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

Cees Diks† and Pietro Dindo‡

CeNDEF Working Paper 06-11 University of Amsterdam

March 2007

Abstract

In this paper we study the properties of an asset pricing model where boundedly rational agents respond to incoming news about economic fundamentals such as future dividends. Our aim is to characterize the resulting fluctuations of the market price around the time-varying underlying fundamental value. The starting point is an asset market in which agents can choose among two different degrees of information regarding future dividends. At the same time agents also try to learn the growth rate of the dividend generating process. Their interaction leads to prices that deviate perpetually from the fundamental value in the short run but stay close to it in the long run. In particular, prices exhibit time-varying nonlinear mean reversion, with parameters determined by the learning process.

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

JEL Classification: G12, D8, E3

Acknowledgments

Earlier versions of this work were presented at the Tinbergen Institute in Amsterdam, the Department of Mathematics of the Politecnico di Milano, the 4th PhD Conference "Research in Economics" (2006) in Volterra, and at the 15th Symposium of the Society for Nonlinear Dynamics and Econometrics in Paris. Stimulating discussions with seminar participants are gratefully acknowledged. We also thank Cars Hommes, Sebastiano Manzan and Roald Ramer for their helpful comments on an earlier draft.

*Research supported by the Netherlands Organisation for Scientific Research (NWO) under a MaGW-Vernieuwingsimpuls grant.

†CeNDEF, Faculty of Economics and Business, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands

‡Corresponding author. CeNDEF, Faculty of Economics and Business, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands, and LEM, St. Anna School of Advanced Studies, Piazza Martiri della Libertà 33, 56127 Pisa, Italy

1 Introduction

The aim of this paper is to investigate how news about fundamentals affects price fluctuations in an asset market populated by boundedly rational agents. The model is set up in such a way that if all agents have full information regarding fundamentals, the price behaves according to the 'classical' asset pricing model described by Gordon (1962). To allow for possible deviations from this benchmark, we assume that agents do not fully know the parameters of the dividend process, and allow them to choose between different information sets regarding future dividends. This leads, among other results, to an explicit model for nonlinear mean reversion.

We are by no means the first to address bounded rationality in this context. Since the discovery of so-called financial anomalies, the validity of the efficient market hypothesis has been questioned on the basis of empirical evaluation. For instance, markets exhibit excess volatility as described by Shiller (1981) and LeRoy and Porter (1981), mean reversion of asset prices, 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, 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.

One class of models in the asset market literature with boundedly rational agents takes explicitly into account the role of news on fundamentals in the price dynamics. Early examples are Bulkley and Tonks (1989) and Barsky and De Long (1993) who investigate the effect of agents trying to learn 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 parameters defining 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) as if the underlying parameters were known. 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 dividends. Convergence of the asset price to the present value of the stream of future dividends is achieved when the agents learn the correct parameters of the dividend process.

A limitation of these models is that they assume the presence of a representative agent, so that interaction of agents with different information or expectations does not play a role. Another class of models with boundedly rational agents does concentrate on the interaction of agents choosing different expectation schemes or different investment strategies (see e.g. the survey paper of Hommes, 2006). An early example is Chiarella (1992) where a model of a stylized financial market with fundamentalists and chartists is shown to generate a number of dynamic regimes which are compatible with the empirical anomalies reported above. In a more recent paper, Brock and Hommes (1998) consider a model with agents who do not know whether it is more profitable to predict prices by relying on fundamental information, or to extrapolate trends. To choose between these two strategies, agents use a performance measure such as realized profits. This repetitive re-evaluation of the performance of the strategies leads to complex price fluctuations. These fluctuations are totally endogenous and they do not require the presence of exogenous influences such as time-varying fundamentals.

A shortcoming of models in this class is that they typically do not take into account the

direct effect of news about the economic fundamentals on the 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 exogenous shocks.

In this paper we construct a framework for examining markets with boundedly rational agents where both the role of news about economic fundamentals and interaction between agents play a role. Because it is impossible to carry out this exercise under all conceivable behavioral assumptions, we limit ourselves to a simple class of models where all agents act upon the information available to them regarding fundamentals, including that revealed by prices. In our case the interaction is triggered by different expectations, due to the different degrees of information regarding future dividends only. This means that agents neither extrapolate price trends or use other chartist rules per se, nor expect other agents to do so, so that second or higher-order expectations play no role.

A convenient feature of our model is that it contains two important benchmarks as special cases. The first benchmark is the classical asset pricing model of Gordon (1962). When one discards both the role of informational differences and of agents learning the growth rate of dividends, the equilibrium price we derive coincides with the correct present value price. The second benchmark is the model developed by Barsky and De Long (1993) who assume that agents are trying to learn the growth rate of dividends. Our equilibrium price coincides with the price derived in the model of Barsky and De Long when we discard the role of informational differences. Our model can thus be seen as an extension of the Barsky and De Long model, where 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, our price 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 its fundamentals value 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 turns out to follow a stationary AR(1) process with time-varying coefficients. Our analysis shows that, whereas the transitory component follows an AR(1) process as a direct consequence of agents' learning the dividend growth rate, as already shown by Barsky and De Long (1993), the fact that the process coefficient are time-varying is due to agents interaction triggered by informational differences. This offers theoretical support to the empirical evidence that the temporary component in a mean reversion model is nonlinear and 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 also 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, in which agents can decide between two different degrees of information about the value of the asset return at the end of the period. We follow GS in assuming that agents operating in the asset market can decide whether or not to be informed about next period's dividend. However, as Bray (1982), Hellwig (1982) and Routledge (1999), we consider a dynamic model rather than a static. Instead of starting off by assuming that agents have rational expectations, we merely see rational expectations as a possible long run outcome of a learning process of boundedly rational agents using simpler rules. Failure of the uninformed agents to learn the relationship between prices and dividends implies deviations of the price from values fully revealing information about fundamentals. Another difference with the framework of GS and the papers mentioned above 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 about the value of the future dividend but also about 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 dividend process to be stationary in differences we assume the asset's dividend to be stationary in log-differences.

The impact of news in a financial market, where heterogeneous boundedly rational agents are operating, has also been investigated by Hong and Stein (1999). Their model reproduces stock prices that are under-reacting to the arrival of new information in the short run, and over-reacting in the long run. In contrast to our framework, they assume that the fractions of agents of each group are constant and that the dividend process is stationary in levels.

Manzan and Westerhoff (2005) conduct another study where the role of news in triggering deviations with respect to a fundamental level is investigated. They focus on exchange rate dynamics and develop a model in which boundedly rational agents place their orders depending on incoming news. Agents over-react to news when past market volatility is high and under-react to news when past market volatility is low. Their simulations show that, as the arrival of news affects the demands, persistent deviations of the exchange rate from its fundamental level can be observed.

This paper unfolds as follows. Section 2 introduces the model in terms of its three founding elements: the asset market, expectation formation and the evolution of the fractions of informed and uninformed agents. 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. 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. There we also relate the price dynamics generated by our model to the classical asset pricing model and to the model developed in Barsky and De Long (1993). Section 5 concludes with a discussion of our main results. The proofs of the propositions can be found in the Appendix.

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" value of the asset. That is, agents are prepared to sell the asset at a price above their "fair" value and to buy the asset at a price below it. 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 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, besides forming expectations regarding dividends, agents must also form expectations about 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 unable to compute the equilibrium relationship between price and dividends that should hold in a market where informed and uninformed agents operate.

The aim of this section is to characterize how, in this setting, the market price at time t of an asset, p_t, and the fraction of informed agents, λ_t, co-evolve given the agents' expectations and the dividend process $\{d_t\}$. In order to arrive at such a relationship we build up our model based as follows.

We consider a stochastic dividend process, $\{d_t\}$ which, in the benchmark case, is given by a geometric random walk. The dividend payed at time t, d_t, is assumed to be a geometric random walk with drift:

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

with $\eta_t := \mathrm{e}^{-\sigma_\eta^2/2+\sigma_\eta\varepsilon_t}$, where $\{\varepsilon_t\}$ represents a sequence of independent standard normal random variables. Consequently, $\{\eta_t\}$ is a sequence of independent $\mathrm{LOGN}(-\sigma_\eta^2/2, \sigma_\eta^2)$ random variables with mean $E(\eta_t) = 1$. The constant g is referred to as the growth rate of dividends.

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 estimate (belief) of the 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 (1). For the moment we consider g^e as being a known constant. In Subsection 2.2 we will discuss how agents actually estimate the growth rate g.

All agents are "fundamentalists" in the sense that they follow the present value model in computing what they consider to be the "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 information set of agents contains past dividend and price realizations and is different across groups of agents. We denote the information set at time t for agents of type H by $\mathcal{F}_t^H$. The fair value, i.e. the expected discounted sum of future dividends, conditional on $\mathcal{F}_t^H$ is denoted by v_t^H:

$$v_t^H = \sum_{i=1}^{\infty} \frac{d_{t,t+i}^{e,H}}{(1+r)^i}, \tag{2}$$

where $d_{t,t+i}^{e,H}$ is the estimate of the dividend to be payed at time $t+i$, taken at time t by the agents of type H. 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 time independent. The latter assumption is made because we want to concentrate on price fluctuations determined by agent interaction and learning rather than agents changing their discount rate. We also assume that the discount rate is always bigger than the agents' estimate of the long run growth rate of dividends.

At any 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, of type I, are fully informed regarding d_{t+1}. This implies that their current expectation of the $t+1$ dividend is

$$d^{e,I}_{t,t+1} = d_{t+1}, \tag{3}$$

We assume that informed agents pay a fixed cost $c > 0$ per period for this information. The uninformed agents, of type 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}. For example, if uninformed agents relied 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 uninformed agents try to get additional information revealed by the demands of the informed agents through the current market price p_t. We assume that the uninformed agents consider the relationship between the dividend and the price to be of the form:

$$d^{e,U}_{t,t+1} = y^e p_t, \tag{4}$$

where y^e is their estimate of the market dividend yield, i.e., the ratio between future dividend and current price. As for g^e, we will initially consider y^e to be a publically known constant. Subsection 2.2 will discuss how agents actually form their beliefs of the market dividend yield. 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 onwards, where we consider a model for the endogenous evolution of λ_t.

We assume that, at each time t, the ex-dividend market price of one share, p_t, is given by the following market equilibrium pricing equation:

$$p_t = \lambda v^I_t + (1-\lambda)v^U_t, \tag{5}$$

where v^I_t and v^U_t denote the "fair" value of the asset conditional on the information of the informed and uninformed respectively, and are derived according to (2). Under (5) the realized price today is a weighted average, weighted by the fractions of agents of each type, of the agent's estimate of the fair share value. Although this market equilibrium pricing equation is admittedly stylized, various settings have been considered in the literature where such a weighted average arises. For instance, Brock and Hommes (1998) derive (5) in a context where agents maximize a *CARA* mean-variance utility function. Dindo (2007) shows that (5) can be derived from maximization of a mean-variance *CRRA* utility function, under the assumption that the discount rate is equal to the asset return. The fraction λ and $(1-\lambda)$ should then be interpreted as the relative fraction of wealth of the informed and uninformed agents respectively, which seems reasonable if the choice of agents of being informed or not is independent of the wealth agents have accumulated so far. A third possible justification of (5) is obtained by assuming that, for the purpose of trading, agents are matched randomly. When two agents of the same type meet they trade at their common "fair" value price. Otherwise, when agents of different types meet, they trade the asset at the price in the middle between their "fair" values. As two informed agents meet with

probability λ^2, two uninformed agents meet with probability $(1-\lambda)^2$, and one informed and one uninformed agent meet with probability $2(\lambda)(1-\lambda)$ the average price emerging from this trading process is

$$p_t = \lambda^2 v_t^I + 2\lambda(1-\lambda)\frac{v_t^I + v_t^U}{2} + (1-\lambda)^2 v_t^U = \lambda v_t^I + (1-\lambda)v_t^U.$$

We next wish to derive the implication of the agents' behavior on the price dynamics in (5). We first consider (2) to derive the fair value for the informed and the uninformed agents.

Informed agents We have assumed 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. At time t the informed agents know the value of d_{t+1} so that their information set at time t is given by $\mathcal{F}_t^I = \{d_{t+1}, d_t, \ldots, p_t, p_{t-1}, \ldots\}$. Hence their expectations of future dividends are:

$$d_{t,t+j}^{e,I} = d_{t+1}(1+g^e)^{j-1}, \quad \text{for} \quad j \geq 1. \tag{6}$$

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 (6) into (2) we arrive at the informed agents' estimate of the value of the stock,

$$v_t^I = \frac{d_{t+1}}{(r-g^e)}. \tag{7}$$

The informed agents' stock valuation is thus in accordance with the Gordon model (Gordon, 1962). Eq. (7) implies that informed agents behave similarly to so-called fundamentalists in the interacting agents literature; for recent surveys see Hommes (2006) and LeBaron (2006). 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 the fair price of the informed agents, v_t^I, as the fundamental price p_t^*, that is, we define:

$$p_t^* = \frac{d_{t+1}}{(r-g^e)}. \tag{8}$$

Uninformed agents Boundedly rational uninformed agents try to infer the value of d_{t+1} from the market clearing price p_t. In doing so, they use their model (4) concerning the relationship between the current realized market price p_t and expected future dividends d_{t+1}. Combining this with agents being fundamentalists, using the estimate g^e as if it is the true value of the growth rate of dividends, we obtain

$$d_{t,t+j}^{e,U} = y^e p_t (1+g^e)^{j-1}, \quad \text{for} \quad j \geq 1, \tag{9}$$

which, using (2), gives the 'uninformed' valuation of the stock:

$$v_t^U = \frac{y^e p_t}{(r-g^e)}. \tag{10}$$

Notice that also in this case, as we have seen for informed agents, there is a correspondence with the literature on interacting agents as summarized by Hommes (2006) and LeBaron (2006). The "fundamentalists" uninformed agents are behaving "as if" they are chartists, that is, they use current prices to estimate the value of the asset. In particular, if at time t, $y^e/(r-g^e)$ is larger than one, uninformed agents behave "as if" they are trend followers, and can drive prices well above the fundamental levels. The converse happens when $y^e/(r-g^e) < 1$. In this case the uninformed agents behave "as if" they are contrarian. 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.

Combining the asset valuation for both types of agents, specified in (7) and (10), and the market equilibrium pricing equation (5) one obtains,

$$p_t = \frac{d_{t+1}}{(r-g^e)} \frac{\lambda(r-g^e)}{(r-g^e-(1-\lambda)y^e)}. \tag{11}$$

This equation shows that the realized price is proportional to the fundamental price p_t^* defined in (8) — a similar result one finds with the Gordon (1962) model — but there is an additional factor due to the presence of the uninformed agents trying to extract information from the market price. One can consider this as a generalization of the Gordon equation to a simple setting where agents do not know the growth rate of dividends and have different degrees of information. Notice that when $\lambda = 1$ (i.e. all agents are informed) the realized prize p_t in (11) is equal to the fundamental price p_t^* in (8). On the other hand when $\lambda = 0$, i.e. all agents are uninformed, either $p_t = 0$ or, if $y^e = r - g^e$, p_t is indeterminate. Thus, if $\lambda = 0$ the pricing equation (11) is not well defined, as the price contains no information to be extracted about future dividends.

Eq. (11) can be rewritten as a relationship between the realized dividend yield $y_{t+1} = d_{t+1}/p_t$, the expected growth rate g^e and dividend yield y^e, and the fraction of informed agents λ:

$$y_{t+1} = \frac{(r-g^e)}{\lambda} - \frac{(1-\lambda)}{\lambda} y^e := f(y^e, g^e, \lambda). \tag{12}$$

We call the map f an expectational feedback map because, given the 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 regarding the dividend yield and dividend yield realizations. Using the expectational feedback map (12), 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 expectational feedback equilibrium or the rational expectations dividend yield. Notice that when $y^e = r - g^e$, using equation (11), one finds that the market price equals the fundamental price p_t^*, which we have defined in (8) as the fair value of the informed agents. For this reason we denote

$$y^* = r - g^e \tag{13}$$

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 update 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 deviate from y^*. In particular the presence of uninformed agents can generate a price that differs from the fundamental price.

2.2 Expectation formation

In our model agents form expectations about the growth rate of dividends and the dividend yield. As far as the growth rate of dividends is concerned we have assumed that both informed and uninformed agents have homogeneous expectations on the dividend growth rate. We follow Barsky and De Long (1993) in assuming that agents estimate the growth rate using an adaptive expectations scheme:

$$g_t^e = \gamma g_{t-1}^e + (1-\gamma)\left(\frac{d_t}{d_{t-1}} - 1\right), \tag{14}$$

where g_t^e denotes the estimate of the future growth rate based on the dividend information available up to time t, and where γ is a memory parameter specifying the rate at which agents discount past information. Naive expectations are obtained in the special case $\gamma = 0$, whereas g_t^e is the mean of all past observations if instead of a fixed value of γ one would take $\gamma = (t-1)/t$ for $t > 1$ with $g_1^e = (d_1/d_0 - 1)$. Note that we make the simplifying assumption here that informed and uninformed agents, although they have different information sets, use the same (publically available) estimate of the growth rate.

If the estimated growth rate is taken to be time-varying, the definition of the fundamental price and of the fundamental dividend yield should be updated from (8) and (13) to, respectively:

$$p_t^* = \frac{d_{t+1}}{r - g_t^e}, \tag{15}$$

$$y_{t+1}^* = r - g_t^e. \tag{16}$$

The same expectation used for the estimation of the future growth rate of dividends is assumed to be used for the estimation of the value of the dividend yield y^e. This assumption is consistent with the existing of a relationship between growth rate of dividends and the dividend yield as given by the Gordon model or equation (16). Notice that only uninformed agents need to form expectations about the dividend yield. In fact, only uninformed agents use estimates of the dividend yield to predict the future dividends as specified in (9). Adaptive expectations for the dividend yield are specified by:

$$y_t^e = \alpha y_{t-1}^e + (1-\alpha) y_t, \tag{17}$$

where y_t^e denotes the belief or estimate of the dividend yield 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. With time varying estimates of the growth rate of dividends and of the dividend yield, the expectational feedback map changes from (12) to:

$$y_{t+1} = f(y_t^e, g_t^e, \lambda) = \frac{(r - g_t^e)}{\lambda} - \frac{(1-\lambda)}{\lambda} y_t^e. \tag{18}$$

When using adaptive expectations, agents are not strictly using optimal estimators. We are aware that, given the data generating process of $\{d_t\}$ in (1) and the endogenous dynamics of

y_t in (18), adaptive expectations in (14) and (17) are not minimizing mean squared errors. We simply assume that agents estimate the value of the growth rate g and of the dividend yield y as geometrically down-weighted averages of past observations, with memory parameters γ and α respectively. This discounting of past observations reflects the fact that agents take into account the possibility that the world they inhabit is nonstationary.

2.3 Evolution of the fraction of informed agents

So far we have assumed that the fractions of informed and uninformed agents are fixed. We next consider a model for the way fractions change over time. The driving force behind changes in the fractions of informed and uninformed agents is the trade-off between the costs of information and the precision of the dividend yield estimator based on public information. Intuitively, given the costs 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 costs of being informed, the smaller the fraction of informed agents.

The evolution of the fraction of informed agents λ is modeled by the replicator dynamics mechanism. We recall that the replicator dynamics can be motivated in the context of boundedly rational agents who are learning and imitating strategies in a strategic environment (see e.g. Weibull, 1995 and Binmore and Samuelson, 1997). As a measure of success of each strategy we take the difference between its implementation costs and the squared error of its estimation of the dividend yield.

Informed agents know the future dividend so that they have zero forecast error. In order to implement their strategy, they pay a cost $c > 0$ per time unit. Uninformed agents try to forecast future dividends by estimating the dividend yield at no cost. At time t the squared forecast error for the realization of the dividend yield is $(y^e_{t-1} - y_t)^2$, where y^e_{t-1} is the uninformed agents estimate of the dividend yield taken at time $t-1$ as defined in (18). In addition, we assume that agents are granted a fixed payoff ρ. As a result, we can define the fitness π^H_t of the strategy H at time t. The fitness of the strategy to buy information is

$$\pi^I_t = \rho - c, \tag{19}$$

while the fitness of the strategy of remaining uninformed is

$$\pi^U_t = \rho - \left(\frac{y^e_{t-1} - y_t}{y^*_t} \right)^2. \tag{20}$$

The time-varying fundamental dividend yield y^*_t defined in (16) is introduced in the denominator of π^U_t 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 $(100\sqrt{c})\%$ is made. For example, if $c = 0.1$, the two fitness measures are equal when errors of approx. 30% around y^*_t are made. The fixed payoff ρ assures that both (19) and (20) remain positive and can be taken as fitnesses in a replicator dynamics framework. In fact, only when fitnesses are positive, the replicator dynamics produces fractions in the interval $[0, 1]$. Notice that in the case of (19) a positive fitness requires $\rho > c$.

We can derive a different formulation of π_t^U by considering that, in the presence of a time varying fraction λ_t, (18) becomes:

$$y_{t+1} = f(y_t^e, g_t^e, \lambda_t) = \frac{r - g_t^e}{\lambda_t} - \frac{(1-\lambda_t)}{\lambda_t} y_t^e. \tag{21}$$

Using this relation between y_{t+1} and $y_{t,t+1}^e$, we can rewrite the fitness of the uninformed agents (20) as:

$$\pi_t^U = \rho - \left(\frac{y_t^* - y_{t-1}^e}{\lambda_{t-1} y_t^*} \right)^2. \tag{22}$$

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

$$\lambda_t = (1-\delta) \frac{\lambda_{t-1} \pi_t^I}{\lambda_{t-1} \pi_t^I + (1-\lambda_{t-1}) \pi_t^U} + \frac{\delta}{2}, \tag{23}$$

where the parameter δ is to be interpreted as a mutation or experimentation parameter (see e.g. Young and Foster, 1991; Droste *et al.*, 2002). The parameter δ is related to what we call evolutionary (or selection) pressure in the following way. When $\delta = 0$ the updating of the fractions is determined only 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. Intermediate values of δ result in a convex combination of selection pressure and experimentation. Based on expressions (19) and (22), the replicator dynamics (23) gives,

$$\lambda_t = (1-\delta) \frac{\lambda_{t-1}(-c+\rho)}{\lambda_{t-1}(-c) - (1-\lambda_{t-1}) \left(\frac{y_t^* - y_{t-1}^e}{\lambda_{t-1} y_t^*} \right)^2 + \rho} + \frac{\delta}{2}. \tag{24}$$

The fixed payoff ρ regulates the speed of adjustment of the replicator dynamics. When ρ is large changes in λ are smaller, everything else being equal, than when ρ is small. Moreover, in the limit $\rho \to \infty$ the trajectories of 24 approach the trajectories of the corresponding continuous dynamical system (see e.g. Weibull 1995). For this reason we refer to ρ also as to the inverse of the speed of adjustment.

It is beyond the scope of this paper to consider other specifications for the dynamics of λ, or other formulations for the fitness measures π^H. Although it would be interesting to consider other definitions, we leave this analysis for future research. Concerning the fitness measure, in particular, one could argue that it should be based on total returns, given by both dividend yields and capital gains, and not on dividend yields forecast errors as we do here. At this first level of investigation, we have not made this choice to avoid complications in the analysis of the model, which should then be re-expressed in terms of capital gains other than dividend yields. Notice, however, that as stock returns from capital gains would equally affect the fitness of both the informed and uninformed agent we can conjecture what can happen if one incorporates capital gains in the fitnesses. This would be equivalent to considering a time varying, instead of a fixed, payoff ρ. For the replicator dynamics (23) this would imply a time varying speed of adjustment.

2.4 Market returns

To summarize, for lognormal shock η_t such that $d_t = d_{t-1}(1+g)\eta_t$ with g as the long run growth rate of dividends, the dynamics of the full model developed so far is given by the following three equations:

$$g_t^e = \gamma g_{t-1}^e + (1-\gamma)\left((1+g)\eta_t - 1\right), \tag{25}$$

$$y_t^e = \frac{(1-\alpha)(r-g_{t-1}^e)}{\lambda_{t-1}} + \frac{\alpha+\lambda_{t-1}-1}{\lambda_{t-1}} y_{t-1}^e, \tag{26}$$

$$\lambda_t = (1-\delta)\frac{\lambda_{t-1}(-c+\rho)}{\lambda_{t-1}(-c) - (1-\lambda_{t-1})\left(\frac{r-g_{t-1}^e-y_{t-1}^e}{\lambda_{t-1}(r-g_{t-1}^e)}\right)^2 + \rho} + \frac{\delta}{2}. \tag{27}$$

Equation (25) gives the common expectation formation regarding the growth rate of dividends g^e as defined in (14) with γ as the memory parameter. Equation (26) is obtained by putting together the expectation formation of the dividend yield y^e by the uninformed traders as defined in (17) and the time varying expectational feedback map (21). The parameter α is a memory and r is the agents' discount rate. Equation (27) expresses the dynamics of the fraction of informed agents λ as in (24) with c being the cost of information per time step, ρ the fixed payoff, and δ the experimentation rate. At every point in time the three state variables $(g_t^e, y_t^e, \lambda_t)$ are a function of their lagged values, $(g_{t-1}^e, y_{t-1}^e, \lambda_{t-1})$, and of the shock η_t. The dividend yield dynamics is given by the evolution of the three state variables through the feedback equation (21).

Before presenting a formal analysis of the system above, notice that whenever y_t converges to its steady state value, that is, whenever y_t^e converges to the rational expectations value $y_{t+1}^* = r - g_t^e$ as in (16), the price follows $p_t^* = d_{t+1}/r - g_t^e$ as in (15). The fundamental price depends on the changing estimates of the growth rate of dividends. This is the same price as derived by Barsky and De Long (1993). If, moreover, $g_t^e \to g$, i.e. if the expected dividend growth rate converges to its true value, the fundamental price converges to the "correct" present value price: $y_t^* \to r-g$ and $p_t^* \to d_{t+1}/(r-g)$. In this case the price follows a geometric random walk. 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 (25–27) to converge to its fixed point, in which case the adaptive expectations do not converge to rational expectations. When this is the case, 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 system (25-27) is stochastic, and the estimate g^e is time varying. In what follows, we analyze these effects separately as well as their interplay. First, in Section 3, we concentrate on the role of informational differences and analyze the conditions of convergence of the deterministic dynamics of y and λ to their equilibrium values. In Section 4 we complement this analysis by investigating the effect of a time varying estimate of the growth rate of dividends g^e.

3 Informational differences

In this section we analyze the impact of informational differences alone on the dynamics of asset prices without the learning of the divided growth rate. That is, we take $g_t^e \equiv g^e$. Technically, we analyze the system of equations (25–27) when the memory parameter $\gamma = 1$. Without loss of generality we consider only the case $g^e = g$. The generalization to $g^e = g'$, where g' is a constant strictly smaller than r, is straightforward and implies only a shift of the level of the steady state dividend yield from $r - g$ to $r - g'$.

If $g_t^e = g$, the system (25–27) reduces to a two dimensional (2-D) system in the variables y_t^e and λ_t:

$$y_t^e = \frac{(1-\alpha)(r-g)}{\lambda_{t-1}} + \frac{\alpha + \lambda_{t-1} - 1}{\lambda_{t-1}} y_{t-1}^e, \tag{28}$$

$$\lambda_t = (1-\delta)\frac{\lambda_{t-1}(-c+\rho)}{\lambda_{t-1}(-c) - (1-\lambda_{t-1})\left(\frac{r-g-y_{t-1}^e}{\lambda_{t-1}(r-g)}\right)^2 + \rho} + \frac{\delta}{2}. \tag{29}$$

The model parameters 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 fixed payoff. Notice that, given the dynamics of (y_t^e, λ_t) specified by (28–29), the dynamics of the dividend yield y_{t+1} can be derived by using the expectational feedback map f defined in (21). Before investigating the full dynamics of (28–29) it is instructive to consider the 1-D system obtained when the fraction λ_t of informed agents is fixed at a constant value λ.

3.1 Dividend yield dynamics

Taking $\lambda_t \equiv \lambda$, Eq. (28) becomes:

$$y_t^e = \frac{(1-\alpha)(r-g)}{\lambda} + \frac{\alpha + \lambda - 1}{\lambda} y_{t-1}^e. \tag{30}$$

Given the linearity of (30), the analysis of the dynamics is straightforward and it is possible to compute the general solution of the difference equation. That is, given y_0^e one can analytically express the value of y_t^e, for all t. The following proposition summarizes this.

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 $r > g$, the solution of the difference equation (30) with initial condition y_0^e is given by:*

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

where

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

If

$$\lambda > \bar{\lambda} \equiv \frac{1-\alpha}{2}, \tag{31}$$

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

Notice that whenever the expected dividend yield y_t^e converges to its steady state y^*, also the realized dividend yield y_t converges to y^* through the feedback map (21). At the steady state y^* the price equals the fundamental price p_t^* in (8) and thus fully reveals the information concerning the next dividend. The shaded area in Fig. 1 shows the stability region of (30) in the parameter space, (α, λ), whereas the white area shows the unstable region. The shaded area is divided in two gray scales. In the lighter region, the convergence of the expected dividend yield to the steady state y^* is oscillatory, whereas in the darker, the convergence is monotone. Notice that the border between the stable and unstable regions is characterized by oscillatory behavior of the expected dividend yield y_t^e, and thus of the realized dividend yield y_t. This implies that failure of the price to fully reveal the fundamental information should be characterized by price fluctuations with negative autocorrelation. Already from this analysis it appears that, in the long run, only when the fraction of informed agents is above a certain threshold, the learning of uninformed agents leads to a fully informative price. If there are too few informed agents, the price diverges from its steady state (fundamental) value and does not reflect all available information.

INSERT FIGURE 1 HERE

3.2 Dividend yield and fractions dynamics

In the more general case the fraction of informed agents, λ_t, is time dependent. In this case the dynamics of the dividend yield and of the fraction of informed agents is given by the nonlinear system (28–29). We use local stability analysis to characterize the behavior of the state variables (y_t^e, λ_t) near the steady state of the system.

First consider the case of zero information costs. The following proposition characterizes the steady states of the system and their local stability.

Proposition 2. *Consider zero information costs, $c = 0$. Given the memory parameter $\alpha \in (0, 1)$, the fixed payoff $\rho > 0$ and the required rate of return $r > g$, one has:*
(i) If the experimentation rate $\delta = 0$, every point (y^*, λ^*) with $y^* = r - g$ and $\lambda^* \in (0, 1]$ is a steady state. For a given $\alpha \in (0, 1)$ the set of steady states such that $\lambda^* > \bar{\lambda} = (1 - \alpha)/2$, as in (31), is globally stable.
(ii) If the experimentation rate $\delta \in (0, 1)$, the unique steady state of the system (28–29) is the point (y^*, λ^*), where

$$y^* = r - g \quad \textit{and} \quad \lambda^* = \frac{1}{2}.$$

This steady state is globally stable for every parameter combination.

Thus, with zero information costs, the expected dividend yield y_t^e always converges to the fundamental value y^* defined in (13) and the fraction of informed agents is either $1/2$, when $\delta > 0$, or any value satisfying (31) when $\delta = 0$. As a consequence, the expectational feedback map (21) converges to the rational expectations equilibrium and prices are fully informative in the long run. This is not always true for positive information costs, in which case the following proposition applies.

Proposition 3. *Consider positive information costs, $c > 0$. Given the memory parameter $\alpha \in (0, 1)$, the ratio between the fixed payoff and the cost of being informed, k, such that*

$k = \rho/c > 1$ and the required rate of return $r > g$, one has:
(i) If the experimentation rate $\delta = 0$, no steady state is well defined.
(ii) If the experimentation rate $\delta \in (0,1)$, the unique steady state of the system (28–29) is the point (y^*, λ^*) where

$$y^* = r - g$$

and

$$\lambda^* = \frac{2 - \delta + 2k\delta - \sqrt{-8k\delta + (2 - \delta + 2k\delta)^2}}{4}. \tag{32}$$

Moreover, $\lambda^* \in (0, 1/2)$. The Jacobian of (28–29) at the steady state is diagonal and given by

$$\mathbf{J}|_{(y^*,\lambda^*)} = \begin{pmatrix} \dfrac{\alpha + \lambda^* - 1}{\lambda^*} & 0 \\ 0 & (1-\delta)\dfrac{k(k-1)}{(k-\lambda^*)^2} \end{pmatrix}. \tag{33}$$

If

$$\delta > \bar{\delta} \equiv \frac{(1+\alpha)}{1 + \frac{2\alpha k}{(1-\alpha)}} = \frac{(1+\alpha)}{1 + \frac{2\alpha\rho}{(1-\alpha)c}}, \tag{34}$$

the steady state (y^*, λ^*) is locally stable. This condition corresponds to the stability condition (31) of the 1-D dynamical system (30). That is, the condition $\delta > \bar{\delta}$, and $\lambda^* > \bar{\lambda}$ are equivalent.

The local stability condition (34) 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. Recall that when the expected dividend yield y_t^e converges to y^*, the realized dividend yield y_t also converges to y^*. Therefore, stability of the steady state (y^*, λ^*) implies convergence of the price to the fundamental price p_t^* in (8) and thus to a fully informative price.

INSERT FIGURE 2 HERE

In a repeated single period model with rational agents, GS show that no equilibrium value of the fraction of informed agents exists for which the price fully reveals the information about future dividends. In such an equilibrium, if it existed, nobody would be willing to pay for the information, hence prices could not possibly reveal any information. In our model, which has a long living asset and boundedly rational agents, the analogue of the paradox described by GS is part (i) of Proposition 3 which states that for positive information costs and no experimentation, no steady state equilibrium dividend yield (and thus price) exists.

The correspondence of our results with those of GS holds also if we add some form of noise or experimentation. Part (ii) of Proposition 3 shows that if the experimentation level is positive, a steady state dividend yield (and thus price) does exist. At the steady state uninformed agents do learn to extract information, and the system is well defined because, due to experimentation the fraction of informed agents is "artificially" kept above zero, and prices converge to being fully informative — the conditions for the paradox no longer exist due to experimentation. In this respect, experimentation in our model has a similar effect as the introduction of noise on

the price equation, considered by GS in their model, which also removes the conditions for the paradox to arise.

It is important to notice that even if a steady state dividend yield is well defined in our model, the learning of boundedly rational agents need not always converge to it. In fact, the second part of statement (ii) in Proposition 3 shows that the equilibrium is (locally) stable only if the experimentation rate is sufficiently large ($\delta > \bar{\delta}$). If, instead $\delta < \bar{\delta}$, small perturbations destabilize the system and fluctuations around the fully informative equilibrium price arise. In particular, the definition of $\bar{\delta}$ in Eq. (34) of Proposition 3 shows that, for a given α, the stability region of the system (28–29) shrinks, if the cost of information c increases, and expands if the fixed payoff ρ increases.

Summarizing, for a positive information cost, prices are not always fully informative in the long run, depending on the experimentation rate as compared to $\bar{\delta}$ in (34). Moreover, prices are never fully informative if the experimentation rate δ is zero.

What happens to the dynamics of the expected dividend yield and the fraction of informed agents if the steady state is unstable and prices do not converge to being fully informative? To answer this question we analyze the global dynamics of the system (28–29) for small experimentation levels, $\delta < \bar{\delta}$. If the stability conditions (31) and (34) 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 a typical evolution of the uninformed agents' expected dividend yield, y_t^e, and of the fraction of informed agents, λ_t, respectively. At time $t = 0$, the fraction of informed agents is above the dotted line, which marks the value $\bar{\lambda}$ in (31) 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 (21), 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_t^* that informed agents are better off. Eventually, paying the cost of being informed leads to a higher fitness than using a freely available estimate with a large error. As a result, the fraction of informed agents λ_t grows sharply, see e.g. the top right plot of Fig. 2 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 progresses this sequence of events repeats itself, with λ_t 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_t^e, λ_t) space and the corresponding dynamics of the log price compared to the log fundamental price.

INSERT FIGURE 3 HERE

Fig. 3 shows phase plots of the system for decreasing values of the experimentation rate δ. The typical concentration of points for smaller values of δ, along lines stretching away in two directions from the fixed point, suggest that the fluctuations of y_t^e and λ_t, are associated with a so-called homoclinic bifurcation. Similar phenomena are encountered in other multi-dimensional nonlinear systems and emerge from the interplay between local instability of the

steady state and global stability of the dynamics. Brock and Hommes (1997) and Droste *et al.* (2002) present other economic frameworks where homoclinic bifurcations arise. They also offer detailed discussions of the mathematical aspects of these interesting phenomena.

If the experimentation rate δ is sufficiently small, by assuming that agents are boundedly rational and that fractions are endogenously determined, our model offers another mechanism for balancing the cost of information and the informational content of the price: the learning process of the uninformed agents. In our framework we thus obtain a "dynamic equilibrium degree of disequilibrium" due to endogenous price fluctuations produced by the interaction of boundedly rational agents. With the insight given by the simulations shown in Fig. 3, we conjecture that these type of price fluctuations also characterize the case of a zero experimentation rate ($\delta = 0$), even if an equilibrium does not exist.

Before turning to the analysis of the full model it is instructive to characterize the convergence of the fraction of informed agents in the stability region. Close to the equilibrium, λ_t turns out to change very slowly.

Proposition 4. *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_t in a neighborhood of (y^*, λ^*), we have*

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

This proposition shows that if the experimentation rate δ is small and the ratio between the speed of adjustment and the cost of information, k, is large, ν_2 is very close to one. As a result, if the system is stable, changes in the value of the fraction of informed agents λ_t are very slow, and hence λ_t is very persistent. This will be confirmed in the next section, where we appraise the effect of agents learning the growth rate of dividends on the dynamics.

4 Informational differences and parameter learning

In the previous section we have assumed that the agents' estimate of the growth rate of dividends is constant. As a result, the equation for the dividend yield was fully deterministic. In this section we analyze the simultaneous impact of informational differences and of agents' learning of the growth rate of dividends as new information about the fundamentals becomes available. In this case the stochastic dividends affect the state variables, and the dynamical system becomes stochastic.

Recalling the results from Section 2, the learning of the long run growth rate of dividends is specified by (14) which, when the dividend follows the process specified in (1), gives:

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

This stochastic equation, together with the evolution of the dividend yield, its expectations, and the fraction of informed agents, lead to a stochastic version of the deterministic model (28–29),

that is, to the equations (25–27) which we rewrite here for convenience:

$$g_t^e = \gamma g_{t-1}^e + (1-\gamma)\left((1+g)\eta_t - 1\right), \tag{35}$$

$$y_t^e = \frac{(1-\alpha)(r-g_{t-1}^e)}{\lambda_{t-1}} + \frac{\alpha+\lambda_{t-1}-1}{\lambda_{t-1}} y_{t-1}^e, \tag{36}$$

$$\lambda_t = (1-\delta)\frac{\lambda_{t-1}(-c+\rho)}{\lambda_{t-1}(-c) - (1-\lambda_{t-1})\left(\frac{r-g_{t-1}^e-y_{t-1}^e}{\lambda_{t-1}(r-g_{t-1}^e)}\right)^2 + \rho} + \frac{\delta}{2}. \tag{37}$$

The shocks $\{\eta_t\}$ on the growth rate of dividends drives the co-evolution of the state variables, consisting of the agents' estimate g_t^e of the growth rate of dividends, the expected dividend yield y_t^e, and the fraction of informed agents λ_t. Given the expected growth rate of dividends, g_t^e, the expected dividend yield, y_t^e, and the fraction of informed agents, λ_t, the dividend yield itself, y_{t+1}, is determined by the feedback map (21). Before we numerically analyze the impact of shocks on the dynamics of (35–37), we show that our model contains two important benchmarks as special cases.

Classical Asset Pricing model The first benchmark is the classical asset pricing model, based on the assumption that all agents know the dividend growth rate g, that agents use rational expectations, and that information costs can be neglected. In this case, if only a few agents are informed about d_{t+1}, the other can correctly extract this information from prices. The market price and the market dividend yield are then given by:

$$p_t^{**} = \frac{d_{t+1}}{r-g}, \quad y_{t+1}^{**} = r - g. \tag{38}$$

where p_t^{**} denotes the "correct" present value price and y_t^{**} the "correct" dividend yield. Our model boils down to the classical asset pricing model if we put the cost of information at $c = 0$ and assume agents know the growth rate of dividends: $g_t^e = g$. As noted in Section 3, in this case the system is deterministic and Proposition 2 applies. In case (i) ($\delta = 0$) as well as (ii) ($\delta \in (0,1)$) there are fixed points, at any of which which $y_t^e = r - g$. Substituting this in (36) and using the feedback map (21) to transform the expected dividend yield y_{t+1}^e into realized dividend yield y_{t+1}, gives $y_{t+1} = r - g$ which corresponds to y_{t+1}^{**} given in (38). The classical asset price model is also obtained if the cost of information is nonzero, provided that the fraction of informed agents is one ($\lambda_t = 1$). Substituting $g_t^e = g$ and $\lambda_t = 1$ into (21) immediately gives $y_{t+1} = r - g$, irrespective of the expected dividend yield (which would usually govern the behavior of the uninformed agents, but these have a fraction zero now).

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 growth rate g. In their case the price and the dividend yield are given by:

$$p_t^{*} = \frac{d_{t+1}}{r-g_t^e}, \quad y_{t+1}^{*} = r - g_t^e, \tag{39}$$

where g_t^e is given by adaptive expectations as in (35). These are respectively what we have defined as the fundamental price and the fundamental dividend yield in (15) and (16). Our model corresponds to the model of Barsky and De Long (1993) under the assumptions that agents fractions are fixed and all agents are informed ($\lambda_t = 1$). In fact, by substituting $\lambda_t = 1$ in (36) and using the feedback map (21) one retrieves $y_{t+1} = r - g_t^e$, with g_t^e defined in (35), which corresponds to y_{t+1}^* defined in (39).

A way of evaluating the differences between our model and the two benchmarks is to write, for all the three models, the evolution equation for the dividend yield as a function of lagged dividend yields and shocks on the growth rate of dividends and compare the resulting expressions. In the classical asset pricing model the dividend yield is constant and given by $y^{**} = r - g$. In the Barsky and De Long model, one can use $y_{t+1}^* = r - g_t^e$ to rewrite (35) as:

$$y_{t+1}^* = (1-\gamma)(1+r) + \gamma y_t^* - (1-\gamma)(1+g)\eta_t. \tag{40}$$

That is, the dividend yield follows an AR(1) process driven by the shocks $\{\eta_t\}$ on the growth rate of dividends. The memory parameter γ determines the AR(1) coefficient and affects 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 classical asset pricing dividend yield $y^{**} = r - g$.

In our model, fixing for the moment the value of λ, the dynamic equation (36) for the expected dividend yield y_t^e rewritten in terms of the realized dividend yield y_{t+1} using the feedback map (21), gives:

$$y_{t+1} = \frac{(r - g_t^e) - \alpha(r_{t-1} - g_{t-1}^e)}{\lambda} + \frac{\alpha + \lambda - 1}{\lambda} y_t. \tag{41}$$

This equation expresses that, when the fraction of informed agents is fixed at λ, the dividend yield follows an AR(1) process with shocks that are correlated (via the estimated growth rates g_s^e) with the shocks of the growth rate of dividends. As for the dividend process implied by the Barsky and De Long model in (41), the long run mean of the dividend process implied by our model is given by the classical asset pricing dividend yield $y^{**} = r - g$. If $\gamma = \alpha$, that is, if agents use the same memory parameter to estimate the growth rate of dividends and the dividend yield, the quantity $g_t^e - \alpha g_{t-1}^e$ reduces to $(1-\gamma)((1+g)\eta_t - 1)$, depending on the latest innovation only, and Eq. (41) takes the following appealing form:

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

Furthermore, if we define

$$\tilde{\gamma} = \tilde{\gamma}(\lambda) \equiv 1 - (1-\gamma)/\lambda = \gamma - (1-\gamma)\frac{1-\lambda}{\lambda}, \tag{42}$$

we can rewrite (41) as

$$y_{t+1} = (1 - \tilde{\gamma}(\lambda))(1+r) + \tilde{\gamma}(\lambda) y_t - (1 - \tilde{\gamma}(\lambda))(1+g)\eta_t. \tag{43}$$

The result is that, when $\gamma = \alpha$, our model specified by Eq. (43) and the model of Barsky and De Long (1993) specified by Eq. (40) 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 $\tilde{\gamma}(\lambda)$ as the

effective memory. Eq. (42) shows that the effective memory has two components, one given by the real memory and the other related to the presence of uninformed agents. This second effect becomes less important as more informed agents are present in the market. The general result is that $\tilde{\gamma}(\lambda)$ in (42) is an increasing function of λ with $\tilde{\gamma}(\lambda) \leq \gamma$ for all λ, and $\tilde{\gamma}(1) = \gamma$. That is, the presence of uninformed agents is equivalent to all agents being informed and using an effective memory which is lower then the real memory. The value of $\tilde{\gamma}$ determines both the AR(1) coefficients and the variance of the shocks but not the long run mean which is always $y^{**} = r - g$. Notice that 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 larger short run effect and a smaller long run effect. Also, since the effective memory $\tilde{\gamma}$ is a function of λ, our model allows for variation of the memory parameter with the fraction of informed agents λ. Changes in λ thus have an impact, through $\tilde{\gamma}$, 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 (43) compared to the two benchmarks in (38) and (40).

4.1 Nonlinear mean reversion

Our model (43) clearly differs, both structurally and regarding parameter values, from that of Barsky and De Long in (40) if the fraction of informed agents is time varying. In this case, our model implies an AR(1) process 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 recently proposed to characterize fluctuations of stock indices. By using the fact that $p_t^{**} \equiv d_{t+1}/(r-g)$, i.e. the price implied by model (38), 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:

$$\log(p_t) \approx \log(p_t^{**}) + z_t, \tag{44}$$

The dynamics of z_t can be easily derived using its definition in terms of x_t, the definition of x_t and Eq. (43) and is given by:

$$z_t = \tilde{\gamma}(\lambda_t) z_{t-1} + \frac{(1-\tilde{\gamma}(\lambda_t))(1+g)}{(r-g)}(\eta_{t-1} - 1). \tag{45}$$

This equation shows that we have a model whose realized log price in (44) is the sum of the persistent component $\log(p_t^{**})$, which follows a random walk with drift, and of the temporary component, z_t, which follows a stationary autoregressive process (45) with a time-varying AR(1) coefficient $\tilde{\gamma}(\lambda_t)$. Empirical investigation of the properties of stock prices are in accordance 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 as would be predicted by the model of Barsky and De Long. 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 autoregressive coefficients and heteroskedasticity can be related to agent interaction triggered by informational differences. In fact, both the AR(1) coefficient and the shocks variance in (45) are a function of $\tilde{\gamma}(\lambda_t)$ which is a nonlinear function of the time varying fraction of informed agents λ_t.

4.2 Simulation study

To illustrate the qualitative effect of the shocks on the growth rate of dividends on the dividend yield and the dynamics of fractions, we first analyze 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 dividend yields and prices generated by our model, but also for dividend yields and prices generated by the classical asset pricing model (38) and by the model of Barsky and De Long (40). 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 refer to this model as the "modified" Barsky and De Long model and its series of prices and dividend yields by $p^*_{\tilde{\gamma}}$ and $y^*_{\tilde{\gamma}}$. We use the "modified" Barsky and De Long model to appraise the role of time variability of effective memory. Notice that when different models are compared, they are simulated with the same realization of the dividend process.

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 p^{**}_t, the price implied by the classical asset pricing model. Fig. 4 shows that in both cases there is an initial over-reaction followed by convergence to the equilibrium value, which is given by p^{**}_t.

INSERT FIGURE 4 HERE

Since, for $\lambda < 1$, the effective memory $\tilde{\gamma}$ 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 over-reaction is more pronounced. At the same time, when $\tilde{\gamma} < \gamma$, 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 $\tilde{\gamma}$ to changes in λ. From the Jacobian of the 2-D system (see Proposition 3) 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 4, 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 $\tilde{\gamma}$ shown in the right panel. A confirmation of the fact that one shock has no considerable consequence on changes of $\tilde{\gamma}$ comes from the time series for $p^*_{\tilde{\gamma}}$ shown in the left panel. The price $p^*_{\tilde{\gamma}}$ is the price obtained using the "modified" Barsky and De Long model, that is equation (40) with $\gamma = \tilde{\gamma}(\lambda_0)$, 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 $p^*_{\tilde{\gamma}}$ 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 $p^*_{\tilde{\gamma}}$. 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 long before such a change has an effect on the value of the effective memory.

INSERT FIGURE 5 HERE

We now turn to investigating the effect of a sequence of shocks. The dividends are generated as a geometric random walk with drift, as described just after (1). In particular $\{\eta_t\}$ is a sequence of i.i.d. log-normal shocks. Figure 5 shows the impact of a sequence of 500 i.i.d. shocks $\{\eta_t\}$. The 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 (source: Shiller database available from R. J. Shiller's homepage). The discount rate r is taken such that $y^{**} = 0.05$, which for the present value model would imply a price of 20 times the the value of the dividend. In our simulations r is always larger than $(1 + g)\eta_t$ so that the stochastic system (35–37) is always well defined. Since dividends are paid quarterly, 500 consecutive dividends correspond to 125 years. The top left panel shows the time series of the dividend yield y generated by our model whereas the top 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, 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.

INSERT FIGURE 6 HERE

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 (42). As a confirmation of our previous results and of our theoretical analysis in Proposition 4, changes in $\tilde{\gamma}$ (through 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 $y^*_{\tilde{\gamma}}$. Notice that deviations of up to more than ten percent arise. Our conclusion is that when subsequent shocks are present, both the fact that the effective memory is lower than the real memory and that the effective memory is 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 time varying informational differences, and more generally agents interaction, might explain the nonlinearity in the mean reversion that has been shown to exist in the empirical literature.

INSERT FIGURE 7 HERE

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 q_t,

$$q_t = \log(p_t + d_t) - \log(p_{t-1}), \tag{46}$$

for a typical run of our model (35–37). 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 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 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 (46) 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 $\tilde{\gamma}(\lambda)$ is time varying. Further research is required to investigate these issues in greater detail and relate them to the statistical properties of empirical market returns.

5 Discussion and conclusion

We have investigated the combined effect of informational differences and learning in a stylized asset market model where agents are boundedly rational. As far as the theoretical guidelines behind our model are concerned, we have shown 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, such as Grossman and Stiglitz (1980), that analyze the impact of learning, such as Barsky and De Long (1993), and that investigate the interaction of agents who are using different predictor schemes or different strategies, such as Brock and Hommes (1998).

We have developed a model where fluctuations of prices are driven by agents learning the growth rate of dividends while the less informed agents among them try to extract information form the price. We have shown that the first effect is responsible for an autoregressive dynamics of the dividend yield and that the second effect is responsible for time-varying coefficients of the parameters regulating this process. These variations are driven by the time-varying informational content of the price which, due to the learning of uninformed agents, switches repeatedly between being nearly fully informative and hardly informative. As a result, our model has the same features of recent empirical investigations of mean reversion conditions for stock prices, such as Gallagher and Taylor (2001) and Manzan (2003) where time-varying parameters of the transitory component have been detected.

It is beyond the scope of this paper to calibrate our model to reproduce the stock price evolution given the historical dividend process. Our theoretical model is based on several simplifying assumptions and in particular on an admittedly ad hoc dynamics of the fraction of informed

agents, λ. Nevertheless it is instructive to note that time-varying parameters in a mean reversion model can be related to what in general may be referred to as interaction of agents, which in our case is triggered by informational differences. The fact that agent interaction can be responsible for nonlinearity in the behavior of stock prices is also argued by Boswijk *et al.* (2007), 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 by linking it to other observable characteristics such as the volume of transactions, might provide a basis for the design of new econometric tests for structure in financial time series.

Appendix

A Proofs

Proof of Proposition 1 Consider the linear difference equation (30) given by

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

and initial condition y_0^e. From the theory of linear dynamical systems it 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 $\lambda > (1-\alpha)/2$, and otherwise diverges to $\pm\infty$. □

Proof of Proposition 2 (i) When both $c = 0$ and $\delta = 0$, solving for the fixed point (y^*, λ^*) of (28–29) leads to $y^* = r - g$ and to any $\lambda^* \in (0, 1]$. For global stability we proceed by showing that for every initial condition the dynamics converges to (y^*, λ^*). Rewriting (28) in terms of deviations from the fixed point $y^* = r - g$ leads to

$$y_t^e - y^* = \frac{\alpha + \lambda_{t-1} - 1}{\lambda_{t-1}} (y_{t-1}^e - y^*). \tag{47}$$

Whenever $\lambda_{t-1} > (1-\alpha)/2 = \bar{\lambda}$, $y^e \to y^*$ as the linear coefficient of equation (47) is always lower than one. Since $\alpha \in (0, 1)$, the condition $\lambda > \bar{\lambda}$ is always satisfied when $\lambda_{t-1} \geq 1/2$. We show that, in the long run $\lambda \geq 1/2$ always holds. First consider $\lambda_{t-1} \geq 1/2$. Using (29) one can derive that it holds

$$\lambda_t - \frac{1}{2} = (1-\delta) \frac{\lambda_{t-1}\rho - (1-\lambda_{t-1}) \left(\rho - \left(\frac{(r-g) - y_{t-1}^e}{\lambda_{t-1}(r-g)} \right)^2 \right)}{2 \left(-(1-\lambda_{t-1}) \left(\frac{(r-g) - y_{t-1}^e}{\lambda_{t-1}(r-g)} \right)^2 + \rho \right)} \geq 0 \tag{48}$$

if one assumes that fractions are always positive, i.e.

$$\rho - \left(\frac{(r-g) - y_{t-1}^e}{\lambda_{t-1}(r-g)} \right)^2 > 0.$$

As a consequence when $\lambda_{t-1} \geq 1/2$ also $\lambda_t \geq 1/2$. If, on the other hand, $\lambda_{t-1} < 1/2$, one can similarly show that the condition $(\lambda_t - \lambda_{t-1}) > 0$ holds, implying that λ converges to one point of the set $[1/2, 1]$. As a consequence also when $\lambda_{t-1} < 1/2$, then $\lambda \geq 1/2$ in the long run. To conclude the proof of global stability notice that when $y^e \to y^*$ also $\lambda \to \lambda^*$ via equation (29). (ii) When $c = 0$ and $\delta > 0$ solving for the fixed point of (28-29) leads to $y^* = r - g$ and to $\lambda^* = 1/2$. Global stability can be shown along the same lines of the proof in (i). □

Proof of Proposition 3 (ii) For $c > 0$ and $\delta \in (0,1)$, solving for the fixed point (y^*, λ^*) of (28–29) leads to $y^* = r - g$ and to λ^* being the solution to 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 the condition

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

reduces to

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

Both inequalities ara always satisfied. The Jacobian is obtained by evaluating the derivatives of (28-29) at the fixed point (y^*, λ^*). Let ν_i, $i = 1, 2$ denote the eigenvalues of the Jacobian. For the stability condition notice that the matrix is diagonal and the second eigenvalue, $\nu_2 \in (0,1)$. Since $\delta < 1$, $k > 1 > 0.5 > \lambda^*$ one has

$$0 < \nu_2 = (1 - \delta)\frac{k(k-1)}{(k - \lambda^*)^2} < \frac{k(k-1)}{k(k-1) + (0.5)^2} < 1. \tag{49}$$

The first eigenvalue, ν_1, depends on λ^*, and is the same as the linear coefficient of equation (30), 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

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

(i) For $c > 0$ and $\delta = 0$, notice that the equation for λ^* gives $\lambda^* = 0$ as a solution. The dynamical system, and in particular (28), is not defined for this value of λ. Thus no steady state is well defined. □

Proof of Proposition 4 The matrix $\mathbf{J}|_{(y^*, \lambda^*)}$ in (33) 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 eigenvalue that governs the dynamics of λ, is given by the entry $(2,2)$ of the Jacobian (33). 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

Barsky, R. B. and De Long, J. B. (1993). Why does the stock market fluctuate? *Quarterly Journal of Economics*, **108,** 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. (2007). Behavioral heterogeneity in stock prices. Technical Report. CeNDEF Working paper 05-12, University of Amsterdam. *Journal of Economic Dynamics and Control*, forthcoming.

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.

Chiarella, C. (1992). The dynamics of speculative behavior. *Annals of Operations Research*, **37,** 101–123.

Dindo, P. (2007). *Bounded Rationality and Heterogeneity in Economic Dynamic Models*. Ph.D. Thesis, Tinbergen Institute and University of Amsterdam,.

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

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.

Gordon, M. J. (1962). *The Investment, Financing, and Valuation of the Corporation*. R. D. Irwin, Inc., Homewood, IL.

Grossman, S. J. and Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *American Economic Review*, **70,** 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, pp. 1109 – 1186. Elsevier, North Holland.

Hong, H. and Stein, J. C. (1999). A unified theory of underreaction, momentum trading and overreaction in asset markets. *Journal of Finance*, **54,** 2143–2184.

LeBaron, B. (2006). Agent based computational finance. In *Handbook of Computational Economics* (eds L. Tesfatsion and K. L. Judd), volume 2, chapter 24, pp. 1187–1234. 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.

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

Manzan, S. and Westerhoff, F. (2005). Representativeness of news and exchange rate dynamics. *Journal of Economic Dynamics and Control*, **29,** 677–689.

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, Cambridge, MA.

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

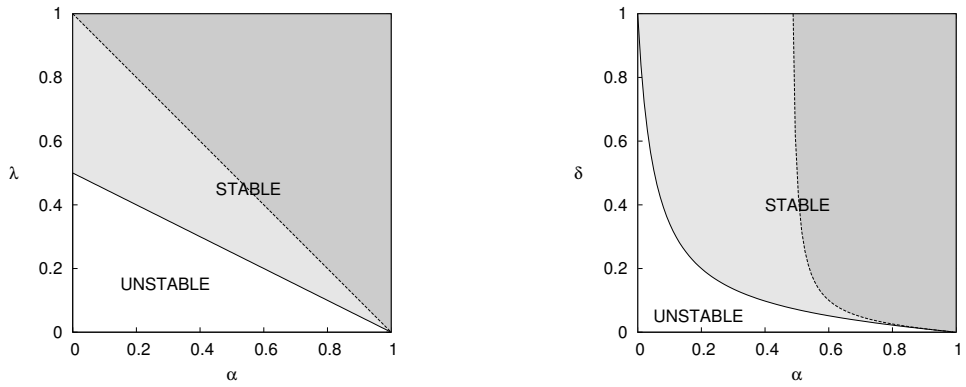

Figure 1: Left panel: Stability region for the 1-D dynamical system (30). The expected dividend yield converges to the steady state $y^* = r - g$ only for values of (α, λ) in the shaded area. In the darker region the convergence to y^* is monotone, whereas in the lighter region the convergence is oscillatory. Right panel: Stability region for the 2-D dynamical system (28–29) as a function of the mutation rate δ and memory parameter α, for $k = \rho/c = 10$. As in the left panel, the solid line marks the border of the stability region, while 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).

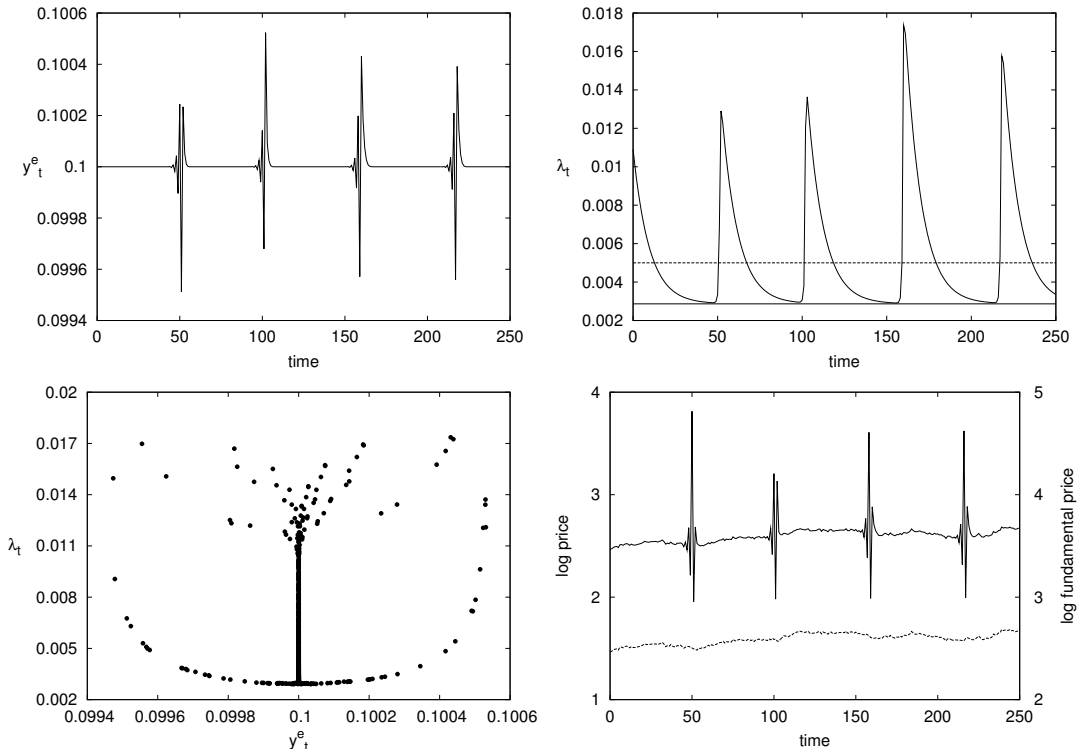

Figure 2: Top panels: time series for the expected dividend yield y_t^e (left panel) and the fraction of informed agents λ_t (right panel) from (28–29). In the top right panel the dotted line corresponds to $\bar{\lambda}$ in (31) (the critical stability border) whereas the solid line corresponds to the steady state λ^* in (32). The bottom left panel shows the state space representation (y_t^e, λ_t). In the bottom right panel, the log price dynamics derived from the dynamics of y_t^e and λ_t is shown. 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: $\alpha = 0.99$, $c = 0.1$, $\rho = 1$ (so that $k = \rho/c = 10$), $r = 0.1$, $g = 0$ and $\delta = 0.000575$.

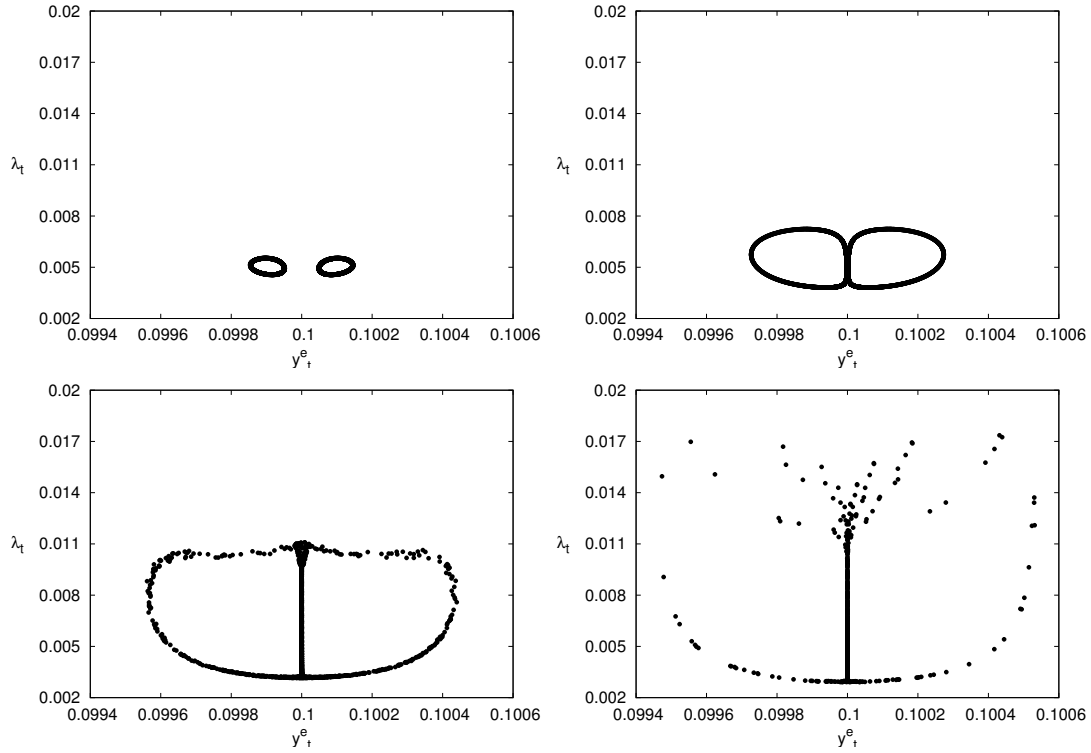

Figure 3: Phase plots of the expected dividend yield y_t^e and the fraction of informed agents λ_t produced by the system of difference equations (28–29). Parameter values: $\alpha = 0.99$, $c = 0.1$, $\rho = 1$ (so that $k = \rho/c = 10$), $r = 0.1$ and $g = 0$. Top left panel: $\delta = 0.000675$, top right panel: $\delta = 0.00066$. Bottom left panel: $\delta = 0.000625$, bottom right panel: $\delta = 0.000575$.

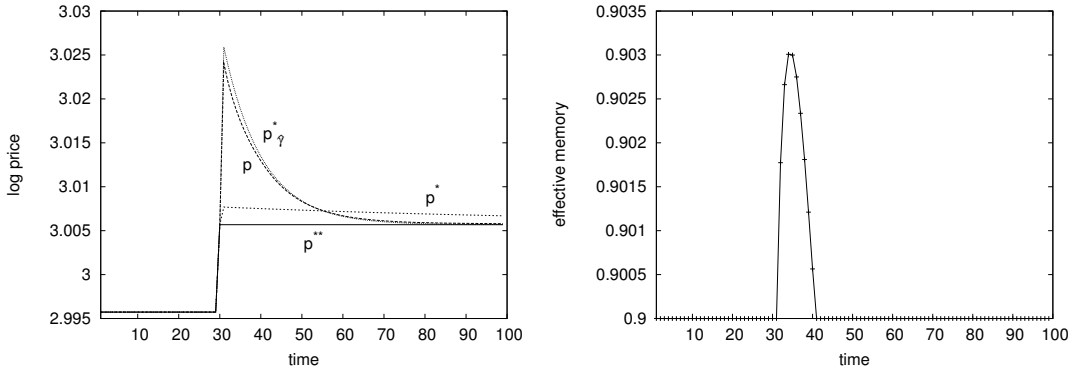

Figure 4: Effect of a single shock on the growth rate g (before and after the shock, $g = 0$). The shock is $\eta_{31} = 0.01$. Left panel: log prices as a function of time. Time series $\{p_t^{**}\}$ corresponds to the log price implied by the classical asset pricing model (38); $\{p_t^*\}$ and $\{p_t\}$ represent, respectively, the value of the log price given by Barsky and De Long model (40) and by our model (35–37), and $\{p_{\tilde{\gamma},t}^*\}$ represents the log price implied by the "modified" Barsky and De Long model (40) with $\gamma = \tilde{\gamma}(\lambda_0)$, where λ_0 is the fraction of informed agents before the shock. Right panel: evolution of the effective memory $\tilde{\gamma}$. Parameter values: $\gamma = \alpha = 0.99$, $\delta = 0.02$, $c = 0.1$ and $\rho = 1.0$.

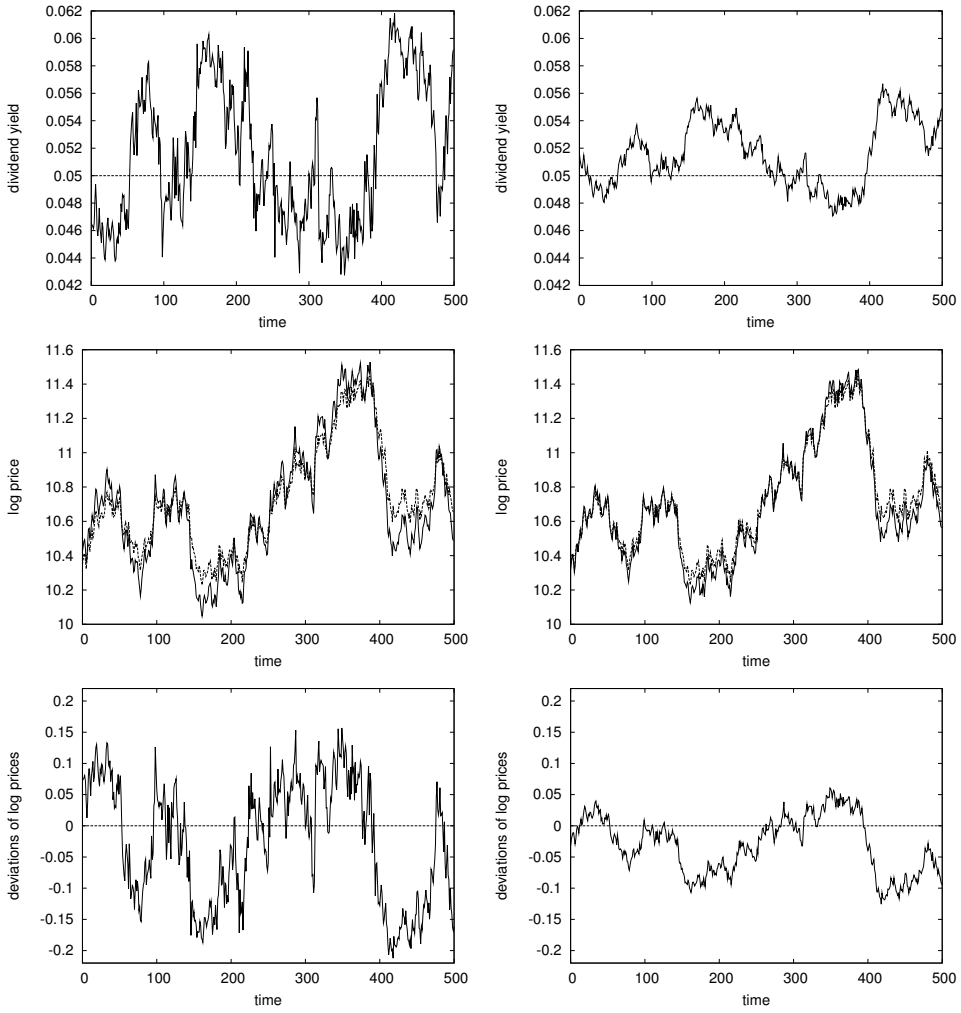

Figure 5: Top left panel: time series of the dividend yield y_t generated by our model (35–37) (solid line) compared with the benchmark y^{**} (38) (horizontal dotted line). Top right panel: time series of the dividend yield y^* as in Barsky and De long (40) (solid line) compared with y^{**}. Middle 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. Parameter values: $\alpha = \gamma = 0.99$, $\rho = 1$, $c = 0.1$, $\delta = 0.02$ (these three parameters imply $\lambda^* \approx 0.09$), $\sigma_\eta = 0.04$, $g = 0.003$. The discount rate is $r = 0.05 + g$. All the simulations shown are obtained with the same realization of the dividend process after a transient of 500 iterations.

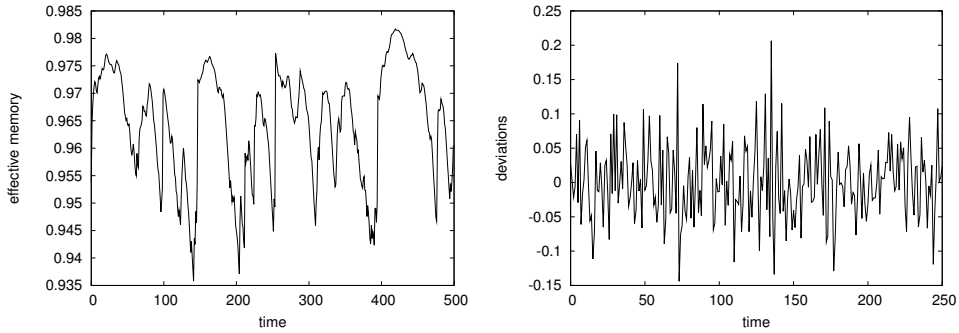

Figure 6: Left panel: values of the effective memory $\tilde{\gamma}$. Right panel: deviations of the two series of log prices generated by our model y and by the "modified" Barsky and De Long model $y^*_{\tilde{\gamma}}$, arising due to fluctuations in our model of $\tilde{\gamma}$ around its mean value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

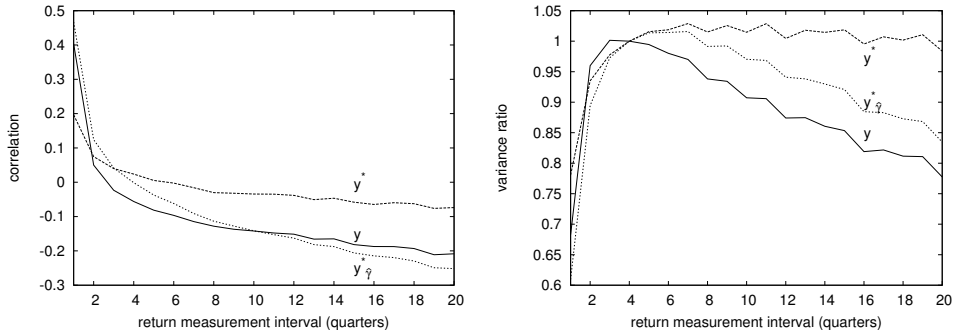

Figure 7: Left panel: autocorrelation of the log return time series for a typical run of our model (35–37). Right Panel, variance test of log lagged returns. $v(\Delta t) = (\sigma^2(q^{\Delta t})/\Delta t)/(\sigma^2(q_y^4)/4)$ where $\sigma^2(q_y^4) = 0.00355)$ and $q^{\Delta t}$ is the total return over a period Δt generated by the model in system (35–37). In both plots three lines refer to data generated with our model, y, to the model of Barsky and De Long (40), 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.