous impact of informational differences and of assuming that agents are learning the growth rate of dividends as new information about the fundamentals becomes available. As a result we have to deal with a stochastic system. A similar analysis has been performed by Barsky and De Long (1993) and Timmermann (1993), among others, in a context where there are no informational differences among agents. In particular Barsky and De Long (1993) also assume that agents use adaptive expectations to estimate g. Recalling the results from Section 2, adaptive expectations are specified by (2.13) which, when the dividend follows a geometric random walk with innovations $(1+g)\eta_t$ gives:

$$g_{t+1}^e = \gamma g_t^e + (1-\gamma)\left((1+g)\eta_t - 1\right). \tag{4.1}$$

This stochastic equation, together with the evolution of the dividend yield, its expectations, and the fraction of informed agents as specified in (2.26-2.29), lead to a stochastic version of the deterministic skeleton (3.1-3.2) namely

$$\begin{aligned} g_{t+1}^e &= \gamma g_t^e + (1-\gamma)\left((1+g)\eta_t - 1\right), && (4.2)\\ y_{t+1}^e &= \frac{(1-\alpha)(r-g_t^e)}{\lambda_{t-1}} + \frac{\alpha+\lambda_{t-1}-1}{\lambda_{t-1}} y_t^e, && (4.3)\\ \lambda_t &= (1-\delta)\frac{\lambda_{t-1}(-c+\rho)}{\lambda_{t-1}(-c) - (1-\lambda_{t-1})\left(\frac{(r-g_t^e)-y_t^e}{\lambda_{t-1}(r-g)}\right)^2 + \rho} + \frac{\delta}{2}. && (4.4)\end{aligned}$$

Shocks $\{\eta_t\}$ on the growth rate of dividends are the stochastic component that drives the co-evolution of agents expectations of the growth rate of dividend and of the dividend yield, and of the fraction of informed agents. Given the evolution of the expected growth rate of dividend, g_{t+1}^e, of the expected dividend yield, y_{t+1}^e, and of the fraction of informed agents, λ_t, the dividend yield itself, y_{t+1}, is set by the feedback map (2.11). Before we start with the analysis of the impact of shocks on the dynamics of (4.2-4.4), we show that our model contains two important benchmarks.

Classical Asset Pricing model The first benchmark is the classical asset pricing model, which assumes that all agents know the long run dividend growth rate g, and that agents use rational expectations. In this case, if some agents are informed about d_{t+1}, the market price and the market dividend yield are given by:

$$p_t^{**} = \frac{d_{t+1}}{r-g}, \quad y_{t+1}^{**} = r - g. \tag{4.5}$$

Previously we have called the price p_t^{**} the "correct" present value price. Notice that the classical asset pricing model corresponds to our model if we assume that agents have rational expectations and use the correct long run dividend growth rate g. Alternatively the classical asset pricing model corresponds to our model when all agents are informed and use the correct long run dividend growth rate.

Barsky and De Long model The second benchmark is the model proposed by Barsky and De Long (1993). They consider agents without informational differences who have to form expectations about the long run growth rate g. In their case the price and the dividend yield are given by:

$$p_t^* = \frac{d_{t+1}}{r - g_{t+1}^e}, \quad y_{t+1}^* = r - g_{t+1}^e, \tag{4.6}$$

where g_{t+1}^e is given by adaptive expectations as in Eq. (4.1). These are respectively the fundamental price and the fundamental dividend yield defined in (2.14) and (2.15). For a given g_{t+1}^e, y_{t+1}^* is also the steady state of the system where uniformed agents are learning the value of the dividend yield, i.e. to the point where price are fully informative. The model of Barsky and De Long (1993) thus corresponds to our model under the assumptions that all the agents use rational expectations or that all agents are informed.

A way of comparing our model with these two benchmarks is to write an evolution equation for the dividend yield as a function of lagged dividend yields and shocks on the growth rate of dividends for each model. In the classical asset pricing model the dividend yield is constant and given by $y^{**} = r - g$. In the Barsky and De Long model, using (4.6) together with (4.1), one can derive:

$$y_{t+1}^* = (1-\gamma)(1+r) + \gamma y_t^* - (1-\gamma)(1+g)\eta_t. \tag{4.7}$$

That is, the dividend yield follows an AR(1) process with shocks given by the shocks $\{\eta_t\}$ on the growth rate of dividends. The memory parameter γ is related to both the AR(1) coefficient and the variance of the innovations, $(1-\gamma)^2(1+g)^2\sigma_\eta^2$. The mean of the process is independent of the memory parameter and equal to the constant classical asset pricing dividend yield $y^{**} = r - g$. In our model, fixing for the moment the value of λ, the map h defined in (2.18) would give give:

$$y_{t+1} = \frac{(r_t - g_{t+1}^e) - \alpha(r_{t-1} - g_t^e)}{\lambda} + \frac{\alpha + \lambda - 1}{\lambda} y_t. \tag{4.8}$$

This equation expresses that, when the fraction of informed agents fixed at λ, the dividend yield follows an AR(1) process with shocks that are correlated with the shocks of the growth rate of dividends. When $\gamma = \alpha$, that is, when agents use the same memory parameter to estimate the growth rate of dividend an the dividend yield, Eq. (4.8) has a simple appealing formulation:

$$y_{t+1} = \frac{(1-\gamma)((1+r) - (1+g)\eta_t)}{\lambda} + \frac{\gamma + \lambda - 1}{\lambda} y_t.$$

As for the dividend process implied by the Barsky and De Long model in (4.8), the long run mean of the dividend process implied by our model is given by the classical asset pricing dividend yield y^{**}. If we define

$$\gamma'(\lambda) = 1 - (1-\gamma)/\lambda = \gamma - (1-\gamma)\frac{1-\lambda}{\lambda}, \tag{4.9}$$

we can rewrite (4.8) as

$$y_{t+1} = (1-\gamma'(\lambda))(1+r) + \gamma'(\lambda)y_t - (1-\gamma'(\lambda))(1+g)\eta_t. \tag{4.10}$$

The result is that, when $\gamma = \alpha$, our model specified by Eq. (4.10) and the model of Barsky and De Long (1993) specified by Eq. (4.7) differ only in the value of the memory parameter γ. Since γ is the real memory agents use to discount new information, we can refer to $\gamma'(\lambda)$ as the effective memory. The definition (4.9) shows that the effective memory has two components, one being given by the real memory and the other by the effect the action of uninformed agents imposes a negative feedback on the evolution of y. This second effect becomes less important as more informed agents are present in the market. The general result is that $\gamma'(\lambda)$ in (4.9) is an increasing function of λ with $\gamma'(\lambda) \leq \gamma$ for all λ, and $\gamma'(1) = \gamma$. That is, the presence of uninformed agents is equivalent to the all agents being informed and using an effective memory which is lower then the real memory agents use. The value of γ' determines both the AR(1) coefficients and the variance of the shocks but not the long run mean which is always $y^{**} = r - g$. In particular the lower the effective memory, the higher the impact of the shocks on the dynamics of the dividend yield and the faster the reversion of the process to its mean. That is, a lower effective memory creates a bigger short run effect and a smaller long run effect. Also, since the effective memory γ' is a function of λ, our model allows for variation of the memory parameter as the fraction of informed agents λ varies. Changes in λ have an impact on γ' and thus on the variance of shocks and on the speed of convergence. In what follows we explore the importance of both the effective memory being lower than the real memory and the effective memory being time varying on the dynamics of the dividend yield implied by our model (4.10) compared to its two benchmarks in (4.5) and (4.8).

4.1 Nonlinear mean reversion

Our model (4.10) clearly differs, both structurally and regarding the parameters, from that one of Barsky and De Long in (4.7) when the fraction of informed agents is time varying. If this is the case, our model implies an AR(1) for the dividend yield where both the rate of convergence of the dividend yield to its mean and the variance of shocks are time varying. This consideration links our model to the econometric analysis of nonlinear mean reversion that has recently been proposed to characterize the fluctuations of stock indices. In fact by using the fact that $p_t^{**} \equiv d_{t+1}/(r-g)$, i.e. the price implied by model (4.5), and defining $x_t = y_{t+1}/(r-g)$, given the definition of the dividend yield one can write:

$$\log(p_t) = \log(d_{t+1}) - \log(y_{t+1}) = \log(p_t^{**}) - \log(x_t).$$

If x_t is close to its long run average of 1 one can rewrite the previous expression and expand the logarithm around one. Using the variable $z_t = 1 - x_t$ one gets

$$\log(p_t) \approx \log(p_t^{**}) + z_t, \tag{4.11}$$

where the dynamics of z_t can be easily derived using its definition in terms of x_t, the definition of x_t and Eq. (4.10). The resulting dynamics of the component z_t of the log price is given by:

$$z_t = \gamma'(\lambda_t) z_{t-1} + \frac{(1-\gamma'(\lambda_t))(1+g)}{(r-g)} (\eta_{t-1} - 1). \tag{4.12}$$

This equation shows that we have a model whose realized log price in (4.11) is the sum of a persistent component $\log(p_t^{**})$, which follows a random walk with drift, and of a temporary component, z_t, which follows a stationary autoregressive process (4.12) with a time-varying AR(1) coefficient $\gamma'(\lambda_t)$. Empirical investigation of the properties of stock prices are in accord with this statement. Both Gallagher and Taylor (2001) and Manzan (2003) reject the null hypothesis that the temporary component in a mean reversion model follows a stationary process with fixed parameters. In particular Gallagher and Taylor (2001) show that quarterly data of the logarithm of the dividend yield of the index SP500 are well fitted by an ESTAR(4) (Exponentially Smooth Transition AR) ARCH(1) model whose two regimes have AR(1) coefficients equal to 0.72 and 0.20 respectively. As the model of Barsky and De Long (1993) suggests, the fact that the dividend yield follows an autoregressive process might be related to the agents' learning of the model parameters. In addition to this effect, our model suggests that changing "learning" coefficients and heteroskedasticity can be related to agent interaction. In fact, both the AR(1) coefficient and the shocks variance in (4.12) are a function of $\gamma'(\lambda_t)$ which is a nonlinear function of the time varying fraction of informed agents λ_t.

It is beyond the scope of this paper to calibrate our model to reproduce the stock prices evolution given the historical dividend process. Our theoretical model is based on several simplifying assumptions and in particular on an ad hoc dynamics of the fraction of informed agents λ. Nevertheless we find it instructive to note that the nonlinearity in a mean reversion model can be related to what in general may be referred to as agents' interaction, which in, our case, is triggered by informational differences. That agent interaction can be responsible for nonlinearity in the behavior of stock prices is also argued by Boswijk *et al.* (2005), who estimate a modified version of the model of Brock and Hommes (1998) using yearly data of the index SP500. Further efforts to characterize the effect of informational differences, for example to link it to other observable characteristics as the volume of transactions, might offer insight to the design of new econometric tests for the evolution of stock prices.

4.2 Simulation study

In presenting the qualitative effect of the shocks on the growth rate of dividends on the dividend yield and fraction dynamics, we proceed by analyzing the impact of a single shock η_t, and then by analyzing the cumulative impact of a sequence of shocks $\{\eta_t\}$. We present results not only for the dividend yield and price generated by our model, but also for the dividend yields and prices generated by the classical asset pricing model (4.5) and by the model of Barsky and De Long (4.6). In addition we also present results for a model similar to that of Barsky and De Long with the difference that the real memory is taken as the average memory of the time varying effective memory generated by our model. We name this model as "modified" Barsky and De Long model and its series of prices and dividend yields as $p^*_{\gamma'}$ and $y^*_{\gamma'}$. We use the "modified" Barsky and De Long model to appraise the role of time variability of the effective memory. We simulate all the models with dividends generated according to Assumption (i), that is, $d_{t+1} = d_t(1+g)\eta_{t+1}$, where $\{\eta_t\}$ is a sequence of i.i.d. log normal shocks with mean zero and variance σ_η.

It is instructive to start the analysis by comparing the effect of a single shock on g and on the realized price that is respectively on the price p_t implied by our model, on the fundamental price p^*_t implied by the model of Barsky and De Long, and on the "correct: present value price implied by the classical asset pricing model. Fig. 4 shows that in both cases there is an initial overreaction followed by convergence to the equilibrium value, which is given by p^{**}, the price implied by the classical asset pricing model.

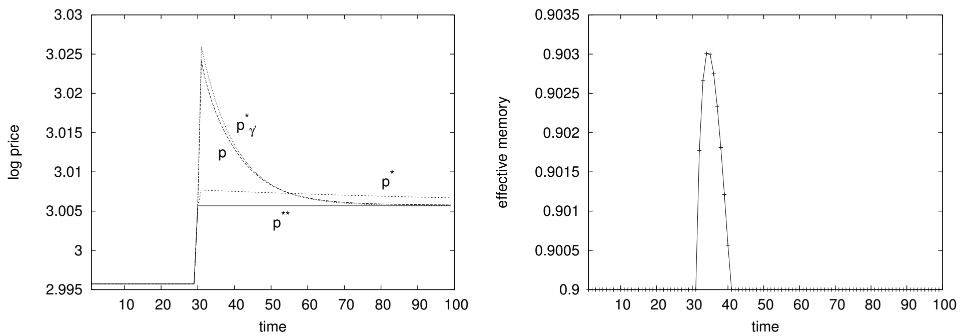

Figure 4: Effect of a single shock. Before and after the shock, $g = 0$. The shock is $\eta_{31} = 0.01$. Left panel: log prices as a function of time. The time series p^{**} gives the value of the log price as implied by the classical asset pricing model (4.5). The time series p^* and p give respectively the value of the log price as implied by Barsky and De Long model (4.7) and by our model (4.2-4.4) respectively. The time series $p^*_{\gamma'}$ gives the log price implied by the " modified" Barsky and De Long model (4.7) when $\gamma = \gamma'(\lambda_0)$, where λ_0 is the fraction of informed agents before the shock. Right panel: evolution of the effective memory γ'. The parameter values are $\gamma = \alpha = 0.99$, $\delta = 0.02$, $c = 0.1$ and $\rho = 1.0$.

Since, for $\lambda < 1$, the effective memory γ' is lower than the real memory γ, the variance of the shocks is larger in our model than in the model of Barsky and De Long, so that the overreaction is more pronounced. At the same time, when $\gamma' < \gamma$, the value of the autoregressive coefficient is closer to zero so that convergence is faster. The overall effect is that the shock has a higher short run impact but a shorter half life for p than for p^*. The right panel of Fig. 4 shows the response of the effective memory γ' to changes in λ. From the Jacobian of the 2-D system (see Proposition 2) we know that changes in y only have second order effects on λ, and as a result changes in λ are negligible in the short run. But, from Proposition 3, we also know that the eigenvalue ν_2 is close to one so that changes in λ are very persistent. Both results are confirmed by the changes in γ' shown in the right panel. A confirmation of the fact that one shock has no considerable consequence on changes of γ' comes from the time series for $p^*_{\gamma'}$ shown in the left panel. The price $p^*_{\gamma'}$ is the price obtained using the "modified" Barsky and De Long model, that is equation (4.7) with $\gamma = \gamma'(\lambda_0)$, where λ_0 is the value of the fraction of informed agents before the shock.

The overall comparison of the dynamics of p, p^* and $p^*_{\gamma'}$ shows that in the single shock case the fact that the effective memory is lower than the real memory plays an important role whereas the fact that the effective memory is time varying is negligible, i.e. p is close to $p^*_{\gamma'}$. Notice also that with informed agents in the market, the price anticipates the shock on the dividend, i.e. the price takes into account the change in the dividend before such a change is realized and much before such change has an effect on the value of the effective memory.

We now turn to investigating the effect of a sequence of shocks. Figure 5 shows the impact of a sequence of 500 i.i.d. shocks $\{\eta\}$. The long run growth rate of dividends, g, and the variance of the growth rate shocks, σ^2_η, are taken in accordance with historical quarterly data of the S&P500 index for the period 1880-2005[1]. The discount rate r is taken such that $y^{**} = 0.05$, that is the price implied by the present value model is 20 times the value of the dividend. If we think of quarters, 500 dividends correspond to 125 years. The top left panel shows the time series of the dividend yield y generated by our model whereas the right panel shows the time series of y^* generated by the model of Barsky and De Long. In both cases the horizontal line represents the long run mean $y^{**} = r - g$. The same results as for a single shock emerge: the dynamics of the dividend yield is less persistent in our model where the fraction of informed agents is time varying and smaller than in the model of Barsky and De Long (1993). Also, deviations from y^{**} are larger. The central and bottom rows offer a comparison of the systems in terms of log prices. How important is the fact that the effective memory is time varying? The right panel of Fig. 6 shows the changes in the effective memory for the same simulation run. These changes are due to changes in the fraction of informed agents λ via the transformation equation (4.9). As a confirmation of our previous results and of our theoretical analysis, changes in γ' (that is changes in λ) are rather persistent. The left panel of the same figure shows deviations of log prices generated by our model and log prices generated by the "modified" Barsky and De Long model. We call this last series $y^*_{\gamma'}$. Notice that deviations of up to more than ten percent arise. Our conclusion is that when subsequent shocks are present, both the effective memory being lower than the real memory and the effective memory being time varying play an important role. Naturally, these properties are dependent on the choice of updating mechanism for λ_t and hence of the fitness measures as presented in Subsection 2.3. We do not claim that the mechanism we propose here to characterize the changes in the fraction of informed agents is more realistic than others. We merely offer a qualitative argument to show that informational differences might explain the nonlinearity in the mean reversion that has been shown to exist in the empirical literature.

Another way of comparing the various models is to check for correlation in the time series of returns produced by the evolution of y^* and y. The left panel of Fig. 7 shows the autocorrelation of the asset log return series R_t,

$$R_t = \log(p_t + d_t) - \log(p_{t-1}), \tag{4.13}$$

for a typical run of our model (4.2-4.4). The autocorrelation of returns shows that our model and the "modified" Barsky and De Long model have higher short term autocorrelation and lower long term autocorrelation. Such results are in accordance with the results shown in the left panel

[1]Source of the data is Shiller database available from R. J. Shiller homepage.

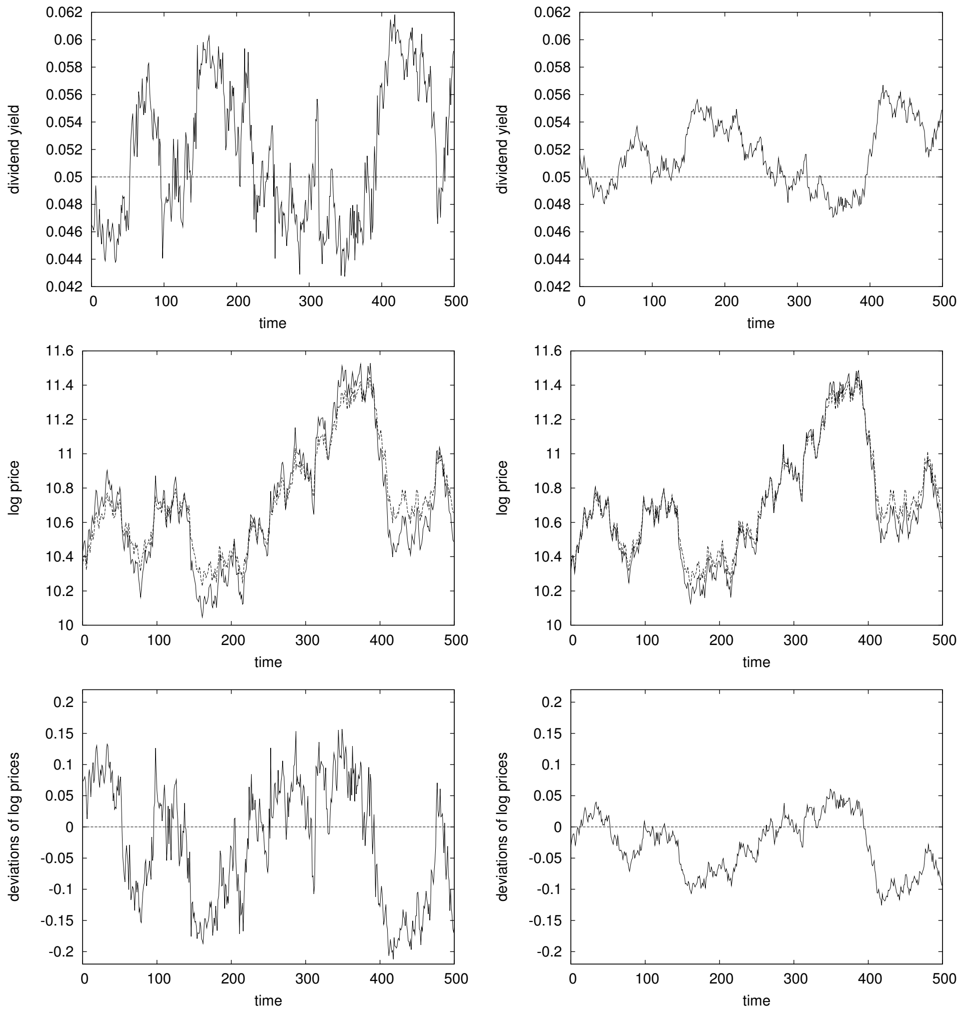

Figure 5: Top left panel: time series of the dividend yield y_t generated by our model (4.2-4.4) (solid line) compared with the benchmark y^{**} (4.5) (horizontal dotted line). Top right panel: time series of the dividend yield y^* as in Barsky and De long (4.7) (solid line) compared with y^{**}. Middle left panel: logarithm of price corresponding to y, $\log(p)$ (solid line), and logarithm of the price corresponding to y^{**}, $\log(p^{**})$ (dotted line). Middle right panel: logarithm of price implied corresponding to y^*, $\log(p^*)$ (solid line), and logarithm of the price implied corresponding to y^{**}. Bottom panels gives the deviations of the log prices series shown in the middle panels. Values of parameters are $\alpha = \gamma = 0.99$, $\rho = 1$, $c = 0.1$, $\delta = 0.02$ (these three parameters imply $\lambda^* \approx 0.09$), $\sigma_\eta = 0.04$, $\mu = 0.003$. The discount rate is $r = 0.05 + g$.

of Fig. 4: if the effective memory is lower than the real memory, shocks have a higher short run impact but a shorter half life for p then for p^*.

A test that has been used in the literature to evaluate the statistical importance of departure of the models from a random walk with drift is the variance ratio test. The variance ratio has

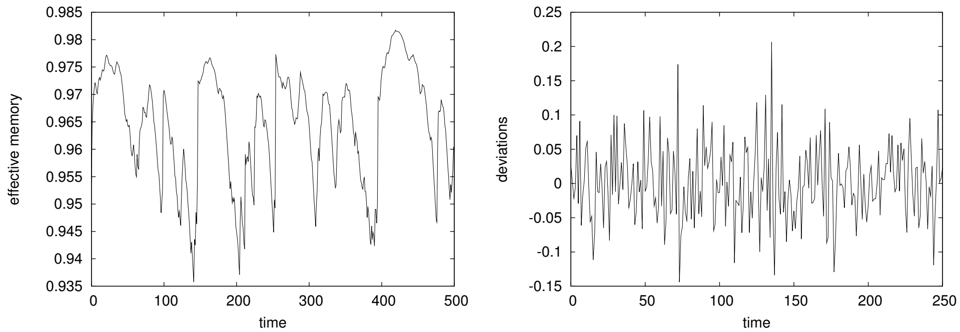

Figure 6: Left panel: values of the effective memory γ'. Right panel: deviations of the two series of log prices generated by our model y and by the "modified" Barsky and De Long model $y^*_{\bar{\gamma}'}$. Deviations are due to fluctuations of γ' around its mean value, $\bar{\gamma}' = 0.9644$. Other parameters as in Fig. 5, in particular the real memory is $\gamma = 0.99$.

been used by Poterba and Summers (1988) to appraise the mean reversion properties of stock prices. Under the null hypothesis that log prices follows a random walk (possibly with drift) the variance of the series of returns in (4.13) is a linear function of the return time span. Results of the variance ratio test for our model and for its restrictions are given in the right panel of Fig. 7. The results suggest that not only the fact that the effective memory is lower than the real memory affects the statistical time series properties of lagged returns, but also that the effective memory $\gamma'(\lambda)$ is time varying. Further research will be devoted to investigating these issues in greater detail and related them to the statistical properties of financial markets empirical returns.

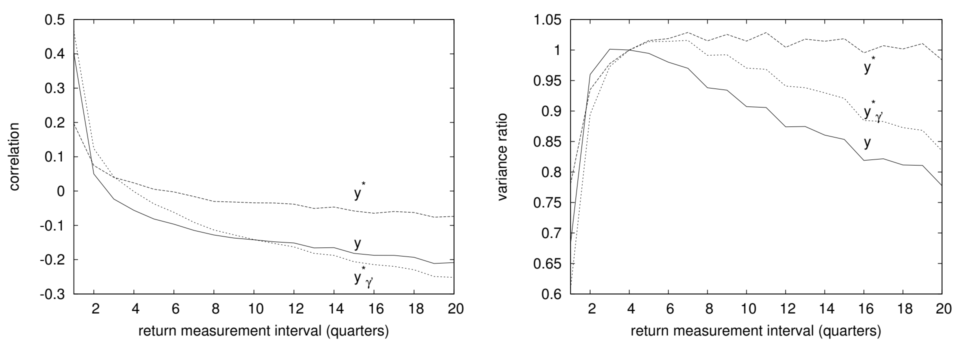

Figure 7: Left panel: autocorrelation of the log return tme series for a typical run of our model (4.2-4.4). Right Panel, variance test of log lagged returns. $v(\Delta t) = (\sigma^2(R^{\Delta t})/\Delta t)/(\sigma^2(R^4_y)/4)$ where $\sigma^2(R^4_y) = 0.00355)$ and $R^{\Delta t}$ is the total return over a period Δt generated by the model in system (4.2-4.4). In both plots three lines refer to data generated with our model, y, to the model of Barsky and De Long in (4.7), y^*, and to the "modified" Barsky and De Long model. Parameters are the same as for Fig. 5. Both panels refer to a simulation of $200,000$ periods.

5 Conclusion

We have built a simple model of an asset market where agents are boundedly rational and can choose between different degrees of information regarding future dividends. As far as the theoretical guidelines behind our model are concerned, we show that our model naturally and parsimoniously extends and links many other contributions in this fields. In particular, we refer to papers that concentrate on informational differences, as Grossman and Stiglitz (1980), that analyze the impact of learning, as Barsky and De Long (1993), and that investigate the interaction of agents who are using different predictor schemes or different strategies, as Brock and Hommes (1998). A first result of theoretical relevance is presented in Section 3. There we extend GS's results and show that a price "dynamic equilibrium degree of disequilibrium" can be achieved in a multi-period market populated with boundedly rational agents. A second result is given in Section 4, where we show that informational differences in a market with boundedly rational agents also have interesting empirical implications. Our framework provides insights into a number of econometric models which have recently been introduced to justify financial anomalies such as mean reversion and return predictability by assuming that the stock price is the sum of a persistent and a of temporary component. In our framework, the persistent component can be directly linked to the dividend process whereas the temporary component can be related to agents' learning of the dividend growth rate and to the agents' attempt to extract information from prices. In particular this second effect shows how a time-varying coefficient of the temporary component, and thus nonlinear mean reversion, can be explained.

Appendix

A Equilibrium price equation

This appendix provides a micro-foundation of the equilibrium price equation (2.4). Consider a group of agents choosing at every time t whether to invest in a risk free asset, whose single period return is r_f, or to invest in a risky asset, whose single period return rate, r_{t+1}, depends on dividend payed at time $t+1$, d_{t+1}, and on the price (or remaining value) of the asset at time $t+1$, p_{t+1}:

$$r_{t+1} = \frac{p_{t+1} + d_{t+1} - p_t}{p_t}. \tag{A.1}$$

At every time t, we assume that each agent maximizes a CRRA mean variance utility function in order to decide which fraction x_t of his wealth to invest in the risky asset. The CRRA utility function to be maximized is

$$U(x_t) = E_t[x_t r_{t+1} + (1 - x_t) r_f] - \frac{\beta}{2} V_t[x_t r_{t+1}],$$

where E_t and V_t denote, respectively, the mean and the variance conditional on the information available at time t and β is the coefficient of risk aversion which we assume constant across agents. Assume also that for each agent $V_t[r_{t+1}] = \sigma^2$. The solution of the maximization of $U(x_t)$ gives:

$$x_t = \frac{E_t[r_{t+1} - r_f]}{\beta\sigma^2}, \tag{A.2}$$

as the fraction of wealth to be invested at time t. Consider now the case where assumptions $(i) - (iv)$ of Section 2 hold. Since informed and uninformed agents have different information, they have different $E_t[r_{t+1}]$ and therefore different demands for risky assets. Call x_t^I (x_t^U) and w_t^I (w_T^U), respectively, the fraction of wealth to be invested at time t and the wealth at time t of the informed (uninformed) agent. Assume a net positive supply of shares s_t and call p_t the price of each share at time t. The Walrasian equilibrium equation at time t is given by:

$$s_t p_t = x_t^I w_t^I + x_t^U w_t^U. \tag{A.3}$$

Now, define λ_t as the fraction of wealth, or market power, of the informed agents at time t. This implies $\lambda_t = w_t^I / w_t$ with $w_t = w_t^I + w_t^U$. Call θ_t the average proportion, at time t, of wealth invested in the risky assets, i.e. $\theta_t = s_t p_t / w_t$. The Walrasian equilibrium equation (A.3) becomes

$$\theta_t = \lambda_t x_t^I + (1 - \lambda_t) x_t^U,$$

which, given the expression for x in (A.2) and rearranging terms, becomes

$$r_f + \theta_t \beta\sigma^2 = \lambda_t E[r_{t+1} | \mathcal{F}_t^I] + (1 - \lambda_t) E[r_{t+1} | \mathcal{F}_t^U], \tag{A.4}$$

where the informed and uniformed agents condition their expectations of the return of the risky asset on different information sets. Eq. (A.4) shows that θ_t is related to the weighted average of the risk premium required by the community of traders to hold the asset.

If we assume that both informed and uninformed agents are fundamentalists, their expectations of the future price is equal to the discounted sum of all future dividends, i.e., as specified in equation (2.1),

$$E\left[p_{t+1} \middle| \mathcal{F}_t^H\right] = v_{t+1}^H = E\left[\sum_{i=1}^{\infty} \frac{d_{t+1+i}}{(1+r)^i} \middle| \mathcal{F}_t^H\right].$$

for the general information set $\mathcal{F}_t^H$. This implies that the informed agents use:

$$E[r_{t+1} | \mathcal{F}_t^I] = \frac{d_{t+1}(1+r)}{(r - g^e) p_t} - 1, \tag{A.5}$$

whereas the uninformed agents use:

$$E[r_{t+1} | \mathcal{F}_t^U] = \frac{p_t y_{t+1}^e (1+r)}{(r - g^e) p_t} - 1. \tag{A.6}$$

As a result the equilibrium equation (A.4) becomes:

$$p_t \frac{1 + r_f + \theta_t \beta\sigma^2}{1 + r} = \lambda_t \frac{d_{t+1}}{(r - g^e)} + (1 - \lambda_t) \frac{p_t y_{t+1}^e}{(r - g^e)}. \tag{A.7}$$

At this point, by fixing $r - r_f = \theta_t \beta \sigma^2$, that is by imposing that the asset excess return required by the agents, which is an exogenous variable of the model, is equal to the endogenous variable $\theta_t \beta \sigma^2$, we get Eq. (2.4) which solved for p_t gives Eq. (2.10).

By fixing $r - r_f = \theta_t \beta \sigma^2$ we are implicitly assuming that the endogenous expected equilibrium return of our model is given by the exogenous parameter r. In fact, when the price is informationally efficient, the resulting expected and realized dividend yield are equal to $y^* = r - g^e$ so that using Eq. (A.6), or equivalently (A.5), to compute the expected equilibrium asset return we get:

$$E[r_{t+1}|\mathcal{F}_t^U] = E[r_{t+1}|\mathcal{F}_t^I] = \frac{y^e(1+r)}{(r-g^e)} - 1 = r.$$

In order to derive the asset return, r_t, endogenously as a function of the exogenous parameters r_f, σ^2, β, s_t and of the discount rate r, one should solve the equilibrium price equation (A.7), without fixing $s_t p_t / w_t$ proportional to $r - r_f$. Levy *et al.* (1994) are, to our knowledge, the first to perform this kind of analysis. They use computer simulations to investigate the evolution of wealth and prices in an asset market where, as in our framework, agents are using CRRA utility function and the underlying dividend process follows a geometric random walk. In a recent paper Anufriev and Dindo (2006) offer analytic support of their simulations. In the present paper, having assumed that $E_t[r_{t+1}] = r$, we are fixing the long run asset return and we concentrate on the properties of the fluctuations induced by agents interaction around this long run level.

B Agents fractions dynamics

The micro-foundation of the price equilibrium price equation (A.7) offers an appealing interpretation of λ_t as the fraction of wealth of the informed agents, and a natural way to endogenize its evolution. In fact the wealth fraction at time t, λ_t, is endogenously determined as a function of the fraction at time $t-1$, λ_{t-1}, the fraction of wealth invested in the risky asset by both agents at time t, x_{t-1}^I and x_{t-1}^U, and of the return of the market at time t, r_t in (A.1). Using the definition of the fraction of wealth of the informed agent $\lambda_t = w_t^I/(w_t^I + w_t^U)$ and wealth evolution

$$w_t^H = w_{t-1}^H(1+r_f) + w_{t-1}^H(r_t - r_f)x_{t-1}^H \quad \text{for} \quad H = I, U$$

one can easily derive the equation that governs the evolution of the fraction of wealth of the informed agents λ_t:

$$\lambda_t = \frac{\lambda_{t-1}((1+r_f) + \pi_t^I)}{(1+r_f) + \lambda_{t-1}\pi_t^I + (1-\lambda_{t-1})\pi_t^U}, \tag{B.1}$$

where

$$\pi_t^I = (r_t - r_f)x_{t-1}^I, \tag{B.2}$$

and

$$\pi_t^U = (r_t - r_f)x_{t-1}^U, \tag{B.3}$$

are the realized excess profit per unit of wealth for informed and for uninformed agents respectively. When the realized single period return of the asset r_{t+1} is higher than r_f if the informed agents invests a higher (lower) share of their wealth compared to the share of uninformed agents, their fraction of wealth increases (decreases) compared to the fraction of wealth of the uninformed agents. Equation (B.1) corresponds to the replicator dynamics equation given in Eq. (2.24) when $\delta = 0$ and when $\rho = 1 + r_f$. To obtain (2.25) one has to further assume that (2.21) and (2.22) can be used as proxies of the realized profits per unit of wealth for respectively the informed agent, as (B.2), and for the uninformed agent, as (B.3). In fact one has to assume that the dynamics of the fractions is driven by the forecasting error of the uninformed compared to the cost of information for the informed, rather than by their realized profits. Investigation of this second framework would lead to a more complicated system due to the presence of r_t, and thus of both y_t and p_t/p_{t-1}, in the expression of agents' profits. To conclude the correspondence between Eq. (B.1) and Eq. (2.24) or Eq. (2.25) we have to discuss the case $\delta \neq 0$. Assume that, in every period, a number of agents which holds a fraction δ of total agents' wealth exits the market and is replaced by new agents with the same amount of wealth. Also assume that these new agents split evenly between being informed and being uninformed. This would mean that at period t the total fraction of informed agents is given by:

$$\lambda_t = (1-\delta)\frac{\lambda_{t-1}((1+r_f)+\pi_t^I)}{(1+r_f)+\lambda_{t-1}\pi_t^I+(1-\lambda_{t-1})\pi_t^U} + \frac{\delta}{2}, \tag{B.4}$$

which, with $1 + r_f = \rho$, is as (2.24) in Subsection 2.3 for every $\delta \in [0, 1]$. Notice that irrespectively of the fitness measure, realized profits or forecasting errors, both expressions (2.25) and (B.4) for the fraction of informed agents λ_t have the same dependence on the previous fraction of informed agents λ_{t-1}.

C Proofs

Proof of Proposition 1 Given the linear difference equation in (3.3), that is

$$y_t^e = \frac{r-g}{\lambda} + \frac{\alpha+\lambda-1}{\lambda} y_{t-1}^e,$$

and the initial condition y_0, from the theory of linear systems follows that

$$y_t^e = (y_0^e - (r-g)) \left(\frac{\alpha+\lambda-1}{\lambda}\right)^t + (r-g).$$

is the unique solution. The solution converges to $y^* = r - g$ as long as $\lambda > (1-\alpha)/2$ otherwise it diverges to $\pm\infty$. □

Proof of Proposition 2 Solving for the fixed point of (3.1-3.2) leads to $y^* = r - g$ and to λ^* solution of the following second order equation

$$c\lambda^2 + (c\delta/2 - c - \delta\rho)\lambda + \delta\rho/2 = 0,$$

which can be shown to have two real roots. Take $\alpha \in (0,1)$, $\delta \in (0,1)$ and $k = \rho/c > 1$. From:

$$2 - \delta + 2k\delta - 1 > 0,$$

it follows that

$$\frac{2 - \delta + 2k\delta + \sqrt{(2 - \delta + 2k\delta)^2 - 8k\delta}}{4} > 1.$$

That is one real root is always larger then 1 and thus not in the co-domain of our state variable λ. The other solution can be shown to be always in the interval $[0, 1/2]$. In fact

$$0 < \lambda^* = \frac{2 - \delta + 2k\delta - \sqrt{(2 - \delta + 2k\delta)^2 - 8k\delta}}{4} < \frac{1}{2},$$

reduces to

$$0 < 8k\delta \quad \text{and} \quad -4(1-\delta) < 0.$$

Both inequalities ara always satisfied. The Jacobian follows from evaluating the derivatives of (3.1-3.2) at the fixed point (y^*, λ^*). For the stability condition notice that the matrix is diagonal and the second eigenvalue, $\nu_2 \in (0,1)$. In fact since $\delta < 1$, $k > 1 > 0.5 > \lambda^*$ one has:

$$0 < \nu_2 = (1-\delta)\frac{k(k-1)}{(k-\lambda^*)^2} < \frac{k(k-1)}{k(k-1) + (0.5)^2} < 1. \tag{C.1}$$

The value of the first eigenvalue, ν_1, depends upon the value of λ^*. This eigenvalue is the same as the linear coefficient of equation (2.18), that is, it is $\nu_1 < 1$ given $\alpha \in (0,1)$ and $\lambda^* \in (0,1)$ and $\nu_1 > -1$ as long as

$$\lambda^* > \frac{(1-\alpha)}{2}.$$

Given the value of λ^* one can check that the previous inequality is satisfied if and only if

$$\delta > \frac{(1+\alpha)}{1 + \frac{2\alpha k}{(1-\alpha)}}.$$

□

Proof of Proposition 3 The matrix $\mathbf{J}|_{(y^*,\lambda^*)}$ in (3.6) is diagonal. As a result the dynamics of y and λ around (y^*, λ^*) can be linearized along the orthogonal basis with eigenvalues given by the diagonal entries of the matrix. Thus the eigenvalues that governs the dynamics of λ, (C.1), is given by the entry $(2, 2)$ of the matrix (3.6). We recall from the previous proof that

$$\nu_2 = (1-\delta)\frac{k(k-1)}{(k-\lambda^*)^2},$$

and we have already shown that $\nu_2 < 1$. The lower bound, $\nu_2 > (1-\delta)(1 - 1/k)$, follows from the previous expression and $\lambda^* > 0$ for all $\delta > 0$ and for all α. □

References

Anufriev, M. and Dindo, P. (2006). Equilibrium return and agents' survival in a multiperiod asset market: analytic support of a simulation model. Technical Report. CeNDEF Working paper 06-03, University of Amsterdam.

Barsky, R. B. and De Long, J. B. (1993). Why does the stock market fluctuate. *Quarterly Journal of Economics*, **58,** 291–311.

Barucci, E., Monte, R. and Renó, R. (2004). Asset price anomalies under bounded rationality. *Computational Economics*, **23,** 255–269.

Binmore, K. and Samuelson, L. (1997). Muddling through: Noisy equilibrium selection. *Journal of Economic Theory*, **74,** 235–265.

Boswijk, H. P., Hommes, C. H. and Manzan, S. (2005). Behavioral heterogeneity in stock prices. Technical Report. CeNDEF Working paper 05-12, University of Amsterdam.

Bray, M. (1982). Learning, estimation, and the stability of rational expectations. *Journal of Economic Theory*, **26,** 318–339.

Brock, W. A. and Hommes, C. H. (1997). Rational routes to randomness. *Econometrica*, **65,** 1059–1095.

Brock, W. A. and Hommes, C. H. (1998). Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control*, **22,** 451–481.

Bulkley, G. and Tonks, J. (1989). Are U.K. stock prices excessively volatile? Trading rules and variance bound tests. *Economic Journal*, **99,** 1083–1098.

Chow, G. C. (1989). Rational versus adaptive expectations in present values models. *Review of Economics and Statistics*, **71,** 376–384.

De Fontnouvelle, P. (2000). Information dynamics in financial markets. *Macroeconomics Dynamics*, **4,** 139–169.

Droste, E., Hommes, C. and Tuinstra, J. (2002). Endogenous fluctuations under evolutionary pressure in Cournot competition. *Games and Economic Behavior*, **40,** 232–269.

Ezekiel, M. (1938). The cobweb theorem. *Quarterly Journal Of Economics*, **52,** 255–280.

Fama, E. F. and French, K. R. (1988a). Dividend yields and expected stock returns. *Journal of Financial Economics*, **22,** 3–25.

Fama, E. F. and French, K. R. (1988b). Permanent and temporary components of stock prices. *Journal of Political Economy*, **96,** 246–273.

Gallagher, L. A. and Taylor, M. P. (2001). Risky arbitrage, limits of arbitrage, and nonlinear adjustment in the dividend-price ratio. *Economic Inquiry*, **39,** 524–536.

Goldbaum, D. (2005). Market efficiency and learning in an endogenously unstable environment. *Journal of Economic Dynamics and Control*, **29,** 953–978.

Grossman, S. J. and Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *American Economic Review*, **70,** number 3, 393–408.

Hellwig, M. F. (1982). Rational expectations with conditioning on past prices: A mean-variance example. *Journal of Economic Theory*, **26,** 279–312.

Hommes, C.H. (2006). Heterogeneous agents models in economics and finance. In *Handbook of Computational Economics* (eds L. Tesfatsion and K. L. Judd), volume 2, chapter 23. Elsevier, North Holland.

Le Baron, B. (2006). Agent based model in finance. In *Handbook of Computational Economics* (eds L. Tesfatsion and K. L. Judd), volume 2, chapter 24. Elsevier, North Holland.

LeRoy, S. F. and Porter, R. D. (1981). The present-value relation: Tests based on implied variance bounds. *Econometrica*, **49,** 555–574.

Levy, M., Levy, H. and Solomon, S. (1994). A microscopic model of the stock market: cycles, booms, and crashes. *Economic Letters*, **45,** 103–111.

Manzan, S. (2003). Nonlinear mean reversion in stock prices. Technical Report. CeNDEF Working paper 03-02, University of Amsterdam.

Muth, J. F. (1960). Optimal properties of exponentially weighted forecasts. *Journal of the American Statistical Association*, **55,** 299–306.

Muth, J. F. (1961). Rational expectations and the theory of price movements. *Econometrica*, **29,** number 3, 315–335.

Poterba, J. M. and Summers, L. H. (1988). Mean reversion in stock prices. *Journal of Financial Economics*, **22,** 27–59.

Routledge, B. R. (1999). Adaptive learning in financial markets. *Review of Financial Studies*, **12,** 1165–1202.

Shiller, R. J. (1981). Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review*, **71,** 421–436.

Shiller, R. J. (1984). Stock prices and social dynamics. *Brookings Papers on Economic Activity*, **2,** 457–510.

Summers, L. H. (1986). Does the stock market rationally reflect fundamental values. *Journal of Finance*, **41,** 591–602.

Timmermann, A. (1993). How learning in financial markets generates excess volatility and predictability of excess returns. *Quarterly Journal of Economics*, **108,** 1135–1145.

Timmermann, A. (1996). Excess volatility and predictability of stock prices in autoregressive dividend models with learning. *Review of Economics Studies*, **63,** 523–557.

Weibull, J.W. (1995). *Evolutionary Game Theory*. MIT Press.

Young, H. P. and Foster, D. (1991). Cooperation in the short and in the long-run. *Games and Economic Behavior*, **3,** 145–156.