Learning, Large Deviations and Rare Events*

Jess Benhabib[†] NYU

Chetan Dave[‡] NYU

Revised: January 9, 2012

Abstract

We examine the asymptotic distribution of the price-dividends ratio in a standard asset pricing model when agents learn adaptively using a constant gain stochastic gradient algorithm. The asymptotic distribution is shown to exhibit fat tails even though dividends follow a stationary AR(1) process with thin tails.We then estimate the deep parameters of our adaptive learning model and show that they are consistent with our model calibration and with the data.

Keywords: Adaptive learning, large deviations, linear recursions with multiplicative noise

JEL Codes: D80, D83, D84

---

*We thank Chryssi Giannitsarou, In-Koo Cho, John Duffy, George Evans, Boyan Jovanovic, Tomasz Sadzik, Benoite de Saporta, and Tom Sargent for helpful comments and suggestions. The usual disclaimer applies.

[†]Address: New York University, Department of Economics, 19 W. 4th Street, 6FL, New York, NY, 10012, USA. E-mail: jess.benhabib@nyu.edu. Telephone: 1-646-567-7316, Fax: 1-212-995-4186

[‡]Address: New York University, Department of Economics, 19 W. 4th Street, 6FL, New York, NY, 10012, USA. E-mail: cdave@nyu.edu.

1. Introduction

Figures 1-2 plot aggregate stock prices and dividends in the U.S. as measured by the S&P 500 and CRSP datasets. The plots show that, as predicted by standard theory, prices and dividends do move in tandem. However the price-dividend ratio, shown in the third panel of each Figure, exhibits large fluctuations, especially in the latter parts of the sample.[1] These large fluctuations in the price-dividend ratio are difficult to explain with the standard rational expectations asset pricing model, for example that of Lucas (1978).[2]

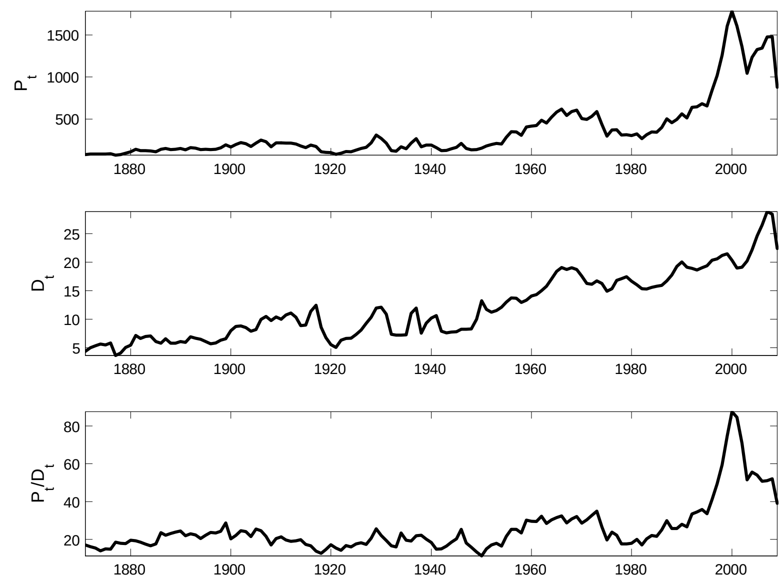

Figure 1. Annual S & P 500 (1871-2009).

[1] The Data Appendix provides details on the series employed.
[2] See for example Carceles-Poveda and Giannitsarou (2008).

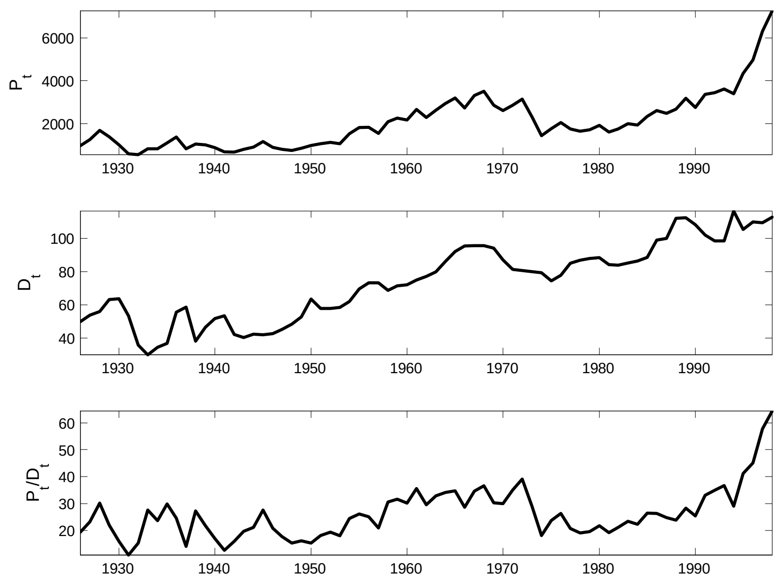

Figure 2. Annual CRSP (1926-1998).

Several modifications of the standard asset-pricing model have been proposed to account for this (and other) departures of the data from the model. A large and rich literature adopts a particular formulation that replaces the rational expectations assumption with that of adaptive learning: agents are assumed to estimate parameters of processes to be forecasted using recursive (adaptive) methods.[3] A shortcoming of this literature however is that the adaptive learning algorithms in general are not optimal in a Bayesian sense. In a stationary model with optimal learning, estimated parameters ultimately converge to their rational expectations equilibrium. In recent work however, Sargent and Williams (2005) introduce a model where agents expect a random walk drift in estimated parameters. They then show that the adaptive "generalized constant gain stochastic gradient" (SGCG) algorithm that assigns more weight to recent observations on account of the underlying drift in the estimated parameters is in fact the optimal Bayesian estimator. Evans et al. (2010) follow Sargent

[3]See for example Timmermann (1993, 1996), Brock and Hommes (1998), Bullard and Duffy (2001), Brennan and Xia (2001), Cogley and Sargent (2006), Branch and Evans (2006), Marcet and Nicolini (2008), Carceles-Poveda and Giannitsarou (2008).

and Williams (2005) and show how a SGCG learning algorithm approximates an optimal (in a Bayesian sense) Kalman filter. Under such adaptive SGCG learning, uncertainty about estimated parameters persists over time and can fuel 'escape' dynamics in which a sequence of rare and unusual shocks propel agents away from the REE of a model, even in the long-run stationary distribution.[4] In an asset-pricing context Weitzman (2007) also shows that if recent observations are given more weight under Bayesian learning of the variance of the consumption growth rate, agents will forecast returns and asset prices using thick-tailed distributions for consumption growth. [5]

Our paper is similar in spirit. We demonstrate, theoretically and empirically, that SGCG learning, which is consistent with optimal Bayesian learning, can account for the data features of the price-dividend ratio. Theoretically, we demonstrate that under adaptive learning the tails of the stationary distribution of the price-dividends ratio will follow a power law, even though the dividend process has thin tails and is specified as a stationary $AR(1)$ process. The tail index or power-law coefficient of the price-dividend ratio can be expressed as a function of model parameters, and in particular of the optimal gain parameter that assigns decaying weights to older observations. In fact, as demonstrated by Sargent and Williams (2005) and more recently by Evans et al. (2010), the optimal gain depends on the variance of the underlying drift in the estimated parameters: the higher the variance of the drift parameter,

[4]See also Holmstrom (1999) for an application to managerial incentives of learning with an underlying drift in parameters.

[5]See also Koulovatianos and Wieland (2011). They adopt the notion of rare disasters studied by Barro (2009) in a Bayesian learning environment. They find that volatility issues are well addressed. Similarly Chevillon and Mavroeidis (2011) find that giving more weight to recent observations under learning can generate low frequency variability observed in the data. See also Gabaix (2009) who provides an excellent summary of instances in which economic data follow power laws and suggests a number of causes of such laws for financial returns. In particular, Gabaix, Gopikrishnan, Plerou and Stanley (2006) suggest that large trades in illiquid asset markets on the part of institutional investors could generate extreme behavior in trading volumes (usually predicted to be zero in Lucas-type environments) and returns.

the higher the gain, and the thicker the tail of the distribution of the price-dividend ratio. Using some new results on random linear recursions with Markov dependent coefficients,[6] we characterize how the power law tail index of the of the long-run stationary distribution of the price-dividend ratio varies as a function of the gain parameter and of the other deep parameters of the model. Under our adaptive learning scheme that approximates optimal Bayesian learning, stationary dividend processes generate distributions for the price-dividend ratio that are not Normal. Rare shocks to exogenous dividends throw off the learning process and lead to large deviations from the rational expectations equilibrium.

However, our simulations indicate that under standard parameter calibrations, to match either the empirical tail index or the variance of the price dividend ratio we require a gain parameter around $0.4 - 0.55$, significantly higher than what is typically used in the adaptive learning literature $(0.01 - 0.04)$. The latter implicitly assumes slowly decaying weights on past observations, and therefore very little underlying drift in the parameters estimated by agents. In order to get an empirical handle on the gain parameter we estimate the parameters of our model, including the gain parameter, by two separate methods. The first is a structural minimum distance estimation method for the tail index. This method puts higher weight on the empirically observed tail of the price-dividend ratio, and produces a gain estimate in the range of $0.35 - 0.53$. The second method computes the gain as Bayesian agents expecting drifting parameters would, using a Kalman filter on the data. This yields a gain parameter in the range of $0.49 - 0.55$, assigning decaying weights on past observations that take the parameter drift into account. Therefore agents who use this gain parameter would have their expectations confirmed by the data.

[6] See Kesten (1973), Saporta (2005) and Roiterstein (2007).

The paper is structured as follows. We first describe the single asset pricing version of Lucas (1978) under learning. We then prove in Section 3 that the model, written as a random linear recursion, predicts that the tails of the stationary distribution of the price-dividends ratio will follow a power law with coefficient κ. In Section 4 we use simulations to study how κ varies with the deep parameters. In Section 5 we provide estimates of the deep parameters of the model, and of the gain parameter in particular, that are consistent with the κ estimated directly from the price-dividends ratio plotted in Figures 1 and 2 above. Section 6 concludes.

2. Learning and Asset Pricing

Consider a discrete time, single asset, endowment economy following Lucas (1978) with utility over consumption given by

$$u(C_t) = \frac{C_t^{1-\gamma}}{1-\gamma},\ \gamma > 0. \tag{1}$$

Under a no-bubbles condition the nonlinear pricing equation is

$$P_t = E_t \left\{ \beta \left(\frac{D_{t+1}}{D_t} \right)^{-\gamma} (P_{t+1} + D_{t+1}) \right\} \tag{2}$$

where $\beta \in (0,1)$ is the usual exponential discount factor and (real) dividends (D_t) follow some exogenous stochastic process. Linearizing the above equation yields

$$p_t = \beta E_t(p_{t+1}) + (1 - \beta - \gamma) E_t(d_{t+1}) + \gamma d_t \tag{3}$$

where all lowercase variables denote log-deviations from the steady state $(\overline{P}, \overline{D}) = \left(\frac{\delta}{1-\delta}, 1\right)$. We assume that the exogenous dividends process follows

$$d_t = \rho d_{t-1} + \varepsilon_t, \ |\rho| < 1 \tag{4}$$

in which ε_t is an $iid(0, \sigma^2)$ random variable (such that $\sigma^2 < +\infty$) with compact support $[-a, a]$, $a > 0$, and a non-singular distribution function F.[7] Since $E_t(d_{t+1}) = \rho d_t$

$$p_t = \beta E_t(p_{t+1}) + \theta d_t, \ \theta \equiv (1 - \beta - \gamma)\rho + \gamma \tag{5}$$

is the fundamental expectational difference equation for prices.[8]

We follow Evans and Honkapohja (1999, 2001) and assume the perceived law of motion (PLM) on the part of the representative agent is

$$p_t = \phi_{t-1} d_{t-1} + \xi_t, \ \xi_t \sim i.i.d.(0, \sigma_\xi^2), \ \sigma_\xi^2 < +\infty \tag{7}$$

which in turn implies

$$E_t(p_{t+1}) = \phi_{t-1} d_t. \tag{8}$$

where ϕ_{t-1} is the coefficient that agents estimate from the data to forecast p_t. Inserting the

[7] The distribution function F is non-singular with respect to the Lebesque measure if there exists a function $f \in R_+$, $\int_R f(t)dt > 0$, such that $F(dt) \geq f(t)dt$.

[8] The rational expectations solution to (5) is

$$p_t = \phi^{REE} d_t, \ \phi^{REE} = \frac{\theta}{1 - \beta\rho} \tag{6}$$

for all $\beta\rho \neq 1$.

above into (5) yields the actual law of motion (ALM) under learning:[9]

$$
\begin{aligned}
p_t &= \beta\phi_{t-1}d_t + \theta d_t = (\beta\phi_{t-1} + \theta)d_t && (9)\\
&= (\beta\phi_{t-1} + \theta)\rho d_{t-1} + (\beta\phi_{t-1} + \theta)\varepsilon_t && (10)
\end{aligned}
$$

By contrast the ALM under rational expectations is

$$p_t = \phi d_t = \phi\rho d_{t-1} + \phi\varepsilon_t. \tag{11}$$

Under SGCG learning, ϕ_t evolves as[10]

$$\phi_t = \phi_{t-1} + g d_{t-1}(p_t - \phi_{t-1}d_{t-1}),\ g \in (0, 1) \tag{12}$$

At this point we take the gain parameter g as given, but in section 5. we will estimate its value under our learning model with Bayesian agents who expect a random walk drift in ϕ. Following the usual practice in the literature for analyzing learning asymptotics, we insert

[9]We note that in the asset pricing context, the ALM is linear in the 'belief' parameter (ϕ_t). In other contexts the ALM might be nonlinear in beliefs. However, the linear forces generating large deviations in the adaptive learning model may drive the dynamics in nonlinear contexts. For example in Cho et al. (2002) adaptive learning leads to non-neglible probablities for large deviations even in the prescence of nonlinearities for the true data generating process.

[10]See Carceles-Poveda and Giannitsarou (2007, 2008) for details and derivations under a variety of learning algorithms.

the ALM under learning in place of p_t in the recursion for ϕ_t in (12) to obtain

$$\phi_t = \lambda_t \phi_{t-1} + \psi_t \tag{13}$$

$$\lambda_t = 1 - (1-\rho\beta) g d_{t-1}^2 + \beta g d_{t-1} \varepsilon_t = 1 - g d_{t-1}^2 + g\beta d_t d_{t-1} \tag{14}$$

$$\psi_t = \theta \rho g d_{t-1}^2 + \theta g d_{t-1} \varepsilon_t = \theta g d_t d_{t-1}. \tag{15}$$

The equation in (13) takes the form of a linear recursion with both multiplicative (λ_t in (14)) and additive (ψ_t in (15)) noise. We show in the next Section that the tail of the stationary distribution of ϕ_t follows a power law and can thus be fat. We characterize this tail and show that under learning the price-dividend ratio can exhibit large deviations from its rational expectations equilibrium value with non-negligible probabilities.

3. Large Deviations and Rare Events

We begin by noting that λ_t is a random variable, generating multiplicative noise, and can be the source of large deviations and fat tails for the stationary distribution of ϕ_t. We use results from large deviation theory (see Hollander (2000)) together with the work of Saporta (2005), Roitershtein (2007) and Collamore (2009) to characterize the tail of the distribution of ϕ_t.[11]

Let $\mathbb{N} = 0, 1, 2...$ and note that the stationary $AR(1)$ Markov chain $\{d_t\}_{t\in Z}$ given by (4) is uniformly recurrent, and has compact support $\left[\frac{-a}{1-\rho}, \frac{a}{1-\rho}\right]$ (see Nummelin (1984), p. 93). We denote the stationary distribution of $\{d_t\}_{t\in\mathbb{N}}$ by π. Since $\{d_t\}_{t\in\mathbb{N}}$ and ε_t for $t = 1, 2...$ are

[11] For an application of these techniques to the distribution of wealth see Benhabib et al. (2011) and to regime switching, Benhabib (2010).

bounded, so are $\{\lambda_t\}_{t\in\mathbb{N}}$ and $\{\psi_t\}_{t\in\mathbb{N}}$. In fact, following the first definition of Roitershtein (2007), $\{\lambda_t,\psi_t\}_{t\in\mathbb{N}}$ constitute a Markov Modulated Process (MMP): conditional on d_t, the evolution of the random variables $\lambda_{t+1}(d_t,d_{t-1})$ and $\psi_{t+1}(d_t,d_{t-1})$ are given by

$$
\begin{aligned}
P\left(d_t\in A,(\lambda_t,\psi_t)\in B\right) &= \int_A K\left(d,dy\right)G\left(d,y,B\right)|_{d=d_{t-1}}, &(16)\\
G\left(d,y,\cdot\right) &= P\left((\lambda_t,\psi_t)\in\cdot\right)\mid d_{t-1}=d,\ d_t=y\right), &(17)
\end{aligned}
$$

where $K\left(d,dy\right)$ is the transition kernel of the Markov chain $\{d_t\}_{t\in\mathbb{N}}$.

Next we seek restrictions on the support of the *iid* noise $\varepsilon_t\in[-a,a]$ to assure that $E\left|\lambda_\infty\right|<1$ where, from equation (14), λ_∞ is the random variable associated with the stationary distribution of d_t . We assume:

$$a<\left(\frac{6\left(1-\rho^2\right)}{g\left(1-\beta\rho\right)}\right)^{0.5}. \qquad (18)$$

Note that

$$
\begin{aligned}
E(\lambda_t) &= E\left(1-g\left(d_{t-1}\right)^2+g\beta\left(d_{t-1}\left(\rho d_{t-1}+\varepsilon_t\right)\right)\right)\\
E(\lambda_t) &= 1-gE\left(d_{t-1}\right)^2+g\beta\rho E\left(d_{t-1}\right)^2\\
E(\lambda_\infty) &= \left(1-gE\left(d_{t-1}\right)^2\left(1-\beta\rho\right)\right)_{t\to\infty}.
\end{aligned}
$$

Since ε_t is *iid* and is uniform with variance σ^2,

$$E(\lambda_\infty) = 1 - g\frac{\sigma^2}{1-\rho^2}(1-\beta\rho) \tag{19}$$

$$E(\lambda_\infty) = 1 - g\frac{\frac{1}{12}(2a)^2}{1-\rho^2}(1-\beta\rho). \tag{20}$$

From equation (20) it follows that $E(\lambda_\infty) < 1$, and solving for a such that $E(\lambda_\infty) > -1$, we obtain the restriction (18) to guarantee that $E|\lambda_\infty| < 1$.

Next, let $S_n = \sum_{t=1}^{n} \log|\lambda_t|$. Following Roitershtein (2007) and Collamore (2009)[12] the tail of the stationary distribution of $\{\phi_t\}_t$ depends on the limit[13]

$$\Lambda(\delta) = \lim_{n\to\infty} \sup \frac{1}{n} \log E \prod_{t=1}^{n} |\lambda_t|^\delta = \lim_{n\to\infty} \sup \frac{1}{n} \log E[\exp(\delta S_n)] \;\forall\; \delta \in \mathbb{R}. \tag{21}$$

Using results in Roitershtein (2007), we can now prove the following about the tails of the stationary distribution of $\{\phi_t\}_{t\in\mathbb{N}}$:

Proposition 1 *For π-almost every $d_0 \in [-a, a]$, there is a unique positive $\kappa < \infty$ that solves $\Lambda(\delta) = 0$, such that*

$$K_1(d_0) = \lim_{\tau\to\infty} \tau^\kappa P(\phi > \tau|d_0) \text{ and } K_{-1}(d_0) = \lim_{\tau\to\infty} \tau^\kappa P(\phi < -\tau|d_0). \tag{22}$$

and $K_1(d_0)$ and $K_{-1}(d_0)$ are not both zero.[14]

[12] For results on processes driven by finite state Markov chains see Saporta (2005).

[13] $\lim_{n\to\infty} \sup \frac{1}{n} \log E[\exp(\delta S_n)]$ is the Gartner Ellis limit that also appears in Large Deviation theory. For an exposition see Hollander (2000).

[14] *We can also show that $\pi(K_1(d_0) = K_{-1}(d_0)) = 1$ if a is large enough. This follows from Condition G given by Roitershtein (2007): Conditon G holds if there does not exist a partition of the irreducible set*

Proof. The results follow directly from Roitershtein (2007), Theorem 1.6 if we show the following:

(i) There exists a δ_0 such that $\Lambda(\delta_0) < 0$. First we note that $\Lambda(0) = 0$ for all n. Note also that

$$\begin{aligned}
\Lambda'(0) &= \lim_{n\to\infty} \sup \frac{1}{n} \frac{d \log E \prod_{t=1}^{n} |\lambda_t|^{\delta}}{d\delta} |_{\delta=0} \\
&= \lim_{n\to\infty} \sup \frac{1}{n} \left(E \prod_{t=1}^{n} |\lambda_t|^{\delta} \right)^{-1} E \left(\prod_{t=1}^{n} |\lambda_t|^{\delta} \log \prod_{t=1}^{n} |\lambda_t| \right) |_{\delta=0} \\
&= \lim_{n\to\infty} \sup \frac{1}{n} E \log \prod_{t=1}^{n} |\lambda_t|
\end{aligned}$$

For large n, as $\{\lambda_t\}_t$ converges to its stationary distribution ω, we have

$$\Lambda'(0) = \lim_{n\to\infty} \sup \frac{1}{n} \log E \prod_{t=1}^{n} |\lambda_t| = E_\omega \left(\log |\lambda_\infty| \right)$$

From equations (18)-(20) we have $E_\omega |\lambda_\infty| < 1$. Therefore $\Lambda'(0) = E_\omega \log(|\lambda_\infty|) < 0$, and

---

$D = \left\{ d \subset \left(\frac{-a}{1-\rho}, \frac{a}{1-\rho} \right) \right\}$ *into two disjoint sets* D_{-1} *and* D_1 *such that:*

$$\begin{aligned}
P(d &\in D_{-1}, \rho d + \varepsilon \in D_1, \lambda < 0) \\
&= P(d \in D_{-1}, \rho d + \varepsilon \in D_{-1}, \lambda > 0) = 0
\end{aligned}$$

where $\varepsilon \subset [-a, a]$ *and* $\rho \subset (0, 1)$. *(See Roitersthein's Definition 1.7 and subsequent discussion, and his Proposition 4.1.) Suppose in fact that* $P(d \in D_{-1}, \rho d + \varepsilon \in D_1, \lambda > 0) = 0$ *for* D_{-1} *with minimal element* d_0 *and maximal element* d_1. *Then* $P(d \in D_{-1}, \rho d + \varepsilon \in D_{-1}, \lambda > 0) = 1$. *Then it must be true, since* d_1 *is the maximum element of* D_{-1}, *that* $\rho d_1 + a \le d_1$ *and so* $\frac{a}{1-\rho} \le d_1$, *implying* $d_1 = \frac{a}{1-\rho}$. *Similarly, it must be true that* $\rho d_0 - a \ge d_0$ *so that* $\frac{-a}{1-\rho} \ge d_0$, *implying* $\frac{-a}{1-\rho} \ge d_0$. *Thus* $D_{-1} = D$, *that is the whole set. Now we can show that for a large enough,* $P(d \in D, \rho d + \varepsilon \in D, \lambda > 0) = 1$ *cannot hold. Since*

$$\lambda = 1 - g\,(d_0)^2 + g\beta d_0\,(\rho d_0 + \varepsilon) = 1 - g\,(d_0^2)\,(1 - \rho\beta) + g\beta d_0 \varepsilon,$$

we attain the smallest possible λ *if we set* $d_0 = \frac{a}{1-\rho}$ *and* $\varepsilon = -a$, *or equivalently* $d_0 = \frac{-a}{1-\rho}$ *and* $\varepsilon = a$. *Then* $\lambda \ge 0$ *with probability 1 if and only if* $a \le \bar{a} = \frac{(1-\rho)}{(g(1+\beta(1-2\rho)))^{0.5}}$. *If* $a > \bar{a}$ *with positive probability, then* $P(\lambda < 0) > 0$, *which contradicts* $P(d \in D_{-1}, \rho d + \varepsilon \in D_{-1}, \lambda > 0) = 1$. *Note also that* $\lambda = 1$ *for* $d_0 = 0$ *so it also follows that the* $P(\lambda > 0) > 0$.

there exists $\delta_0 > 0$ such that $\Lambda(\delta_0) < 0$.

(ii) There exists a δ_1 such that $\Lambda(\delta_1) > 0$. As in (i) above, we can evaluate, using Jensen's inequality,

$$\Lambda(\delta) = \lim_{n\to\infty} \sup \frac{1}{n} \log E \prod_{t=1}^{n} |\lambda_t|^{\delta} = \lim_{n\to\infty} \sup \frac{1}{n} \log E[\exp(\delta S_n)] \tag{23}$$

$$= \lim_{n\to\infty} \sup \log \left(E[\exp(\delta S_n)]\right)^{\frac{1}{n}} \geq \lim_{n\to\infty} \sup \log \left(E[\exp(\delta \frac{S_n}{n})] \right) \tag{24}$$

so that at the stationary distribution of $\{\lambda_t\}_{t\in\mathbb{N}}$

$$\Lambda(\delta) \geq \log E_{\omega}[\exp(\delta \log |\lambda_{\infty}|)] = \log \int_{\lambda} [\exp(\delta \log |\lambda_{\infty}|)] d\omega\left(\lambda\right). \tag{25}$$

As $\delta \to \infty$ for $\log |\lambda| < 0$ we have $[\exp(\delta \log |\lambda_t|)] \to 0$, but if $P_{\omega}\left(\log |\lambda| > 0\right) > 0$ at the stationary distribution of $\{\lambda_t\}_t$, then $\lim_{\delta\to\infty} \Lambda(\delta) = \log \int_{\lambda} [\exp(\delta \log |\lambda_t|)] d\omega\left(\lambda\right) \to \infty$. Therefore if we can show that $P_{\omega}\left(\log |\lambda_t| > 0\right) > 0$, it follows that there exists a δ_1 for which $\Lambda(\delta_1) > 0$. Since $\Lambda(\delta)$ is convex[15], it follows that there exists a unique κ for which $\Lambda(\kappa) = 0$. To show that $P_{\omega}\left(|\lambda| > 1\right) > 0$, define $A = \left\{ d \in \left(0, \frac{\mu a \beta}{1-\rho\beta}\right) \right\}$, $\mu \in (0,1)$ so that $\frac{\mu a \beta}{1-\rho\beta} < \frac{a}{1-\rho}$. At its stationary distribution $\{d_t\}_{t\in\mathbb{N}}$ is uniformly recurrent over $\left[\frac{-a}{1-\rho}, \frac{a}{1-\rho}\right]$ which implies that $P_{\pi}\left(d_{t-1} \in A\right) > 0$. We have $\lambda_t = 1 - \beta g d_{t-1}\left(\beta^{-1}(1-\rho\beta) d_{t-1} - \varepsilon_t\right)$, so for $d_{t-1} \in A$ and $\varepsilon_t \in (\mu a, a]$, it follows that $\lambda_t > 1$. Thus $P_{\omega}\left(|\lambda_t| > 1\right) = P_{\pi}\left(d_{t-1} \in A_t\right) P\left(\varepsilon_t \in (\mu a, a]\right) > 0$.

(iii) The non-arithmeticity assumption required by Roitershtein (2007) (p. 574, (A7))

[15] This follows since the moments of nonnegative random variables are log convex (in δ); see Loeve (1977, p. 158).

holds[16]: There does not exist an $\alpha > 0$ and a function $G : \mathcal{R} \times \{-1, 1\} \to R$ such that

$$P\left(\log|\lambda_t| \in G\left(d_{t-1}, \eta\right) - G\left(d_t, \eta \cdot sign\left(\lambda_t\right)\right) + \alpha\mathbb{N}\right) = 1. \tag{26}$$

We have

$$\log|\lambda_t| = \log\left|\left(1 - gd_{t-1}^2 + g\beta d_t d_{t-1}\right)\right| = \log\left|\left(1 - (1-\rho\beta)gd_{t-1}^2 + \beta g d_{t-1}\varepsilon_t\right)\right| = F\left(d_{t-1}, \varepsilon_t\right), \tag{27}$$

which contains the cross-partial term $d_t d_{t-1}$. Therefore in general $F\left(d_{t-1}, \varepsilon_t\right)$ cannot be represented in separable form as $R\left(d_{t-1}, \eta\right) - R\left(d_t, \eta\right) + \alpha\mathbb{N}\ \ \forall\left(d_{t-1}, d_t\right)$ where $d_t = \rho d_{t-1} + \varepsilon_t$. Suppose to the contrary that there is a small rectangle $[D, D^*] \times [E, E^*]$ in the space of (d, ε), over which λ remains of constant sign, say positive, such that $F(d, \varepsilon) = R(d) - R(\rho d + \varepsilon)$, d is in the interior of $[D, D^*]$, and ε is in the interior of $[E, E^*]$, up to a constant from the discrete set $\alpha\mathbb{N}$, which we can ignore for variations in $[D, D*] \times [E, E*]$ that are small enough. Now fix d, d' close to one another in the interior of $[D, D^*]$. We must have, for $\varepsilon \in [E + \rho|d - d'|,\ E^* - \rho|d - d'|]$, that

$$\begin{aligned} F(d, \varepsilon) - R(d) &= -R(\rho d + \varepsilon) = -R(\rho d' + \varepsilon + \rho(d - d')) & (28) \\ &= F(d', \varepsilon + \rho(d - d')) - R(d'), & (29) \end{aligned}$$

or $F(d, \varepsilon) - F(d', \varepsilon + \rho(d - d')) = R(d) - R(d')$. However the latter cannot hold since the cross-partial term $d_{t-1}\varepsilon_t$ in $F(d_{t-1}, \varepsilon_t) = 1 - (1 - \rho\beta)gd_{t-1}^2 + \beta g d_{t-1}\varepsilon_t$ is non-zero except of

[16] See also Alsmeyer (1997). In other settings $\{\lambda_t\}_t$ may contain additional *iid* noise independent of the Markov Process $\{d_t\}_t$, in which case the non-aritmeticity is much more easily satisfied.

a set of zero measure where d or ε are zero.[17,18]

(iv) To show that $K_1(d_0) = \lim_{\tau\to\infty} \tau^\kappa P(\phi > \tau|d_0)$ and $K_{-1}(d_0) = \lim_{\tau\to\infty} \tau^\kappa P(\phi < -\tau|d_0)$ are not both zero, we have to assure, since ψ_t and λ_t are not assumed to be independent, that ϕ is not a deterministic function of the initial d_{-1}. We invoke (a) and (c) of Proposition 8.1 in Roitersthein (2007): Condition 1.6, $\pi(K_1(d_0) + K_{-1}(d_0) = 0) = 1$, holds if and only if there exists there exists a measurable function $\Gamma : \left[\frac{-a}{1-\rho}, \frac{a}{1-\rho}\right] \to R$ such that

$$P(\psi_0 + \lambda_0\Gamma(\rho d_{-1} + \varepsilon_0) = \Gamma(d_{-1})) = 1.$$

However

$$\psi_0 + \lambda_0\Gamma(\rho d_{-1} + \varepsilon_0) = \theta g d_{-1}\rho d_{-1} + \theta g d_{-1}\varepsilon_0 + \left(1 - g d_{-1}^2 + g\beta d_{-1}(\rho d_{-1} + \varepsilon_0)\right)\Gamma(\rho d_{-1} + \varepsilon_0)$$

is a random variable that depends on ε_0 while $\Gamma(d_{-1})$ is a constant, so

$$P(\psi_0 + \lambda_0\Gamma(\rho d_{-1} + \varepsilon_0) = \Gamma(d_{-1})) < 1$$

---

[17] We thank Tomasz Sadzik for suggesting this proof for (iii).

[18] We can avoid possible degeneracies that may occur if λ_t and ψ_t have a specific form of dependence so that

$$P(\phi|\lambda_t\phi + \psi_t = \phi) = 1.$$

Note

$$\begin{aligned}\phi &= \frac{\psi_t}{1-\lambda_t} = \frac{\theta\rho g d_t^2 + \theta g d_t\varepsilon_{t+1}}{1-(1-\rho\beta)g d_t^2 + \beta g d_t\varepsilon_{t+1}} \\ &= \frac{\theta}{\beta}\frac{\beta\rho g d_t^2 + g b g d_t\varepsilon_{t+1}}{1-(1-\rho\beta)g d_t^2 + \beta g d_t\varepsilon_{t+1}}\end{aligned}$$

Differentiating with respect to ε_t, the right side is zero only if $\beta\rho g d_t^2 = 1-(1-\rho\beta)g d_t^2$, or $\beta\rho g = 1 - g + g\rho\beta$. This holds only if $g = 1$. So in general, for any d_t, there exists a constant ϕ such that $P(\phi|\lambda_t\phi + \psi_t = \phi) = 1$ only if $g = 1$, which we ruled out by assumption.

and Condition 1.6 in Roitersthein (2007) cannot hold. Then from Roitersthein (2007) Proposition 1.8 (c), $K_1(d_0)$ and $K_{-1}(d_0)$ are not both zero.[19] ■

The Proposition above characterizes the tail of the stationary distribution of ϕ as a power tail with exponent κ. It follows that the distribution of ϕ has moments only up to the highest integer less than κ, and is a 'fat tailed' distribution rather than a Normal. The results are driven by the fact that the stationary distribution of $\{\lambda_t\}_{t\in\mathbb{N}}$ has a mean less than one, which tends to induce a contraction towards zero, but also has support above 1 with positive probability, which tends to generate divergence towards infinity. The stationary distribution arises out of a balance between these two forces. Then large deviations as strings of realizations of λ_t above one, even though they may be rare events, can produce fat tails.

In the asset price model ϕ relates the dividends to assets prices. Under adaptive learning, the results above show how the probability distribution of large deviations, or 'escapes' of ϕ from its REE value is characterized by a fat tailed distribution, and will occur with higher likelihood than under a Normal.[20]

We now briefly discuss the case where $\{d_t\}_t$ is an $MA(1)$ process. Proposition 1 still applies and we obtain similar results to the $AR(1)$ case. Let

$$d_t = \varepsilon_t + \zeta\varepsilon_{t-1}, \ |\zeta| < 1, \ t = 1, 2... \tag{30}$$

[19] In models where the driving stochastic process is *iid* or is a finite stationary Markov chain, the exponent κ can be analytically derived using the results of Kesten (1973) and Saporta (2005). In the case where λ is *iid* in equation (13), κ solves $E(\lambda^\kappa) = 1$. In the finite markov chain case, under appropriate assumptions, κ solves $\varsigma(PA^\kappa) = 1$ where P is the transition matrix, A is a diagonal matrix of the states of the Markov chain assumed to be non-negative, and $\varsigma(PA^\kappa)$ is the dominant root of PA^κ.

[20] In the model of Cho et al. (2002), the monetary authority has a misspecified Philips curve and sets inflation policy to optimize a quadratic target. The learning algorithm using a constant gain however is not linear in the recursively estimated parameters (the natural rate and the slope of the Philips curve).

Then at its stationary distribution $d_t \in [-a(1+\zeta), a(1+\zeta)]$. Under the PLM

$$p_t = \phi_{0t}\varepsilon_t + \phi_{1t}\varepsilon_{t-1}, \tag{31}$$

after observing ε_t at time t but not ϕ_{1t+1}, the agents expect

$$E_t(p_{t+1}) = \phi_{0t}E_t(\varepsilon_{t+1}) + \phi_{1t}E_t(\varepsilon_t) = \phi_{1t}\varepsilon_t \tag{32}$$

Then the ALM is

$$p_t = \beta\phi_{1t}\varepsilon_t + \gamma(\varepsilon_t + \zeta\varepsilon_{t-1}) = [\beta\phi_{1t} + \gamma]\,\varepsilon_t + \gamma\zeta\varepsilon_{t-1}$$

and the REE is given by

$$\phi_0 = \gamma(1+\beta\zeta) \tag{33}$$

$$\phi_1 = \gamma\zeta. \tag{34}$$

Under the learning algorithm in equation (12) we obtain

$$\phi_{1t} = \phi_{1t-1} + gd_{t-1}(p_t - \phi_{1t-1}d_{t-1}) \tag{35}$$

$$\phi_{1t+1} = \lambda_{t+1}\phi_{1t} + \psi_{t+1} \tag{36}$$

$$\lambda_{t+1} = 1 - gd_t^2 + g\beta\varepsilon_{t+1}d_t \tag{37}$$

$$\psi_{t+1} = g\gamma\varepsilon_{t+1}d_t + \gamma\zeta g d_t\varepsilon_t \tag{38}$$

It is straightforward to show that at the stationary distribution of $\{\lambda_t\}_t$, $E(\lambda_t) < 1$, and that $P(\lambda_t > 1) > 0$. It is also easy to check that $\lambda_t > 0$ if $a < ((1+\zeta)(1+\zeta-\beta))^{-0.5}$. With the latter restriction, it is easy to check that the other conditions in the proof of Proposition 1 are satisfied.

4. Model Simulations and Comparative Statics

The theoretical results above indicate that rare but large shocks to the exogenous dividend process can throw off forecasts for the price-dividend ratio away from its rational expectation value. Of course escapes are more likely if the variance of the shocks to dividends are high. More critically, escapes in the long-run are possible if agents put a large weight on recent observations and discount older ones. The decay of the weights on past observations depends on the gain parameter g.[21] The size of the Bayesian optimal g will in turn depend on the drift that agents expect in the estimated parameter ϕ. We will estimate g in the next section, both directly, and also from the perspective of Bayesian agents expecting a random walk drift in ϕ.

In this section we explore how κ is related to the underlying parameters of our model. We can simulate the learning algorithm that updates ϕ, and then estimate κ from the simulated data using a maximum likelihood procedure following Clauset et al. (2009). We can then explore how κ varies as we vary model parameters. We simulate 1000 series, each of length 5000, for ϕ_t under the $AR(1)$ assumption for dividends with *iid* uniform shocks. We then feed the simulated series into the model to produce $\{P_t\}$ and $\{P_t/D_t\}$. We estimate κ for

[21] Under constant gains the decay in weights on past observations dating i periods back is given by $(1-g)^{i-1}$.

each simulation and produce an average κ.

Escapes or large deviations in prices will take place when sequences of large shocks to dividends throw off the learning process away from the rational expectations equilibrium. Such escapes will be more likely if dividend shocks can produce values of λ_t above 1, as we can see from equations (35-38). We expect lower κ, or fatter tails, as the support of λ_t that lies above 1 gets larger.

In the $AR(1)$ case for dividends we have $\lambda_{t+1} = 1 - (1 - \rho\beta) g d_t^2 + \beta g d_t \varepsilon_{t+1}$. Given the stationary distribution of $\{d_t\}_t$ and that of $\{\varepsilon_t\}_t$, the support of λ_t above 1 unambiguously increases if β increases. In principle increasing ρ can have an ambiguous effect: while the term $(1-\beta\rho)$ declines and tends to raise λ_t for realizations of d_t and ε_{t+1}, the support of the stationary distribution of $\{d_t\}_t$ gets bigger with higher ρ. While this can increase $(1-\rho\beta)gd_t^2$ and reduce the support of λ that is above 1 for large realizations of d_t^2, in our simulations the former effect seems to dominate. Finally we expect that decreasing g will shrink the support of λ_t that is above 1 so that κ increases with g: as the gain parameter decreases, the tails of the stationary distribution of $\{\phi_t\}$ get thinner.[22]

We use a baseline parameterization, $(\rho, g, \beta, \gamma) = (0.80, 0.4, 0.95, 2.5)$ based on estimates that we obtain in the next section. The estimated parameters, except for g, are in line with standard calibrations. The discount factor of $\beta = 0.95$ is consistent with annual data and an annual discount rate of about 5%. While empirical estimates of g are hard to come by, the

[22] This of course is in accord with the Theorem 7.9 in Evans and Honkapohja (2001). As the gain parameter $g \to 0$ and $tg \to \infty$, $\{\phi_t^g - \varkappa\}/g^{0.5}$ converges to a Gaussian variable where $\varkappa$ is the globally stable point of the associated ODE describing the mean dynamics. More generally, as $g \to 0$, the estimated coefficient under learning with gain parameter g, ϕ_t^g, converges in probability (but not uniformly) to $\varkappa$ for $t \to \infty$. However, there will always exist arbitrarily large values of t with ϕ_t^g taking values remote from $\varkappa$ (See Benveniste, Métivier and Priouret (1980), pp. 42-45). Note however that our characterization of the tail of the stationary distribution of $\{\phi_t\}_t$ and of κ is obtained for fixed $g > 0$.

usual values of g used in theoretical models are much smaller, in the order of 0.01 or 0.04, suggesting a very slow decay in the weights attached to past observations. Values of g in the range of $0.3-0.5$ indicate a high decay rate, suggesting a propensity for the agents to think that "this time it's different". As noted above, we attempt to estimate g in the context of our model by two separate methods in the next section. However, as the comparative statics in Figure 4 below demonstrate, for the learning model to explain the fat tails and the high variance of the P/D ratio, the gain parameter has to be large enough. This also implies, as discussed further in the next section, that the expected drift in the estimated parameters should have a large variance.

For the baseline parametrization we set the value of $a = 0.33$ to match the standard deviation of the detrended dividends in the data. We find that the average κ is 5.0210, the average (P_t/D_t) is 20.6274 and the average standard deviation of (P_t/D_t) is 9.8934. We then vary each element of $(\rho, g, \beta, \gamma, \alpha)$ while keeping the others at their baseline values. The results of varying each parameter around the baseline values are plotted in Figures 3 and 4 below.[23]

[23] The restriction given by equation (18) implies a maximum value of $a = \hat{a} = 4.2733$. For all parameter values used to produce Tables 3 and 4, the restriction is easily satisfied.

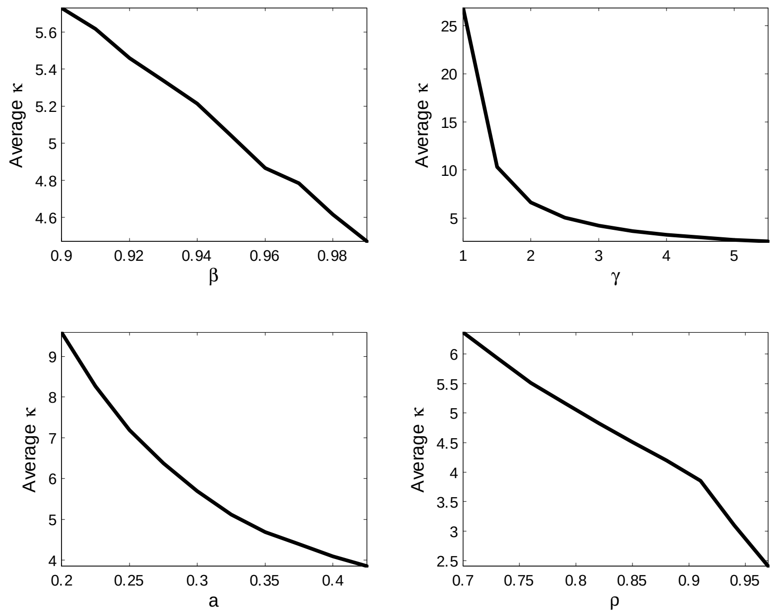

Figure 3. Simulation Results.

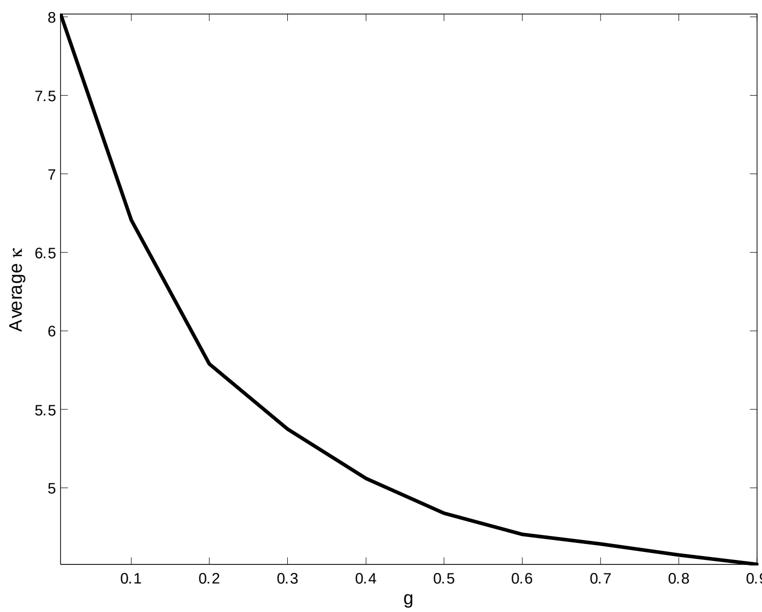

Figure 4. Simulation Results (cont'd.).

The simulation results confirm the notion that the average κ's should decline with β, γ and a. Figure 4 plots the results of the critical learning parameter g; it clearly demonstrates that as the learning gain falls, that is, the horizon for learning increases, the average κ

rises. In summary, SGCG learning leads to large deviations of (P_t/D_t) from its rational expectations value.

5. Empirics

We first check whether real world data on price-dividend ratios have fat tails. We use the maximum likelihood procedure following Clauset et al. (2009) to estimate κ associated with P_t/D_t for both S&P and CRSP dividend series plotted in Figures 1 and 2 above. The results provided in Table 1 below show fairly small values of κ for both series, suggesting that only the first few moments of P_t/D_t exist irrespective of the data source. Table 1 also reports the estimated persistence ρ under an $AR(1)$ specification for the two linearly detrended dividends series, alongside the average price-dividends ratio (P_t/D_t) and its standard deviation.

Table 1. Data Characteristics

	S & P 500	CRSP
$\widehat{\kappa}$	3.6914	5.5214
$s.e.(\widehat{\kappa})$	0.3828	2.6046
$\widehat{\rho}$	0.7891	0.7519
$s.e.(\widehat{\rho})$	0.0523	0.0777
Mean (P_t/D_t)	25.5211	26.1805
Std. Dev. (P_t/D_t)	13.1758	9.3298

We use two separate approaches to get estimates for the gain parameter g. First we feed the actual S&P and CRSP dividend series into our learning model and estimate the parameters, $\vartheta = [g \ \gamma \ \beta \ \rho]$ by minimizing the squared difference between the empirical κ's reported in Table 1 and those generated by our model. That is, we implement a simulated minimum distance method to estimate ϑ as[24]

$$\min_{\vartheta} \; [\kappa - \kappa(\vartheta)]^2. \tag{39}$$

This estimation process necessarily puts a great deal of emphasis on the tail of the empirical data given by κ. Since the puzzle lies in the fat tail and high variance of P/D, emphasizing the empirical the tail in the estimation method may be justified. The parameter estimates other than g are certainly in line with basic calibrations in the literature, but the value of g, as expected from our model, is higher than the usual values of $0.01 - 0.04$ that we find in the literature.

The minimization procedure proceeds as follows. For candidate parametrization of ϑ we employ the S&P or CRSP series dividends d_t to calculate ϕ_t as per (13)-(15). The ALM (9) then produces a corresponding p_t series which in turn delivers a price-dividend ratio P_t/D_t. We then estimate the κ associated with the 'simulated' P_t/D_t, using the methods of Clauset et al. (2009) to produce the $\kappa(\vartheta)$. The minimization procedure searches over the parameter space of ϑ to implement (39). Table 2 below reports the estimates and associated standard errors for each of the S&P or CRSP dividend series. We also report associated κ values obtained by simulating prices using the estimated parameters and the actual dividend data.[25]

---

[24] Minimization was conducted using a simplex method and standard errors were computed using a standard inverse Hessian method.

[25] Starting values for the minimization procedure were $\vartheta_0 = [0.5 \; 2.5 \; 0.95 \; 0.75]$.

Table 2. Parameter Estimates

	S & P 500		CRSP	
Parameter	Estimate	Std. Err.	Estimate	Std. Err.
g	0.3468	2.7158	0.5257	0.4722
γ	2.6503	1.7481	2.4598	0.6259
β	0.9615	0.3870	0.8984	0.4576
ρ	0.8729	0.0552	0.7959	0.1355
Associated κ	2.4128		5.5214	

The point estimates of g, ranging from 0.35 to 0.53 are high, although the standard errors are quite large in the case of the S & P 500 dataset. Carceles-Poveda and Giannitsaraou (2008) discuss possible values of g. Looking at standard deviations of the price-dividend ratios for the Lucas asset pricing model, they report that the standard deviations generated by the rational expectations or the learning models are smaller than the standard deviations in the actual data by factors of about 20 to 50. Note that our estimates of the parameter values, including g, are very close to those used by Carceles-Poveda and Giannitsaraou (2008) in their simulations except for γ, the CRRA parameter: they set $\gamma = 1$ while we have it at $\gamma = 2.5$. Note also that for our simulations of Figure 3 κ drops dramatically with γ.

For our second approach to pin down the gain parameter we let the agent optimally determine g by estimating the standard deviations of the parameter drift, the noise in the P/D ratio, and the shock to the dividend process.[26] Recall that under SGCG learning ϕ_t

[26] See Sargent, Williams and Zha (2006) and others for a more complex version of this approach requiring dynamic tracking estimation.

evolves as

$$\phi_t = \phi_{t-1} + g d_{t-1}(p_t - \phi_{t-1} d_{t-1}), \quad g \in (0,1) \tag{40}$$

Consider the case in which the agents assume that the PLM is

$$p_t = \phi_{t-1} d_{t-1} + \xi_t,\ \ \xi_t \sim iid(0, \sigma_\xi^2),\ \ \sigma_\xi^2 < +\infty \tag{41}$$

with the coefficient ϕ drifting according to a random walk:

$$\phi_t = \phi_{t-1} + \Lambda_t,\ \ \Lambda_t \sim iid(0, \sigma_\Lambda^2),\ \ \sigma_\Lambda^2 < +\infty \tag{42}$$

In this case, the Bayesian agent would use (40) and estimate σ_Λ, σ_d and σ_ξ to set an optimal estimate of the gain in the limit as

$$g = \frac{\sigma_\Lambda \sigma_d}{\sigma_\xi} \tag{43}$$

where σ_d denotes the standard deviation of d_t (see Evans et al (2010)). Under this approach, the long-run value of g that generates $\{p\}$ and $\{\phi\}$ under adaptive learning would be self-confirming in the sense that agents would estimate g using (43).

To compute (43) an estimate of σ_d is of course readily obtained from the dividend data. However we need to specify a method for the agents to compute estimates of σ_Λ and σ_ξ. If we recognize the system above as being analogous to a time varying parameter formulation, then employing the methods laid out in Kim and Nelson (1999) we can obtain estimates of σ_Λ and σ_ξ.[27] We report these results in Table 3 below.

[27] Given our estimate of $\beta = 0.95$ we convert the CRSP data to annual, summing dividends quarterly dividends for each year. For the S&P 500 we use the annual data reported by Shiller (1999), pp.439-441.

Table 3. Drifting Beliefs Model Parameter Estimates

	S & P 500		CRSP	
Parameter	Estimate	Std. Err.	Estimate	Std. Err.
σ_Λ	0.8122	0.7718	0.8588	0.2963
σ_ξ	0.3157	0.0230	0.2596	0.0291
$\log L$	-61.4102		-17.5256	
σ_d	0.1892		0.1649	
Associated g	0.4866		0.5455	

These estimates suggest values of the gain significantly larger than those usually assumed in the literature. Looking at Figure 4, a value of $g = 0.4866$ yields a tail estimate κ of about 4.9 while a value of $g = 0.5455$ yields a κ of about 4.75, compared to κ in the data ranging from 3.7 to 5.5 in Table 1. We also simulated the model with baseline parameter values but with gains of 0.4866 and 0.5455. These simulations resulted in average price-dividend ratios of 20.6324 and 20.6965 respectively with corresponding standard deviation values of 10.0051 and 10.5870.

Finally, instead of using actual P and D data series, we generate data by simulating our model with our benchmark values $(\rho, g, \beta, \gamma) = (0.80, 0.4, 0.95, 2.5)$, and then compute g from (43) using the methods in Kim and Nelson (1999).[28] The average g is 0.3826, which is quite close to and confirms the benchmark value of $g = 0.4$ that is used in generating the simulated data.

[28] We run 1000 simulations each with 5000 periods, and obtain the average g from (43) across the 1000 simulations.

6. Conclusion

An important and growing literature replaces expectations in dynamic stochastic models not with realizations and unforecastable errors, but with regressions where agents 'learn' the rational expectations equilibria. When such agents employ constant gain learning algorithms that put heavier emphasis on recent observations (which is optimal when there is drift in estimated parameters), escape dynamics can propel estimated coefficients away from the REE values. In an asset pricing framework 'bubbles', or asset price to dividend ratios that exhibit large deviations from their REE values (even though our model has presumed a no-bubble condition) can occur with a frequency associated with a fat tailed power law, as observed in the data. The techniques used in our paper can be generalized to higher dimensions, to finite state Markov chains, to continuous time,[29] and can be applied to other economic models that use constant gain learning.

References

[1] Adam, K., Marcet, A., and J. P. Nicolini. 2008. Stock Market Volatility and Learning, *European Central Bank Working Paper Series*, No. 862.

[2] Adam, K. and Marcet, A., 2011.Internal rationality, imperfect market knowledge and asset prices *Journal of Economic Theory*, 146, 1224-1256..

[3] Alsmeyer, G. 1997. The Markov Renewal Theorem and Related Results. *Markov Process Related Fields* 3 103–127.

[29] See for example Saporta (2005), Saporta and Yao (2005), and Ghosh et al. (2010).

[4] Barro, R. J. 2009. Rare Disasters, Asset Prices, and Welfare Costs, *American Economic Review*, 99:1, 243–264.

[5] Benhabib, J. Bisin, A. and S. Zhu, 2011. The Distribution of Wealth and Fiscal Policy in Economies with Finitely Lived Agents. *Econometrica* 79, 123-158.

[6] Benhabib, J., 2010. A Note Regime Switching, Monetary Policy and Multiple Equilibria, *NBER Working Paper* No. 14770.

[7] Benveniste, A., Métivier, M. and P. Priouret, 1980. *Adaptive Algorithms and Stochastic Approximations,* Springer-Verlag, New York.

[8] Branch, W., and Evans, G. W., 2010, Asset Return Dynamics and Learning, *Review of Financial Studies*, 1651-1680.

[9] Brennan, M. J., Xia, 2001, Stock price volatility and equity premium, *Journal of Monetary Economics,* 47, 249-283.

[10] Bullard, J., Duffy, J., 2001. Learning and excess volatility. *Macroeconomic Dynamics* 5, 272-302.

[11] Campbell, J. Y., 2003. Consumption-Based Asset Pricing, *Handbook of the Economics of Finance*, George Constantinides, Milton Harris, and Rene Stulz eds., North-Holland, Amsterdam.

[12] Carceles-Poveda, E., Giannitsarou, C., 2007. Adaptive Learning in Practice. *Journal of Economic Dynamics and Control* 31, 2659-2697.

[13] Carceles-Poveda, E., Giannitsarou, C., 2008. Asset pricing with adaptive learning. *Review of Economic Dynamics* 11 629–651.

[14] Chevillon, G., Mavroeidis, S., "Learning Generates Long Memory," http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1969602

[15] Cho, I-K., Sargent, T. J., Williams, N., 2002. Escaping Nash Inflation. *Review of Economic Studies* 69, 1-40.

[16] Clauset, A., Shalizi, C. R. and M. E. J. Newman, 2009. Power-law distributions in empirical data, *SIAM Review* 51(4), 661-703

[17] Collamore, J. F. , 2009. Random Recurrence Equations and Ruin in a Markov-Dependent Stochastic Economic Environment, *Annals of Applied Probability* 19, 1404–1458.

[18] Cogley, T. and Sargent, T. J., 2008, The market price of risk and the equity premium: A legacy of the great depression?, *Journal of Monetary Economics*, 55, 454–476.

[19] Evans, G., Honkapohja, S., 1999. Learning Dynamics. *Handbook of Macroeconomics*, Vol.1, eds. J. Taylor and M. Woodford, 1999, North-Holland, pp.449-542.

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

[21] Evans, G., Honkapohja, S., and N. Williams, 2010. Generalized Stochastic Gradient Learning, *International Economic Review*, 51, 237-262.

[22] Gabaix, X., Gopikrishnan, P., Plerou, V. and Stanley, H. E. 2006. Institutional Investors and Stock Market Volatility, *Quarterly Journal of Economics*, 121 (2), p. 461-504.

[23] Gabaix, X., 2009. Power Laws in Economics and Finance, *Annual Review of Economics*, 1, p. 255-93.

[24] Goldie, C. M., 1991. Implicit Renewal Theory and Tails of Solutions of Random Equations, *Annals of Applied Probability*, 1, 126–166.

[25] Ghosh, A. P, Haya, D., Hirpara, H., Rastegar, R., Roitershtein, A.,Schulteis, A., and Suhe, J, (2010), Random Linear Recursions with Dependent Coefficients, *Statistics and Probability Letters* 80, 1597 1605.

[26] Holmstrom, B., 1999, Managerial Incentive Problems: A Dynamic Perspective, The Review of Economic Studies, 66, 169-182.

[27] Hollander, F. den, (2000), *Large Deviations*, Fields Institute monographs,American Mathematical Society, Providence, Rhode Island.

[28] Kesten, H., 1973. Random Difference Equations and Renewal Theory for Products of Random Matrices, *Acta Mathematica*. 131 207–248.

[29] Koulovatianos, C. and V. Wieland, 2011. Asset Pricing under Rational Learning about Rare Disasters, Manuscript.

[30] Loeve, M. 1977. Probability Theory, 4^{th} Ed., Springer, New York.

[31] Lucas, R. E. Jr., 1978. Asset Prices in an Exchange Economy, *Econometrica*, Vol. 46, No. 6. (Nov., 1978), pp. 1429-1445.

[32] Marcet, A. and Sargent, T. J., 1989. Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models, *Journal of Economic Theory*, 48, 337-368.

[33] Nummelin, E., 1984. *General irreducible Markov chains and non-negative operators*. Cambridge Tracts in Mathematics 83, Cambridge University Press.

[34] Roitershtein, A., 2007. One-Dimensional Linear Recursions with Markov-Dependent Coefficients, *The Annals of Applied Probability*, 17(2), 572-608.

[35] Saporta, B., 2005. Tail of the Stationary solution of the Stochastic equation $Y_{n+1} = a_n Y_n + \gamma_n$ with Markovian Coefficients, *Stochastic Processes and their Applications*, 115(12), 1954-1978.

[36] Saporta, B. and Yao, J-F, 2005, Tail of a Linear Diffusion with Markov Switching, *The Annals of Applied Probability*, 992–1018.

[37] Sargent, T. J., 1999. *The Conquest of American Inflation*. Princeton University Press.

[38] Sargent, T. J. and Williams, N., 2005. Impacts of Priors on Convergence and Escape from Nash Inflation, *Review of Economic Dynamics,* 8(2), 360-391.

[39] Shiller, R. J., 1999, *Market Volatility*, 6th printing, MIT Press, Cambridge.

[40] Shiller, R. J., 2005. *Irrational Exuberance*, 2^{nd} edition, Broadway Books.

[41] A. Timmermann, 1993 How Learning in Financial markets generates Excess Volatility and Predictability in Stock Prices, *Quarterly Journal Economics,*108,1135–1145.

[42] A. Timmermann, 1996, Excess Volatility and Predictability of Stock Prices in Autoregressive Dividend Models with Learning, *Review of Economic Studies*, 63, 523–557.

[43] M. L. Weitzman, 2007, Subjective Expectations and Asset-Return Puzzles, American Economic Review, 97, 1102–1130.

[44] Williams, N. 2009. Escape Dynamics, Manuscript.

[45] Woodford, M., 1990. Learning to Believe in Sunspots, *Econometrica*, 58(2), 277-307.

7. Data Appendix

1. Annual S&P 500 Dataset

   (a) The following time series are extracted/constructed for 1871 through 2009 (note that $t = 1, \ldots T$ where $T = 2009.12$):

      i. Extract S & P Comp ($\widetilde{P}(t)$).

      ii. Extract Dividend ($\widetilde{D}(t)$).

      iii. Extract Consumer Price Index ($CPI(t)$).

      iv. Construct Real Price ($P(t)$) as $P(t) = [\widetilde{P}(t) \times \ CPI(T)]/CPI(t)$.

      v. Construct Real Dividend ($D(t)$) as $D(t) = [\widetilde{D}(t) \times \ CPI(T)]/CPI(t)$.

   (b) Construct the Price to Dividends Ratio (ratio) as $P(t)/D(t)$.

2. Quarterly CRSP Dataset

   (a) Download the quarterly data from http://scholar.harvard.edu/campbell/data, where the particular data being used is associated with "Replication Data for: Consumption Based Asset Pricing". The relevant file is titled USAQE.ASC, note that this is effectively a CRSP dataset with the relevant variables being VWRETD and

VWRETX. The text below is an extract from the explanations for this dataset on the above website.

(b) The following quarterly time series are extracted/constructed for 1926.1 through 1998.4 from the above dataset (note that $t = 1, \ldots T$ where $T = 1998.4$):

i. Extract Col. 2: $\widetilde{P}(t)$. For each month, the price index is calculated as $\widetilde{P}(t) = (VWRETX(t) + 1) \times \widetilde{P}(t-1)$. (Note that time t in this equation is in months.) The price index for a quarter, as reported in this column, is the price index for the last month of the quarter. The original data, which goes up to 1996.4 was not altered. The new data, which goes up to 1998.4, was created as described here starting from 1997.1.

ii. Extract Col. 3: $\widetilde{D}(t)$. Dividend in local currency, calculated as follows. The dividend yield for each month is calculated as $\widetilde{DY}(t) = [1+VWRETD(t)]/[1+VWRETX(t)]-1$. Note that if the return index is calculated from $VWRETD$ as above, then this formula agrees with the formula for the dividend yield given earlier. As before, the dividend for each month is calculated as $\widetilde{D}(t) = \widetilde{DY}(t) \times \widetilde{P}(t)$. The dividend for a quarter, as reported in this column, is the sum of the dividends for the three months comprising the quarter.

iii. Extract the Consumer Price Index from Shiller's Monthly Data ($CPI(t)$) which is monthly and associate the last month of a quarter as a quarterly $CPI(t)$.

iv. Construct Real Price ($P(t)$) as $P(t) = [\widetilde{P}(t) \times \ CPI(T)]/CPI(t)$. Take the last price of a quarter as the annual price.

v. Construct Real Dividend ($D(t)$) as $[\widetilde{D}(t) \times\ CPI(T)]/CPI(t)$ and then take quarterly sums to get $D(t)$ at an annual frequency.

vi. Construct the Price to Dividends Ratio (ratio) as $P(t)/D(t)$.