Learning about Risk and Return: A Simple Model of Bubbles and Crashes[*]

William A. Branch
University of California, Irvine

George W. Evans
University of Oregon

January 31, 2008

Abstract

This paper demonstrates that an asset pricing model with least-squares learning can lead to bubbles and crashes as endogenous responses to the fundamentals driving asset prices. When agents are risk-averse they generate forecasts of the conditional variance of a stock's return. Recursive updating of the conditional variance and expected return implies two mechanisms through which learning impacts stock prices: occasional shocks may lead agents to lower their risk estimate and increase their expected return, thereby triggering a bubble; along a bubble path recursive estimates of risk will increase and crash the bubble.

JEL Classifications: G12; G14; D82; D83
Key Words: Risk, asset pricing, bubbles, adaptive learning.

Thus, this vast increase in the market value of asset claims is in part the indirect result of investors accepting lower compensation for risk. Such an increase in market value is too often viewed by market participants as structural and permanent ... Any onset of increased investor caution elevates risk premiums and, as a consequence, lowers asset values and promotes the liquidation of the debt that supported higher asset prices. This is the reason that history has not dealt kindly with the aftermath of protracted periods of low risk premiums.

Alan Greenspan (2005).

1

1 Introduction

In his classic study of financial crises, Kindleberger (1977) provides an accounting of historical episodes of manias and panics. Kindleberger's conjecture for why bubbles – and, their subsequent crashes – arise places primary emphasis on abrupt and unanticipated changes in expectations, in part a response to a sudden economic event. This explanation is in line with the view of many financial market observers that during the mid to late 1990's U.S. stock prices were excessively high – a "bubble". The existence and detection of bubbles in asset prices has long been of interest to economists and, recently, monetary policymakers (Bernanke (2002)). Early empirical evidence in favor of stock market bubbles were provided by Shiller (1981) and LeRoy and Porter (1981).

Despite popular agreement that asset prices are susceptible to large run-ups in prices above the value warranted by observed fundamentals, in the economics literature there is no such consensus. Blanchard and Watson (1982) propose a theory of rational bubbles in which agents' (rational) expectations are influenced in part by extrinsic random variables whose properties accord to historical bubble episodes. West (1987), Froot and Obstfeld (1991) and Evans (1991) construct rational bubbles that periodically explode and collapse.[1] Other recent approaches include Hong, Scheinkman, and Xiong (2005,2006), Santos and Woodford (1997), and Allen, Morris, and Postelwaite (1993). A controversial issue for rational bubbles is that the trigger for the bubble collapse is modeled by an exogenous sunspot process. Moreover, Evans and Honkapohja (2001) show that rational bubbles would not be attainable if rational expectations were replaced with private-sector least-squares learning.

In this paper, we reconsider the issue of recurrent bubbles and crashes and demonstrate that a model, based on econometric learning, can generate bubbles and crashes as endogenous responses to fundamental shocks. We follow Evans and Honkapohja (2001) and replace rational expectations (RE) in a simple linear asset pricing model with a perceived law of motion that has a reduced form consistent with RE and the parameters of which are updated using recursive least squares. We extend the conventional model to include a motive for agents to estimate risk – measured as the conditional variance of net stock returns. We show that the dynamic properties of the economy are altered in a surprising and interesting way once agents must account for, and adaptively learn, the riskiness of stocks.

Figure 1 previews our results by plotting stock prices generated from our simple asset pricing model in which rational expectations are replaced by a (constant gain) econometric forecasting rule. In a fundamentals based rational expectations equilibrium the mean stock price is parameterized to be about 8.7, and along an

---

[1]Diba and Grossman (1988) show that, under rational expectations, bubbles can never restart if they collapse to zero. Thus periodically collapsing rational bubbles are constructed so that they always remain positive.

equilibrium path price is simply a constant plus white noise. With constant-gain learning, however, the dynamics may undergo an abrupt change leading to recurrent bubbles and crashes. The purpose of this paper is to demonstrate that a simple asset pricing model, under adaptive learning, generates the recurrent dynamics illustrated in Figure 1.

Figure 1: Simulated stock price dynamics.

$\beta = 0.95$, $a = 0.75$, $\sigma_u = 0.9$, $\sigma_v = 0.5$, $y_0 = 1.5$, $s_0 = 1$, $\gamma_1 = 0.01$, $\gamma_2 = 0.04$
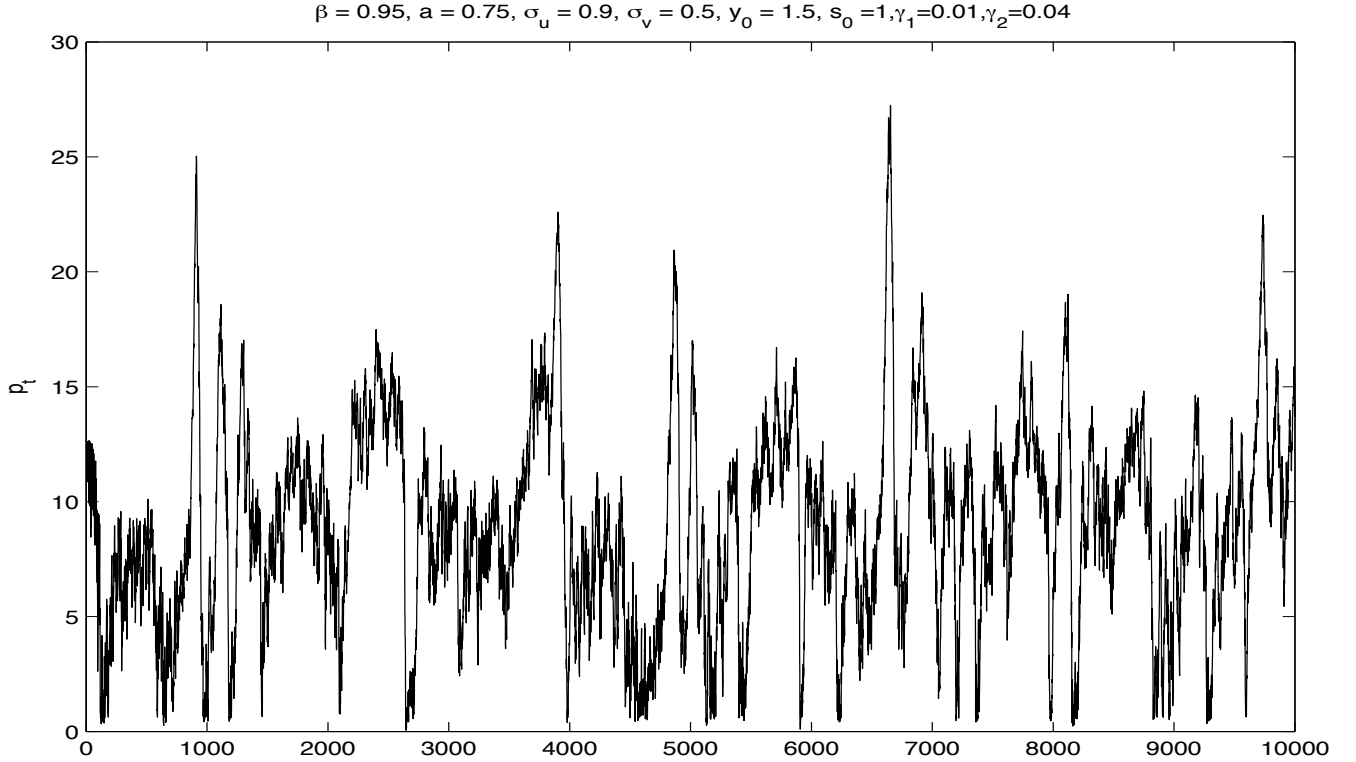
The analysis in this paper identifies two channels through which risk and return estimates affect stock prices. First, occasional shocks to fundamentals may lead agents to suddenly lower their estimate of risk and increase their expected return; combined these two forces cause stock prices to quickly increase above their fundamental values. Second, real-time estimates of risk are also useful in explaining how bubbles crash suddenly. As stock price follows a bubble path, estimates of risk will increase until the perceived riskiness of the asset is too high, at which point demand collapses and price crashes. The primary result of this paper is to demonstrate that a simple model that incorporates adaptive learning can generate recurrent bubbles and crashes.

The onset of bubbles and crashes, as illustrated in Figure 1, is reminiscent of the escape dynamics identified by Sargent (1999), Cho, Williams, and Sargent (2002), and Williams (2004). We assume agents form expectations of stock prices via a simple autoregression – whose form nests the fundamentals and bubbles based REE – with

3

parameters updated adaptively via discounted least squares. Following Sargent, we show that adaptive beliefs introduce serial correlation that would not otherwise exist, and that for some sequences of shocks, agents' forecasting rule begins to track this serial correlation via a random walk forecasting model. Importantly, this "escape" from a serially uncorrelated process to a serially correlated time-series, well approximated by a random walk, arises endogenously. It is this sudden shift in beliefs that leads to recurrent bubbles and crashes.

This paper is closely related to other recent approaches that employ adaptive learning as a mechanism for generating sudden changes in an economy's dynamical behavior. Most relevant are papers by Marcet and Nicolini (2003) and Sargent, Williams, and Zha (2006b) who show that a model with constant gain, or perpetual, learning can be useful in explaining recurrent hyperinflations. In these models a hyperinflation is similar to a bubble in the sense that hyperinflationary paths are unstable under least-squares learning. The key to our results is that agents also estimate the riskiness of the asset by updating in real-time their forecast of the conditional variance of stock returns.

We believe that recurrent bubbles and crashes are a feature of the data. Figure 2 plots the monthly S&P 500 price index from 1973-2006. In order to map them into our analytical framework, the data (in logs) have been detrended using the Hodrick-Prescott filter.[2] De-trended log of monthly stock prices are marked by repeated large run-ups above trend and sudden crashes. The goal of this paper is to derive a model where such recurrent bubbles and crashes, i.e. changes in price not explained by fundamentals, arise as an endogenous response to the fundamental shocks of the economy.

Although the model presented here is too simple to be taken directly to data, we argue that it has features consistent with empirical regularities. It is well-known that in U.S. stock data explosive rational bubbles are not detected using unit root and cointegration tests. However, Wu and Xiao (2002) derive an alternative Recursive Estimate Statistic (RES) test meant to capture the temporary expansionary phases that occur when stocks are subject to recurrent bubbles and crashes. Thus, one implication of our model is that we would expect, for data generated by our model under learning dynamics, that bubbles would not be detectable using unit root tests, but that they would be detectable using RES tests. We confirm this intuition through Monte Carlo simulations.

Our channel for deriving bubbles and crashes as responses to fundamentals is constant-gain least-squares learning. There has been an explosion of research applying constant-gain learning in monetary models, e.g. Sargent (1999), Cho, Williams, and Sargent (2002), Sargent, Williams, and Zha (2006a,b), Orphanides and Williams

---

[2]This figure should be taken as only illustrative. The non-detrended time-series of prices show similar periods of rapid increases and crashes.

Figure 2: Monthly Real S&P 500 price index (in logs) from 1973.1-2006. Data have been detrended using the Hodrick-Prescott Filter
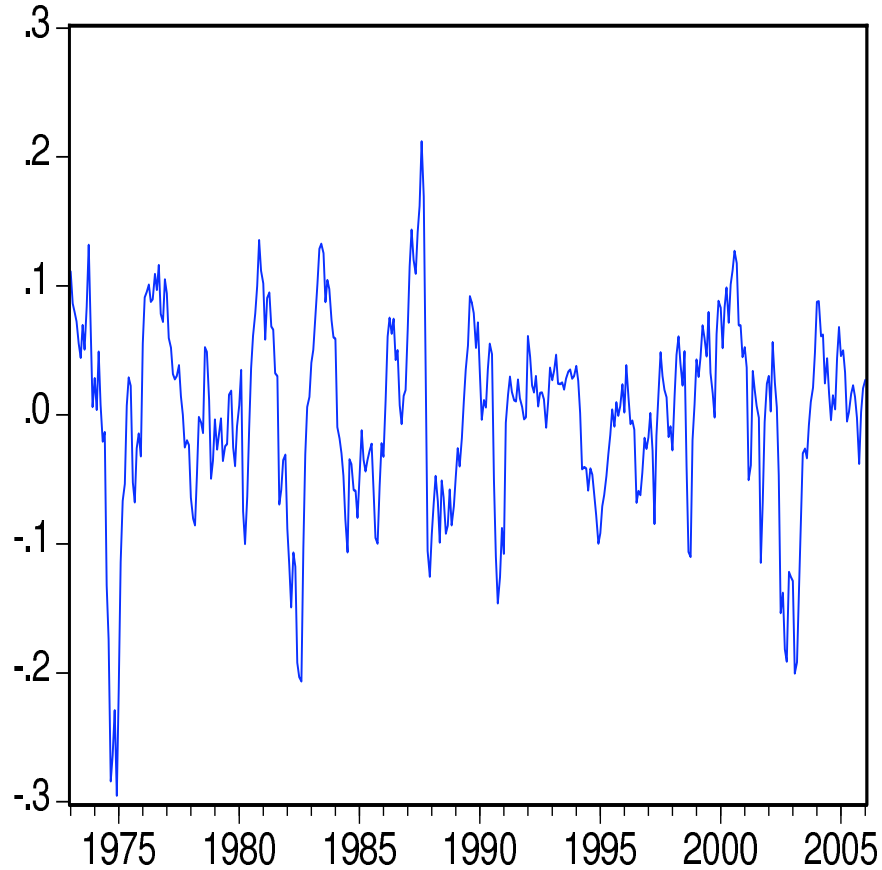
(2005b), McGough (2006). With constant-gain least-squares, agents adopt a linear forecasting rule that is consistent in reduced-form with a rational expectations equilibrium and update their parameter estimates as the economy evolves. Since the model is self-referential it is not obvious whether the dynamics will settle down. The articles cited above demonstrate that it is possible for the learning dynamics to generate persistent escapes from their equilibrium values while eventually returning to a neighborhood of the equilibrium. Along these escape paths (see Williams (2004)) the qualitative nature of the dynamics may be very different from RE or a decreasing gain least-squares learning framework.

We consider a simple asset pricing model in which the stock price today depends on expected *cum* dividend price next period and negatively on share supply, meant to proxy asset float as discussed in Cochrane (2005). Share supply and dividends both

5

follow exogenous iid processes. We assume agents are risk-averse so that they seek to forecast both the expected net return tomorrow as well as the conditional variance of stock returns. It turns out that by requiring agents to also estimate conditional variance, the global dynamics of the model under adaptive learning are dramatically different and indeed can lead to recurrent bubbles and crashes.

Our central insight is that an aversion to risk and adaptive learning can lead to both bubbles and crashes. We show that an aversion to risk is critical for igniting, as well as popping, a bubble. Agents' risk-aversion implies that they are concerned with the variance of stock returns, not just the mean. However, we also assume agents do not know the precise equilibrium values of these moments and are forced to estimate them in real-time. Because stock prices are self-referential, risk estimates feed back onto price, and this in turn affects forecasts of expected returns. A critical insight of this paper is that under econometric learning, the forecasts of risk and expected return are jointly determined through the self-referential nature of the asset pricing model. Agents that engage in constant-gain learning place greater emphasis on recent forecast errors. Occasional shocks to dividends and share supply can lead traders to revise their risk estimates. Through the subsequent effect on stock price, agents revise their estimate of mean stock returns in such a way that it places the economy onto a bubble-like path. Importantly, risk also plays a central role in the collapse of bubbles. As the stock price follows a bubble path, the risk-estimate grows to excessively high levels. As a result, demand for the risky asset eventually collapses, leading to a stock-price crash. Thus, our simple model of adaptive learning creates a coherent story of recurrent bubbles and crashes as endogenous responses to fundamental shocks.

There is an extensive literature that models the existence of bubbles in settings with heterogeneous expectations. Usually these models segment the market participants into insiders and outsiders (see Kindleberger (1977) and Hong, Scheinkman, and Xiong (2005)). Hong, et al. also assume mean-variance preferences and asset float (increases in publicly tradeable shares) and short-sales constraints. In their setting, bubbles arise because insiders holding I.P.O. shares before a lock-up have superior information over outsiders who are overconfident in their beliefs about the asset's value. Thus, a bubble arises as outsiders bid up the value of the stock and it crashes once the lock-up expires since insiders know the asset is overvalued. While we acknowledge the importance of asymmetric information and heterogeneous agents, our principal insight is that a simple model of econometric learning with homogenous agents can generate escape-like behavior that triggers recurrent bubbles and crashes. In the present setting, asset float are iid shocks that can occasionally precipitate large swings in agents' econometric estimates of risk and expected returns.

Our paper is also closely related to a recent paper on learning and asset pricing by Adam, Marcet, and Nicolini (2006). They adopt a consumption-based asset pricing model and replace rational expectations with least squares learning. They find that the model does a better job at matching several quantitative features of stock price

time series data. Our paper differs from their model in the channel through which learning generates bubbles and crashes. In Adam et al. the channel is through the non-linearity in which beliefs enter the law of motion. Our paper emphasizes the escape-like behavior that can arise in a model where agents estimate risk in addition to expected return.

This paper proceeds as follows. Section 2 presents the model. Section 3 states the basic stability results and Section 4 studies global dynamics. Section 5 presents the numerical results illustrating the recurrent bubbles and crashes. Finally, Section 6 concludes.

2   Model

We employ a simple linear asset pricing model. There is one risky asset that yields a dividend stream $\{y_t\}$ and trades at the price p_t, net of dividends. There is also a risk-free asset that pays the rate of return $R = \beta^{-1}$, where β is the discount factor. In this environment, the risky asset price is assumed to follow

$$p_t = \beta \hat{E}_t \left(p_{t+1} + y_{t+1} \right) - \beta a \sigma^2 z_{st} \tag{1}$$

where $\hat{E}$ are (possibly) non-rational expectations. For $a = 0$, (1) can be derived from the Lucas asset pricing model with risk-neutrality. The Lucas model is an endowment economy in which consumers choose sequences of consumption, equity and bond holdings, to maximize the expected present value of lifetime utility. Provided agents are risk-neutral and financial markets are complete $\beta R = 1$, where β is the discount rate.[3]

The second term in (1) captures two key features to our analysis: the outside supply of shares of the risky asset follows a stochastic process z_{st}; the presence of risk-averse agents ($a > 0$) implies that asset price also depends on agents' perceptions of the conditional variance of excess returns $\sigma^2 = Var_t \left(p_{t+1} + y_{t+1} - Rp_t \right)$. A stochastic process for share supply is meant to proxy for asset float and I.P.O. lock-up expiration. In the presence of short sales constraints variations in the outside share supply can affect stock price, an issue of increasing empirical importance in the financial economics literature (see Ofek and Richardson (2003), Cochrane (2005), and Branch and Evans (2006)). Here we motivate the presence of asset share supply by appealing to this literature and emphasizing that incompleteness in markets can give rise to an important role to supply variation in asset pricing. However, (1) makes it clear that with risk-neutral agents share supply does not have an effect on price. This is in line with DeLong, Schleifer, Summers, and Waldmann (1990) who note that

---

[3]Of course, below we motivate the model as not being a complete markets model and so we might expect $\beta R \neq 1$. Our analysis does not hinge on this restriction.

risk-averse traders may not take aggressive short positions in a risky-asset, thereby preventing full arbitrage of profitable trading opportunities.

When $a > 0$, equation (1) can be derived from an overlapping generations version of the consumption CAPM model with mean-variance preferences (Bohm and Chiarella (2005)). This assumption gives rise to a mean-variance maximizing setting in which agents optimize their portfolio by maximizing risk-adjusted expected wealth. Mean-variance preferences are a frequently employed approach to tractably modeling limited risk tolerance (downward sloping asset demand). See, for example, DeLong, et al (1990) and Lewellen and Shanken (2002). The novelty of our approach is that we assume agents estimate the value of this risk. Risk-aversion implies that agents' welfare declines with the variance of returns, σ^2. It is agents' concern with risk that makes σ^2 an equilibrium object of the model and this is a key ingredient to our finding of recurrent bubbles and crashes. In the learning section below, time-varying estimates of σ^2 will arrest explosive bubbles and can lead to crashes.

We assume the exogenous process for dividends is as follows

$$y_t - y_0 = \rho \left(y_{t-1} - y_0 \right) + u_t$$

We assume that share supply follows a multiplicative process of the form

$$z_{st} = \{\min(s_0, \Phi p_t)\} \cdot V_t$$

where u_t, V_t are uncorrelated white noise shocks, $EV_t = 1$, $y_0, s_0 > 0$, $0 \le \rho < 1$ and we set

$$\Phi = \frac{10 s_0}{\bar{p}}$$

where $\bar{p}$ is the mean stock price in a fundamentals based REE. In this formulation, share supply is exogenous except at very low prices (i.e. when price falls to 10% of its fundamentals value). The endogeneity of share supply at low prices is meant to capture that asset float will dry up in financial markets that perform poorly. In the real-time learning simulations, the endogeneity of share supply will ensure that price remains non-negative. In the analysis below, we assume for simplicity that $\rho = 0$. This has the advantage that all of the asset price dynamics are reflective of the learning process.

It is well-known that in asset pricing models of this form there are (broadly) two classes of rational expectations solutions: the "fundamentals" solution and a "bubbles" class of solutions. A rational expectations equilibrium (REE) is a stochastic process $\{p_t\}$ that solves (1) with $\hat{E} = E$. The fundamentals-based REE can be found by iterating (1) forward

$$p_t = \sum_{j=1}^{\infty} \beta^j E_t y_{t+j} - \sum_{j=0}^{\infty} \beta^j a \sigma^2 E_t z_{st+j}$$

There is additionally a class of bubbles REE of the form

$$p_t = \sum_{j=1}^{\infty} \beta^j E_t y_{t+j} - \sum_{j=0}^{\infty} \beta^j a\sigma^2 E_t z_{st+j} + \beta^{-t}\eta_t$$

where η_t is a martingale, i.e. $E_t\eta_{t+1} = \eta_t$. Notice that for $0 < \beta < 1$ the bubbles REE is explosive. To generate empirically meaningful time-series it is often assumed that η_t follows a Markov process that periodically collapses the bubble (Blanchard (1979), Blanchard and Watson (1982), Evans (1991)).

There is a wide literature that catalogs theoretical objections to bubbles. For instance, explosive bubbles may violate the consumers' transversality condition unless η_t is constructed in a very specific manner. Under complete financial markets, an explosive bubble that does not violate the transversality constraint would induce short-sellers to crash the bubble, making the presence of bubbles theoretically impossible. Additionally, Diba and Grossman (1988) show that, since free disposal implies price can never be negative, a collapsing bubble implies price goes to zero and from there a bubble can never again arise.

Our aim in this paper is to provide a model that yields the periodic bursts and collapses of bubbles as was the goal in Blanchard and Watson (1982). However, we avoid some of these theoretical issues by taking a boundedly-rational viewpoint. We assume that agents attempt to learn, in real-time, about the underlying stochastic process followed by the stock price, in particular about the reduced-form parameters of the model and about the conditional variance of the rate of return. Because the model is self-referential it is conceivable that agents' learning will produce, as endogenous reactions to the intrinsic fundamental shocks, periodic bubbles and crashes.

3 Stability under Learning

In this section we turn to an examination of the stability of the fundamentals and bubbles REE under adaptive learning. First, we follow the section above and take σ^2 as given and study the stability under learning of the parameters in the agents' forecasting model. We then show how σ^2 is pinned down in equilibrium and specify a recursive algorithm for estimating the conditional variance in real-time. Finally, we study the stability properties of the REE with endogenous σ^2.

3.1 Expectational Stability

As explained above we now set $\rho = 0$. We begin by noting that the bubbles solution has an alternative representation of the form,

$$p_t = -y_0 + \beta^{-1}p_{t-1} + a\sigma^2 z_{st-1} + \xi_t$$

where ξ_t is an arbitrary martingale difference sequence, i.e. $E_t\xi_{t+1} = 0$.

To address expectational stability we follow Evans and Honkapohja (2001) and study a perceived law of motion (PLM) that allows for both the fundamentals and bubbles REE:

$$p_t = k + cp_{t-1} + \varepsilon_t. \tag{2}$$

With this perceived law of motion, subjective conditional expectations are[4]

$$E_t p_{t+1} = k(1 + c) + c^2 p_{t-1}$$

To ensure stock prices remain non-negative, we also impose that $k(1 + c) \geq -y_0$. Letting $V_t = 1 + v_t$, where v_t is an iid mean-zero disturbance, and plugging these beliefs into (1) yields the actual law(s) of motion (ALM)

$$p_t = \beta(y_0 + k(1 + c) - a\sigma^2 s_0) + \beta c^2 p_{t-1} - \beta a\sigma^2 v_t, \text{ if } s_0 \leq \Phi p_t \tag{3}$$

$$= \frac{\beta(k(1 + c) + y_0)}{1 + \beta a\sigma^2 \Phi(1 + v_t)} + \frac{\beta c^2}{1 + \beta a\sigma^2 \Phi(1 + v_t)} p_{t-1}, \text{ if } s_0 > \Phi p_t \tag{4}$$

Notice that if beliefs are sufficiently close to a rational expectations equilibrium, asset share supply will be exogenous and the actual law of motion can be re-written in terms of a T-map:

$$p_t = T(k, c)(1, p_{t-1})' - \beta a\sigma^2 v_t$$

The T-map

$$T(k, c) = \left(y_0 + k(1 + c) - a\sigma^2 s_0, \beta c^2\right)$$

defines a map from the PLM to the ALM. A fixed point to the T-map is an REE. Notice that there are two fixed points $(\beta(y_0 - a\sigma^2 s_0)/(1 - \beta), 0)$ and $(a\sigma^2 s_0 - y_0, \beta^{-1})$. Therefore, the fixed points of the T-map correspond to the fundamentals and bubbles REE. When analyzing the global dynamics below we consider the case where beliefs are sufficiently far away from the REE so that supply is endogenous.

We follow Evans and Honkapohja (2001) and examine the stability of the fundamentals and bubbles REE under a natural learning rule. The E-stability principle states that locally stable rest points of the ordinary differential equation

$$\frac{d(k, c)'}{d\tau} = (T(k, c) - (k, c))'$$

will be obtainable under reasonable learning algorithms such as least squares. The conditions for E-stability are that the eigenvalues of the Jacobian of $(T(k, c) - (k, c))$ have negative real parts. It has been established by Evans and Honkapohja (2001) in a closely related model that

---

[4]For convenience we adopt the timing assumption that no contemporaneous variables, including z_{st}, are observable at t. The instability of the bubbles solutions under learning does not hinge on this assumption.

1. The fundamentals REE $(\beta(y_0 - a\sigma^2 s_0)/(1-\beta), 0)$ is E-stable provided $0 < \beta < 1$.

2. The bubbles REE $(a\sigma^2 s_0 - y_0, \beta^{-1})$ is not E-stable.

That the bubbles REE is not E-stable has been another cited objection to rational bubbles. Since a slight deviation from the bubbles path would lead to instability, observing such equilibria seems unlikely. In particular, for initial conditions for c that are greater than β^{-1} many learning rules will imply that $p_t \to \infty$. Below we show that with learning about σ^2 it is possible to trigger "escape" like paths from neighborhoods of the fundamentals REE: interestingly, these escape dynamics do not venture off toward the bubbles REE but instead to a time-series well approximated by a random walk.

3.2   Stability of the REE with Risk

Above, we argued that σ^2 is actually an equilibrium object. Taking this into account we now define an REE as a fixed point (k, c, σ^2) of the T-map for k, c and the corresponding map for perceived conditional variance σ^2. Recalling that $\sigma^2 = Var_t(p_{t+1} + y_{t+1} - \beta^{-1}p_t)$, it follows that

$$\sigma^2 = E_t \left(p_{t+1} - E_t p_{t+1} + y_{t+1} - E_t y_{t+1} \right)^2$$

In the case of the fundamentals REE,

$$\sigma^2 = E_t \left(-a\beta\sigma^2 v_{t+1} + u_{t+1} \right)^2 \tag{5}$$

This is analogous to the T-map for the other parameters: given a perceived value of σ^2, (5) gives the implied actual value of σ^2. Thus, a fixed point delivers the REE value for the fundamentals REE:

$$\sigma^2 = \frac{1 \pm \sqrt{1 - 4a^2\beta^2\sigma_v^2\sigma_u^2}}{2a^2\beta^2\sigma_v^2}$$

Below, we will see that the smaller root will be stable under a recursive learning algorithm. For the bubbles REE straightforward calculations show that $\sigma^2 = \sigma_u^2 + \sigma_\xi^2$.

We turn now to a specification of the learning algorithm. Agents are assumed to use recursive least squares to form parameter estimates of k, c, and to use a similar stochastic recursive algorithm, given below, to estimate σ^2. Define $\theta_t = (k_t, c_t)'$ to be the time t estimates of (k, c) and let σ_t^2 be the time t estimate of σ^2 using data through time t. Assuming that at time t agents use parameters estimated using data through time $t-1$ and they condition on variables dated $t-1$ or earlier, real-time expectations are given by

$$E_t p_{t+1} = k_{t-1}(1 + c_{t-1}) + c_{t-1}^2 p_{t-1}.$$

Under learning the price process follows (3)-(4) with k, c, σ^2 replaced with $k_{t-1}, c_{t-1}, \sigma^2_{t-1}$. Let $X_t = (1, p_t)'$. The real-time learning algorithm is thus assumed to be

$$\theta_t = \theta_{t-1} + \gamma_{1,t} R_t^{-1} X_{t-1} \left(p_t - \theta'_{t-1} X_{t-1} \right) \tag{6}$$

$$R_t = R_{t-1} + \gamma_{1,t} \left(X_{t-1} X'_{t-1} - R_{t-1} \right) \tag{7}$$

$$\sigma_t^2 = \sigma^2_{t-1} + \gamma_{2,t} \left(\left(p_t - \theta'_{t-1} X_{t-1} + u_t \right)^2 - \sigma^2_{t-1} \right) \tag{8}$$

The first two equations in (6)-(8) are recursive least squares. For the stability analysis we will assume the gains $\gamma_{1,t} = \gamma_{2,t} = t^{-1}$. Below in the numerical simulations we will instead assume constant gains and allow $\gamma_{1,t} = \gamma_1 \neq \gamma_2 = \gamma_{2,t}$. With a constant gain, referred to as constant-gain learning, the recursive algorithm becomes a form of discounted least squares.

The issue of stability of the REE is whether θ_t, σ_t^2 converge to the fundamentals and/or the bubbles REE. We have already seen that with fixed σ^2 the fundamentals REE will be stable. Evans and Honkapohja (2001) provide conditions that ensure convergence of the recursive system above. These conditions draw on convergence theorems for stochastic recursive algorithms (SRA). For now, we assume initial beliefs lie in the region in which share supply is exogenous. Below we illustrate how the weak convergence results are impacted by (possibly) endogenous supply. To draw on these results requires some re-defining of variables. Let $S_t = R_{t-1}$, $\gamma_{1,t} = \gamma_{2,t} = t^{-1}$ and define $z_t = \left(p_t - \theta'_{t-1} X_{t-1} + u_t \right) = \left(T(\theta_{t-1}) - \theta_{t-1} \right) X_{t-1} - a\beta\sigma^2_{t-1} v_t + u_t$. Then (6)-(8), for the case of exogenous supply, can be re-written as

$$\theta_t = \theta_{t-1} + t^{-1} S_{t-1}^{-1} X_{t-1} \left(X'_{t-1} (T(\theta_{t-1}) - \theta_{t-1})' - a\beta\sigma^2_{t-1} v_t \right) \tag{9}$$

$$S_t = S_{t-1} + t^{-1} \left(\frac{t}{t+1} \left(X_{t-1} X'_{t-1} - S_{t-1} \right) \right) \tag{10}$$

$$\sigma_t^2 = \sigma^2_{t-1} + t^{-1} \left(z_t z'_t - \sigma^2_{t-1} \right) \tag{11}$$

An important difference from the standard parameter learning set-up is that in the first line the stochastic element is the product of the estimated σ^2_{t-1} and the iid shock v_t. For the case of decreasing gain, the "mean dynamics" will, as Cho, Williams, and Sargent (2002) demonstrated, be independent of the iid term. In the numerical simulations below, however, this will provide an important interaction between the learning of θ_t and σ_t^2 that will make it possible for there to be recurrent bubbles and crashes. Defining $\phi_t = (\theta_t, vec(S_t), \sigma_t^2)'$, and then using the framework of Evans and Honkapohja (2001), it is straightforward to verify that the ODE (ordinary differential equation) associated with the asymptotic behavior of the SRA is defined as

$$\frac{d\phi}{d\tau} = h(\phi) \tag{12}$$

and that local stability of this differential equation governs the local stability of the

12

REE under (9)-(11).[5] where the components of h are

$$
\begin{align}
h_\theta &= S^{-1}M(\theta, \sigma^2)(T(\theta) - \theta)' \tag{13}\\
h_S &= M(\theta, \sigma^2) - S \tag{14}\\
h_{\sigma^2} &= \left((T(\theta) - \theta)M(\theta, \sigma^2)(T(\theta) - \theta)' + \sigma_u^2 + (a\beta\sigma^2)^2\sigma_v^2 - \sigma^2\right), \tag{15}
\end{align}
$$

and where $M = EX_{t-1}(\theta, \sigma^2)X_{t-1}(\theta, \sigma^2)'$. That is, the SRA will locally converge (with probability approaching one) to a stable rest point of the ODE. The Jacobian of this ODE, evaluated at the REE, provides the relevant stability conditions:

$$
\begin{pmatrix}
\beta(1+c) - 1 & \beta k & 0 & 0 & 0 & 0 & 0 \\
0 & 2\beta c - 1 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & -1 & 0 & 0 & 0 & 0 \\
\frac{\partial M(1,2)}{\partial k} & \frac{\partial M(1,2)}{\partial c} & 0 & -1 & 0 & 0 & 0 \\
\frac{\partial M(1,2)}{\partial k} & \frac{\partial M(1,2)}{\partial c} & 0 & 0 & -1 & 0 & 0 \\
\frac{\partial M(2,2)}{\partial k} & \frac{\partial M(2,2)}{\partial c} & 0 & 0 & 0 & -1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 2a\beta\sigma_v^2\sigma^2 - 1
\end{pmatrix}
$$

The Jacobian has eigenvalues $-1 + 2c\beta, -1 + \beta + c\beta, -1 + 2a\beta\sigma_v^2\sigma^2$, and repeated values of -1.

We have the following result.

Proposition 1 *For the model (1) with the adaptive learning algorithm (6)-(8) and gains $\gamma_{1,t} = \gamma_{2,t} = t^{-1}$:*

1. *The fundamentals REE is locally stable under learning if $0 < \beta < 1$ and $0 < \sigma^2 < \frac{1}{2a\beta\sigma_v^2}$.*

2. *The bubbles REE is unstable under learning.*

There are various interpretations in this setting for the phrase "locally stable under learning," as discussed at length in Evans and Honkapohja (1998). For example, Marcet and Sargent (1989) point out that probability one convergence obtains provided the stochastic recursive algorithm is augmented with a "projection facility" that restricts parameter estimates to a suitable compact set around the equilibrium of interest. The use of projection facilities has been criticized by Grandmont (1998) and clearly its use rules out some potentially interesting global dynamics.[6] As we will

---

[5]There is also a connection between the global behavior of the O.D.E. and of the learning algorithm. See the discussion below.

[6]Without the use of a projection facility it is possible to establish local convergence with probability close to one for gain sequences with sufficiently slow adaption rates. In contrast, unstable under learning implies convergence with probability zero.

now see, with constant-gain learning, bubble-like global dynamics periodically emerge as temporary escapes from the fundamentals REE. Furthermore the increases in perceived risk along these bubble paths eventually acts like an endogenous projection facility to return the price process to a neighborhood of the fundamentals REE.

We now turn to the analysis of the global learning dynamics in our model.

4   Global Properties

The results above demonstrate that the fundamentals REE is locally stable under learning, while the bubbles REE are not. This result obtains both with and without learning of the conditional variance of stock returns. Thus, the onset of recurring bubbles and crashes will arise from the global dynamic properties of the model under learning.

Figure 3 illustrates the central intuition of the local and global stability results. Figure 3 plots the $T_c = \beta c^2$ component of $T(\theta)$ holding k and σ^2 fixed. There are clearly two REE, the fundamentals at $c = 0$ and the bubbles at $1/\beta$. The arrows in the figure show the direction of adaptation under the ODE h_θ. For initial values $c > 1/\beta$ the ensuing estimated values of c will explode without limit. For initial $c < 1/\beta$ there is convergence to the fundamentals REE. Below we consider a setting where agents estimate θ with a constant-gain least squares algorithm.

Although all trajectories originating in the basin $[0, 1/\beta)$ will eventually settle down at the fundamentals REE, the global dynamics along a convergent path could still be interesting. This point was made forcefully by Sargent (1999), Cho, Williams, and Sargent (2002), and Williams (2004). In particular, at $c = 0$, in a fundamentals REE, the stock price is equal to a constant plus noise. Away from this REE, even as the dynamics are on a convergent path, estimates for $c_t > 0$ introduce serial correlation into the model. This serial correlation may be self-reinforcing leading agents to (temporarily) believe price is serially correlated.

These insights can be seen by studying the mean dynamics of the recursive learning algorithms. We first present results for the case of exogenous share supply and leave the details of the case where share supply may become endogenous to an Appendix.

We now replace t^{-1} updating with a constant gain so that agents' beliefs are updated according to
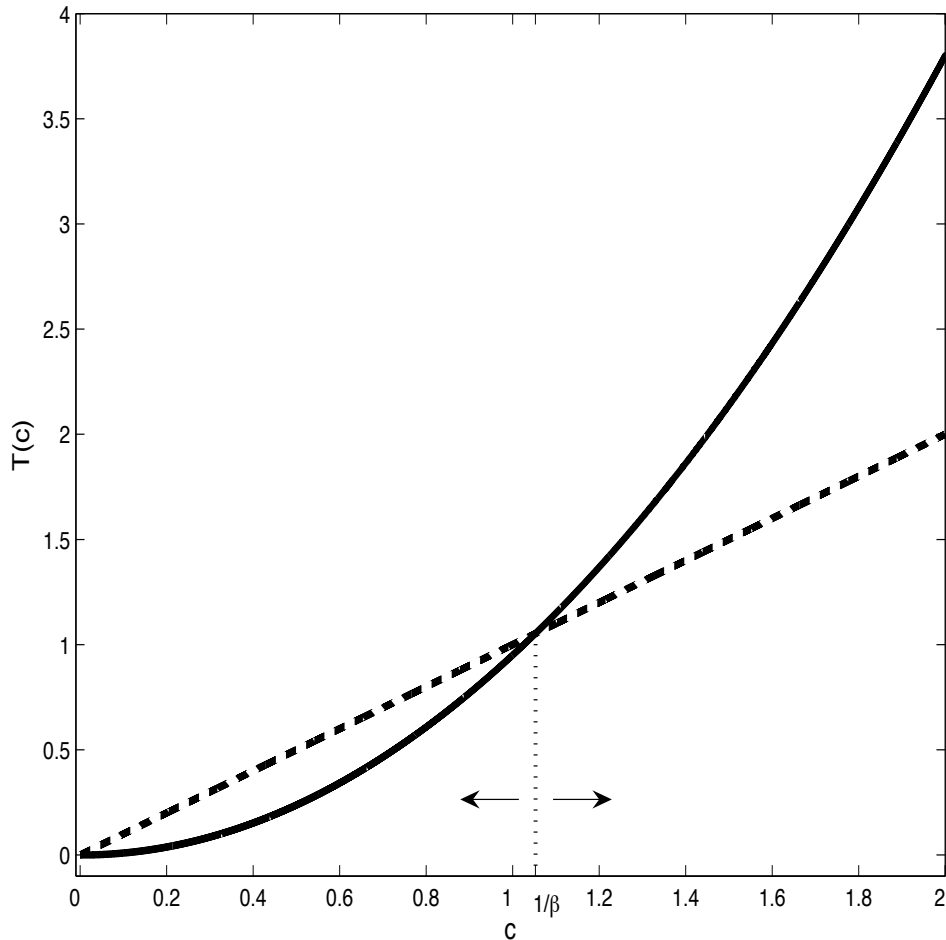
$$\theta_t = \theta_{t-1} + \gamma_1 S_{t-1}^{-1} X_{t-1} \left(X_{t-1}'(T(\theta_{t-1}) - \theta_{t-1})' - a\beta\sigma_{t-1}^2 v_t \right) \tag{16}$$

$$S_t = S_{t-1} + \gamma_1 \left(\left(X_{t-1} X_{t-1}' - S_{t-1} \right) \right) \tag{17}$$

$$\sigma_t^2 = \sigma_{t-1}^2 + \gamma_2 \left(z_t z_t' - \sigma_{t-1}^2 \right) \tag{18}$$

where γ_1, γ_2 are the constant gains and here we allow $\gamma_1 \neq \gamma_2$. The main motivation for believing that private agents would use constant gains is that they would be concerned

Figure 3: T-map (T_c component).

with the possibility of structural change. Constant gain least squares places greater weight on current relative to older data, and so is better able to track the form of the stochastic process generating the data when there is structural change taking an unknown form. Implications of using constant-gain learning have been explored in the literature cited in the Introduction. Empirical support for constant-gain learning can be found, for example, in Orphanides and Williams (2005a), Branch and Evans (2006), Sargent, Williams, and Zha (2007), and Milani (2007).

The idea behind mean dynamics are that solutions to the ODE in (12) provide a good approximation to the expected dynamics as $\gamma_1, \gamma_2 \to 0$. In Proposition 1 we provided conditions under which the REE is locally stable or unstable. That (12) governs local stability was established in Evans and Honkapohja (2001). However, the mean dynamics also can shed light on the global stability of the REE as well, in the sense that for small gains the sample paths from the SRA weakly converge to the

path defined by the ODE.

Below, we further develop the intuition for the origin of bubbles and crashes, by numerically solving the ODE for the mean dynamics. First, we formalize the notion that the mean dynamics approximate the constant-gain learning dynamics for small gain parameters. We proceed by first noting that it is possible to re-write the constant-gain learning algorithms (16)-(18) in the form[7]

$$\phi_t^\gamma = \phi_{t-1}^\gamma + \gamma \mathcal{H}(\phi_{t-1}^\gamma, \bar{X}_t)$$

where $\bar{X}_t' = (X_t, X_{t-1}, u_t)'$, and the components of $\mathcal{H}$ are given by (16)-(18). The parameter estimates ϕ^γ have been re-written to emphasize their dependence on the gain γ. In order to make a comparison between the solutions to the continuous time ODE and the discrete time recursive algorithm, we need to define a corresponding continuous time sequence for ϕ_t^γ, which we denote as $\phi^\gamma(t)$.[8]

When γ is replaced by γ_t, a sequence of gains that decreases over time, Proposition 1 shows that the fundamentals REE is stable. With a constant gain, the learning dynamics do not settle down over time but may converge to a stationary distribution. In general, mean dynamics provide a good approximation to the mean path under constant-gain learning. The following proposition formally establishes this result for the case of exogenous share supply in a neighborhood of the fundamentals REE.

Proposition 2 *Assume $z_{st} = s_0 V_t$. Define, for any ϕ_0 sufficiently close to the fundamentals REE, $\tilde{\phi}(t, \phi_0)$ as the solution to the ODE $d\phi/dt = h(\phi)$, with initial condition ϕ_0, where h is defined in (12)-(15). Consider the random variable, indexed by the constant gain γ, $U^\gamma(t) = \gamma^{-1/2} \left(\phi^\gamma(t) - \tilde{\phi}(t, \phi_0) \right)$. As $\gamma \to 0$, $U^\gamma(t), 0 \leq t \leq T$, converges weakly to a zero mean random variable.*

The proof is contained in an Appendix. We remark that the "neighborhood" of validity of this proposition need not be small and, as shown in the Appendix, can include a wide range of values for ϕ. This result establishes that, over finite periods of time, the constant-gain learning dynamics will converge weakly to the solution of the ODE; these "mean dynamics" then provide a good approximation to the actual real-time learning dynamics with a small constant gain. It is important to emphasize that this convergence result is across sequences of ϕ_t, for alternative gains $\gamma \to 0$, and not along a particular realization.[9] It is worth remarking that in the case where share supply may become endogenous a similar result can be expected to hold since the

---

[7]To capture $\gamma = \gamma_1 \neq \gamma_2$, the σ^2 component must be divided by γ_1, implying the mean dynamics (15) will depend on the relative gain γ_2/γ_1.

[8]It is straightforward to construct a continuous time process for ϕ^γ: denote s as real-time and let $t_s^\gamma = s\gamma$; define $\phi^\gamma(t) = \phi_s^\gamma$ if $t_s^\gamma \leq t < t_{s+1}^\gamma$.

[9]See Williams (2004) and McGough (2006) for further discussion.

weak convergence theorems presented in Evans and Honkapohja (2001) do not depend on the conditional linearity of the price process that is exploited in the proof to this proposition. However, verification of the technical conditions in this case are difficult and beyond the scope of the present paper. Instead, we turn to an approximation and then present numerical results.

4.1   Mean Dynamics

The result in Proposition 2 shows that for sufficiently small constant gain parameters, the mean dynamics provide a good approximation to the constant-gain learning dynamics. Furthermore, Proposition 1 establishes that the fundamentals REE is a rest point of the mean dynamics. One way to think of the mean dynamics is that it gives the mean path followed by the learning dynamics if there is a shock to the parameter estimates. Under constant-gain learning, agents place greater weight on recent observations and may cause agents' beliefs to respond to shocks and trigger these mean dynamics. How responsive agents' beliefs are to these shocks depends on the constant-gain parameters. For sufficiently small gains the economy will, with high probability, remain in a neighborhood of the REE, as indicated by Proposition 2. However, for larger gains interesting global dynamics may arise.

The key to this intuition is that estimates for (k, c) are likely to shift in a systematic manner in response to shocks, since with a constant gain agents place some weight on the possibility that the fundamentals REE is not the true model. The size of the constant gain influences the space of models that agents cannot statistically reject. To illustrate this reasoning Figure 4 plots the 90% and 50% confidence ellipses for the fundamentals REE assuming a constant gain $\gamma_1 = .05$. To compute this figure we follow Evans and Honkapohja (2001, Chp.7) who show that, under constant-gain learning, the parameter estimates are distributed asymptotically normal around the REE value, with the variance increasing in the gain.[10] This figure was generated by assuming the following baseline parameterization: $\beta = 0.95, a = 0.75, \sigma_u^2 = 0.9. \sigma_\nu^2 = 0.5, y_0 = 1.5, s_0 = 1$. These parameter values were chosen so that the fundamentals REE is locally stable according to the condition in Proposition 1. Figure 4 illustrates that the confidence ellipses have a decreasing principal axis, indicating that one can expect many trajectories moving in the direction of this axis. Notice that the ellipses are pointed in the direction of a random walk without drift, with c increasing and k decreasing along the principal axis. The relative size of these ellipses depend on the size of the constant gain. From this figure, we conclude that under constant-gain learning, agents will remain alert to possible misspecification in their forecasting rule – thereby, raising the possibility that they may come to believe price is a random walk.

---

[10]See the Appendix for further details.

Figure 4: Confidence ellipses around fundamentals REE for constant gain learning version of the model.
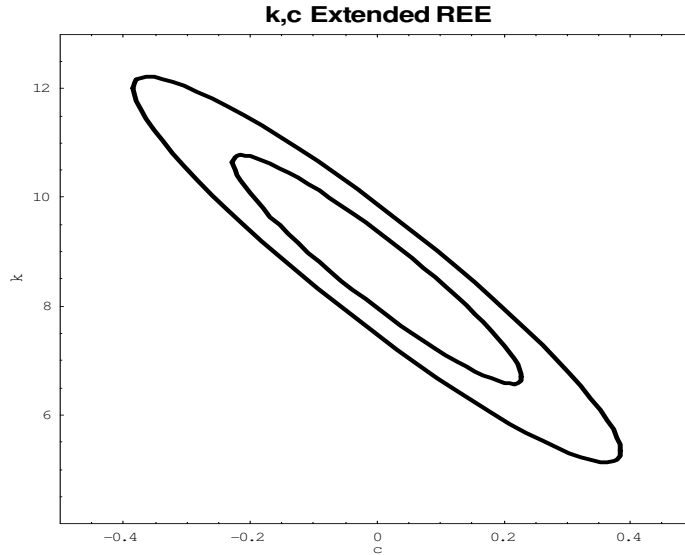
One can think of points on the principal axes of these ellipses as providing the triggers for escapes. In the current setting, the confidence ellipsoid consists of the (k, c) ellipse in Figure 4 and a confidence region for the risk aversion parameter σ^2 that is a small interval around the fundamentals REE value. The numerical examples below illustrate that increases in c (and decreases in k) combined with small changes in risk aversion can trigger escape-like dynamics, and that gain parameters $\gamma_2 > \gamma_1$ make escapes more likely. The trigger for escapes then are some combination of increases in expected return and changing risk estimates.

In this subsection, we numerically solve the mean dynamics to illustrate the global stability properties of the model. These dynamics provide the intuition for the real time dynamics in the next Section. We choose various starting values for the belief parameters θ, σ^2, R, chosen from the 90% confidence ellipsoid, and solve the non-linear

18

differential equation (12). The Appendix contains further details on computing the ODE approximation in the case that share supply is endogenous. We show that the fundamentals REE is a stable rest point of the mean dynamics for all initial values $c \in [0, 1/\beta)$.

The global stability properties of the mean dynamics have interesting implications for the constant-gain learning dynamics (16)-(18). The mean dynamics represent the expected path of beliefs (for small gain) given an initial condition. With a constant-gain recursive algorithm, agents' beliefs react more strongly to recent than past data. It is possible for a large shock to induce agents to revise upwards their beliefs of the slope parameter c, placing the economy along a bubble-like path. The global stability of the fundamentals REE suggests that a large run-up of price will be possible, but will then collapse as the mean dynamics take over.

One novelty to the approach here is the importance of risk in these dynamics. Estimates of σ^2 play two roles: first, a constant-gain algorithm for σ^2 will introduce drift into the model and make 'escapes' more likely; second, along an explosive bubble path estimates of risk will increase, and as σ^2 increases sufficiently the bubble will collapse.

To illustrate these points we turn to a numerical investigation of the mean dynamics. We parameterize the model as above (reproduced here for convenience): $\beta = 0.95, a = 0.75, \sigma_u^2 = 0.9. \sigma_\nu^2 = 0.5, y_0 = 1.5, s_0 = 1$. We randomly draw from values along the ellipse. One can think of constant-gain learning dynamics as re-initializing the mean dynamics.

Figure 5 illustrates the mean dynamics for various initial values of c, k lying in the fundamentals region, and along the confidence ellipse, and for initial σ^2 below its stable REE value. For this figure the gains are chosen so that the gain on parameter updating, i.e. for c, k, is twice the gain on σ^2. Setting σ^2 below and c above their REE values corresponds to a decrease in perceived risk and to a perceived serial correlation in price. The figure plots (moving clockwise and starting in the northwest) the belief parameters k, c, the constant in price expectations $k(1 + c)$, the state covariance matrix estimates R_t (with the (1,2), (2,2) elements shown), and the perceived risk estimate σ^2). Figure 5 demonstrates that the fundamentals REE is a stable rest point for the mean dynamics, as indicated by Proposition 1.

Figure 6 plots the mean dynamics for the same parameter values, with initial values drawn from the confidence ellipse, but now with the relative gains $\gamma_2/\gamma_1 = 2$, so that beliefs are relatively more alert to changes in risk than expected return. Notice again that the fundamentals REE is a stable rest point of the mean dynamics. However, the transition path for the mean dynamics, in particular the behavior of the autoregressive parameter c, are now very interesting. At first the estimates of c move toward the fundamentals REE, but then they reverse course and increase to a value of $c = 1$, where it remains for a short time before converging to $c = 0$. Importantly,

Figure 5: Mean Dynamics, various initial values for slope parameter c, k and perceived risk σ^2, drawn from the confidence ellipsoid. Small gain on perceived risk (γ_2).
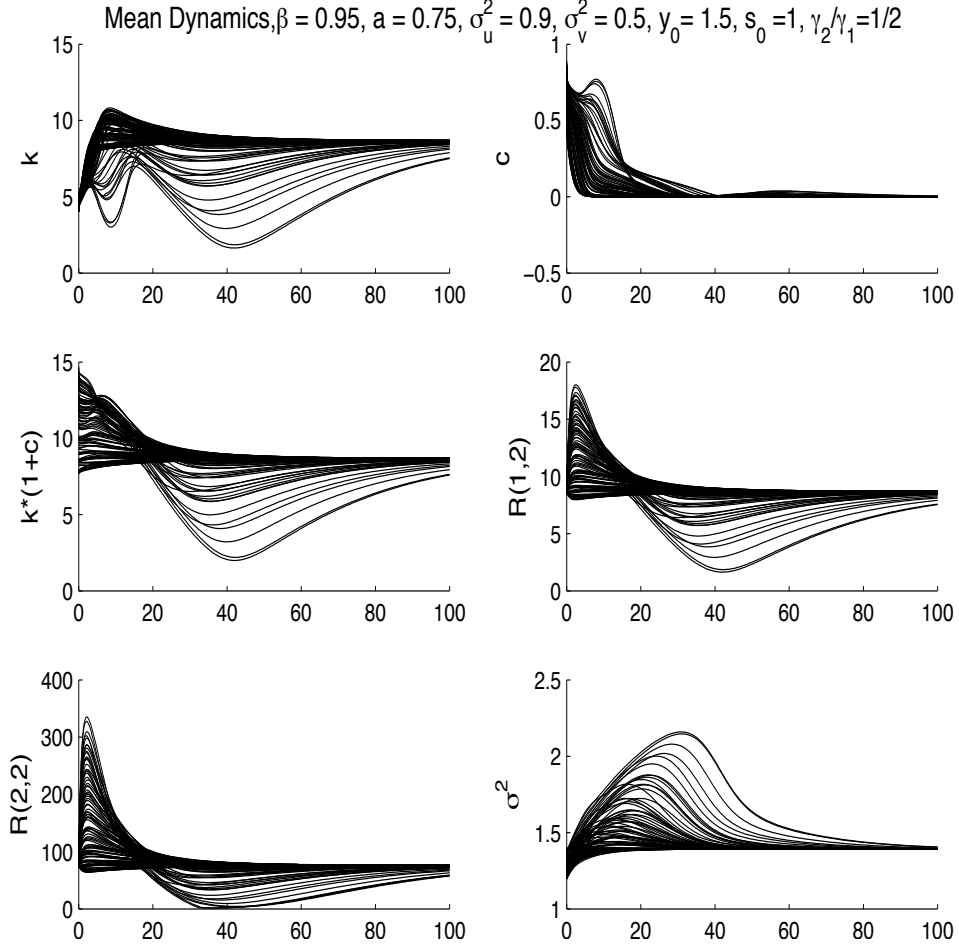
note that $k \to 0$ as $c \to 1$. Thus the the mean dynamics predict that agents will come to believe stock prices follow a random walk. Eventually, the risk estimate will increase, leading agents to again forecast price as a serially uncorrelated process.

It is worth briefly remarking on the manner in which the gain parameters γ_1, γ_2 affect these dynamics. Essentially, the gain affects the relative size of the 90% confidence ellipsoid: larger gains imply greater alertness to possible model misspecification. Through numerical explorations, we found that greater sensitivity in updating estimates of σ^2 were more likely to trigger escape paths in the mean dynamics. For this reason, in the real time dynamics below, we choose values of $\gamma_2 > \gamma_1$.

That agents might forecast price to follow a random walk is intuitive and has implications for the dynamics under constant-gain learning. The mean dynamics in

Figure 6: Mean Dynamics, various initial values for slope parameter c, k and perceived risk σ^2, drawn from the confidence ellipsoid. Large gain on perceived risk (γ_2).

Mean Dynamics,$\beta = 0.95$, a $= 0.75$, $\sigma_u^2 = 0.9$, $\sigma_v^2 = 0.5$, $y_0 = 1.5$, $s_0 = 1$, $\gamma_2/\gamma_1 = 2$
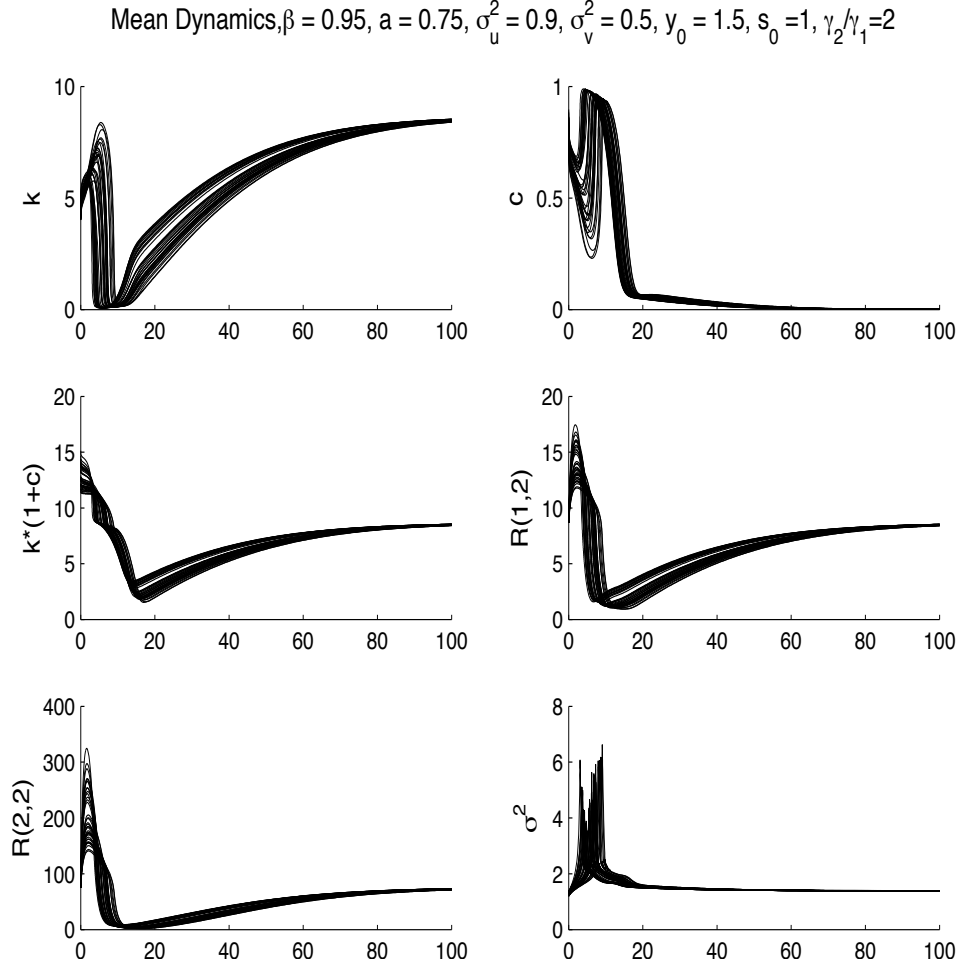
Figure 6 show that if agents detect some serial correlation in the data, and their risk perception changes precipitously, this can lead agents to believe that recent changes in price are permanent shifts and not mean-reverting fluctuations. In a perceived low-risk environment traders will act on these beliefs and further drive up price – inducing a bubble-like path. Eventually, however, agents will revise their risk estimates and crash the bubble.

Figure 6 illustrates that shifts in estimates of σ^2 are important for generating escapes. If the gain γ_2 is sufficiently small (e.g. Figure 5), even as γ_1 increases, mean dynamics do not exhibit random-walk beliefs. Thus, the drift induced by σ^2 is critical for generating bubbles and crashes.

In this model, bubbles and crashes arise under private-sector learning because

agents come to believe that prices follow a random walk: shocks to stock prices are perceived to have permanent effects. These dynamics arise through a complicated interaction between expected-return parameters (k, c) and expected risk perceptions σ^2. Figure 6 shows how it is possible with constant-gain learning for the complicated interaction of these learning features to lead to an escape episode. In the next Section we show that when beliefs escape in this manner, recurrent bubbles and crashes can arise. The second role played by risk is to crash bubbles: along an explosive price path, risk estimates will increase and eventually cause price to collapse. The relative gain γ_2/γ_1 is important for ensuring that bubbles will crash. Larger values of γ_2/γ_1 imply that risk estimates, along a bubbles path, will increase faster than expected return. This insight is illustrated in Figure 7, which plots mean dynamics when k, c are initialized in the explosive bubbles REE region (i.e. $c > 1/\beta$). Notice again that beliefs are temporarily drawn to a random walk before converging to the fundamentals REE.

Figure 7: Mean dynamics, initial k, c drawn from explosive region.
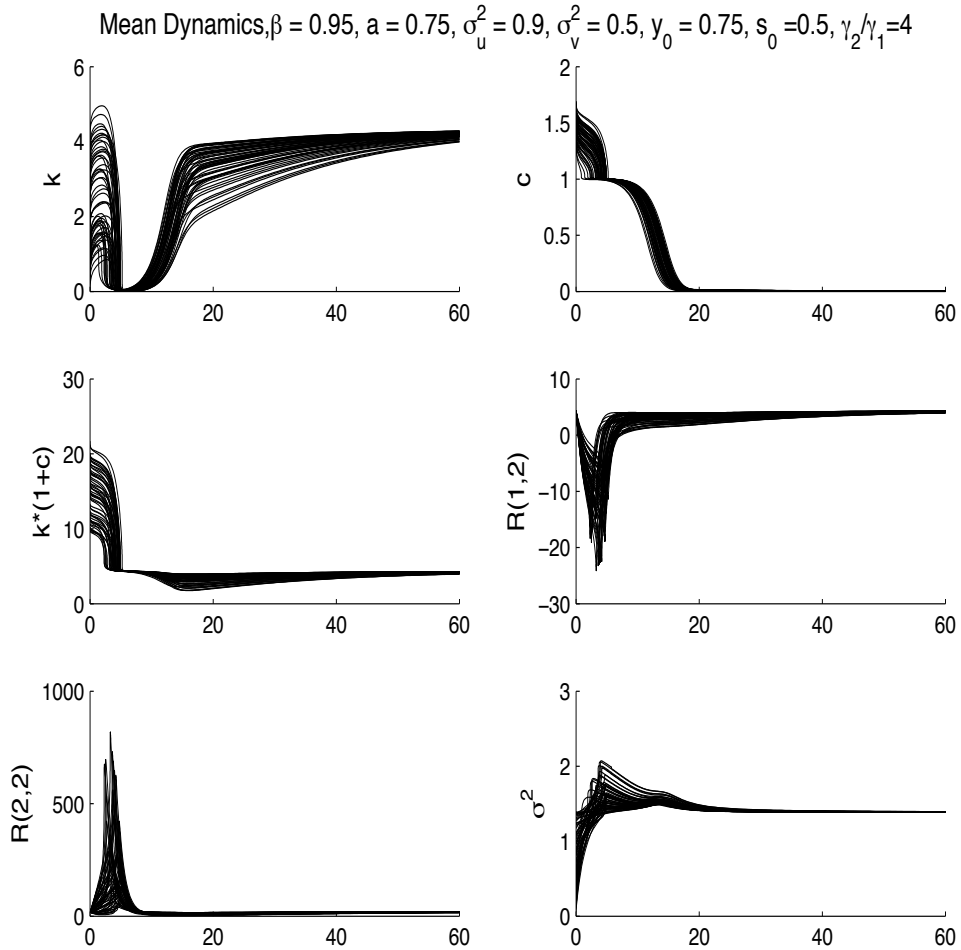
4.2 Discussion

The onset of random walk approximations for stock prices was anticipated in a very different setting by Sargent (1999, Chp. 6). Sargent assumed a simple model in which the state is driven by expectations and a serially uncorrelated shock. Under rational expectations, the economy is white noise deviations from a mean. Sargent demonstrates that under adaptive expectations (in the Cagan sense) the dynamics will be quite different. Adapting an insight from Muth (1960) that adaptive expectations are an optimal predictor for a stochastic process that is MA(1) in first differences, he shows that if agents hold random walk beliefs then the resulting temporary equilibrium stochastic process can be self-fulfilling in the sense that the deviation from rational expectations, of the random walk approximation, may not be detectable. The basic idea is that a random walk model approximates well a model with time-varying means.

The mean dynamics of this Section extend this reasoning to a model with multiple REE and learning about the conditional variance. Sudden decreases in the perceived risk, and adaptive learning via the constant-gain RLS, can lead to self-fulfilling serial correlation that is approximated well by a random walk specification for stock price. The previous subsection shows that random walk beliefs might arise along a path that converges to the REE. This subsection further elaborates on these dynamics, by demonstrating that Sargent's intuition applies in the present setting.

Suppose that agents hold random walk beliefs based on a perceived law of motion of the form

$$p_t = p_{t-1} + \varepsilon_t$$

This arises under the PLM (2) provided $c = 1, k = 0$. Plugging these beliefs into the expectational equation for stock price (1) leads to the actual law of motion (under random walk beliefs):

$$p_t = \beta(y_0 - a\sigma^2 s_0) + \beta L p_t - \beta a \sigma^2 v_t \qquad (19)$$

where L is the lag operator. Following Sargent, one could re-write the PLM and ALM as MA(∞) processes. The PLM can be written[11]

$$p_t = \frac{1}{1 - L}\varepsilon_t \equiv g(L)\varepsilon_t \qquad (20)$$

and the ALM takes the form

$$p_t = \mu + f(L)v_t \qquad (21)$$

where $\mu = (y_0 - a\sigma^2 s_0)\beta/(1 - \beta)$ and $f(L) = \beta a \sigma^2/(1 - \beta L)$.

---

[11]The MA(∞) for the random walk can be derived by taking the limit $c \to 1$.

Sargent (1999) emphasizes two features of misspecified adaptive beliefs that are evident in (20) and (21). First, random walk beliefs introduce serial correlation into a model that is not serially correlated under rational expectations. The moving average processes (20),(21) demonstrate that the perceived serial correlation becomes almost self-fulfilling. Second, random walk beliefs can track constants well. In (20) there is no constant but in (21) there is. Thus, the random walk uses higher-order moments to track low frequency movements (i.e. the mean) in the price process. It is these features that Sargent (1999) mentions as properties of optimally misspecified adaptive beliefs.

It is possible to demonstrate, for the simple asset-pricing model of this paper, that when an escape to random walk beliefs happen, the misspecified forecasting model is optimally tracking these low frequency movements in price. To demonstrate this, we follow Sargent (1999) in plotting the spectral density for the random walk perceived model and the spectral density for the actual law of motion given these random walk beliefs. Figure 8 illustrates the results for the baseline parameterization. The random walk model matches the ALM well except at very low frequencies. This figure is very similar to Figure 7.2 in Sargent (1999), though the model at hand has considerably more serial correlation at all frequencies. Sargent attributes the good match between spectral densities as evidence that a unit root can approximate first moment properties of a stochastic process.

We, therefore, conclude that escapes to a random walk specification for beliefs are likely to happen in a simple asset pricing model with learning about expected return and risk. Moreover, these beliefs are optimally misspecified so that agents will not be likely to detect their misspecification for realistic sample sizes. These findings motivate the next Section which presents real-time simulations of the learning dynamics.

5 Recurrent Bubbles and Crashes: Real-time Learning

The results in the previous section demonstrate that the fundamentals REE is locally stable under real-time learning and that under a suitable recursive algorithm it may be globally stable as well. The mean dynamics results showed that estimates of risk sufficiently away from the fundamentals equilibrium value can induce an escape sending beliefs into a random-walk region. In the mean dynamics, the bubble will eventually collapse, inducing a crash in price, and an eventual return to the fundamental rational expectations equilibrium. Under constant-gain learning we anticipate seeing this process of bubbles arising and then crashing repeatedly. There are two central elements to generating recurrent bubbles and crashes: constant-gain parame-

24

Figure 8: Spectral densities (in logs) for random walk beliefs and the associated actual law of motion.
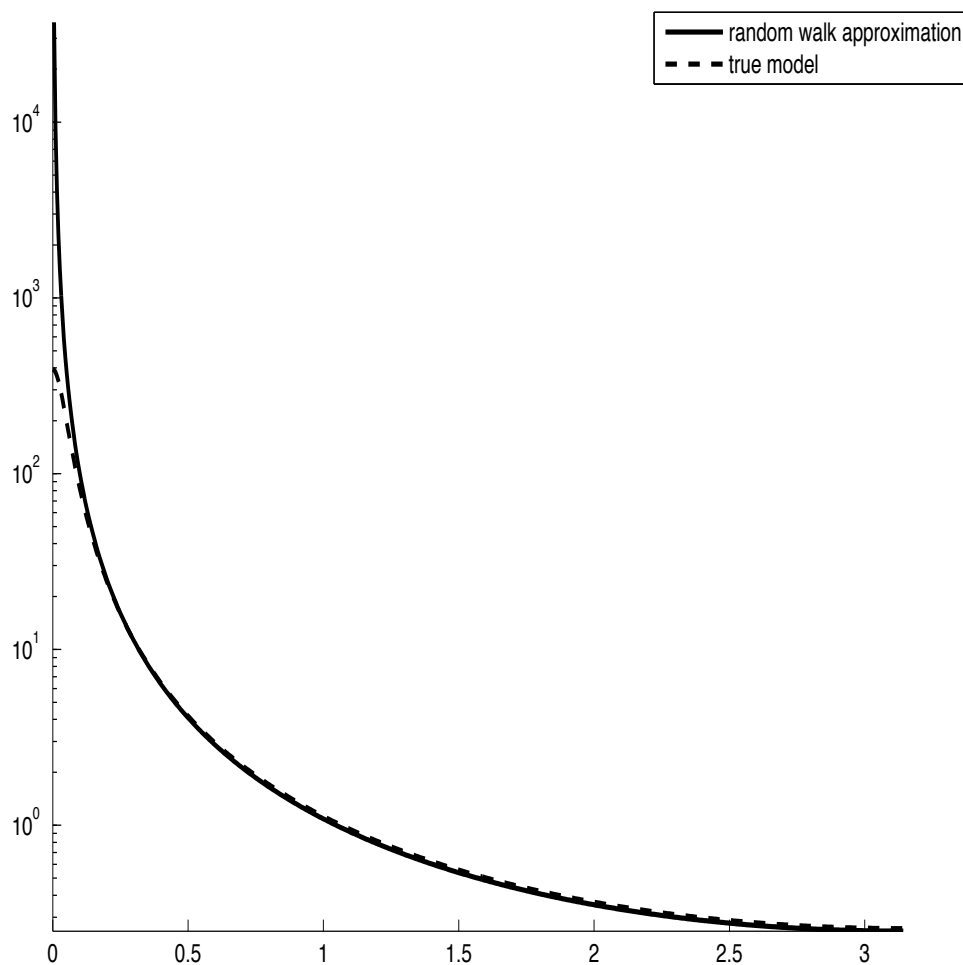
ter updating so that large shocks induce agents to revise upward their beliefs; second, along a bubble path risk estimates will increase sufficiently and eventually crash the bubble. The key player in these channels is the risk estimate σ^2: it generates the drift necessary for an escape; and it acts like an endogenous projection facility to arrest explosive paths. This section turns to a real-time constant-gain learning analysis to illustrate these properties.

Figures 9-13 present the numerical results. We choose the same parameter values as above and consider the effects of different constant gains. The dynamical system,

reproduced here for convenience, is

$$
\begin{aligned}
p_t &= \beta\left(y_0 + k_{t-1}(1 + c_{t-1})\right) + \beta c_{t-1}^2 p_{t-1} - \beta a \sigma_{t-1}^2 z_{st} \\
\theta_t &= \theta_{t-1} + \gamma_1 R_t^{-1} X_{t-1}\left(p_t - \theta_{t-1}' X_{t-1}\right) \\
R_t &= R_{t-1} + \gamma_1\left(X_{t-1} X_{t-1}' - R_{t-1}\right) \\
\sigma_t^2 &= \sigma_{t-1}^2 + \gamma_2\left(\left(p_t - \theta_{t-1}' X_{t-1} + u_t\right)^2 - \sigma_{t-1}^2\right)
\end{aligned}
$$

In each figure we report the results from a typical simulation of length 10,000, which follows a 5,000 length transient period.[12]

Figure 9: Constant gain learning with $\gamma_1 = .01, \gamma_2 = .001$.

$\beta = 0.95$, $a = 0.75$, $\sigma_u = 0.9$, $\sigma_v = 0.5$, $y_0 = 1.5$, $s_0 = 1$, $\gamma_1 = 0.01$, $\gamma_2 = 0.001$
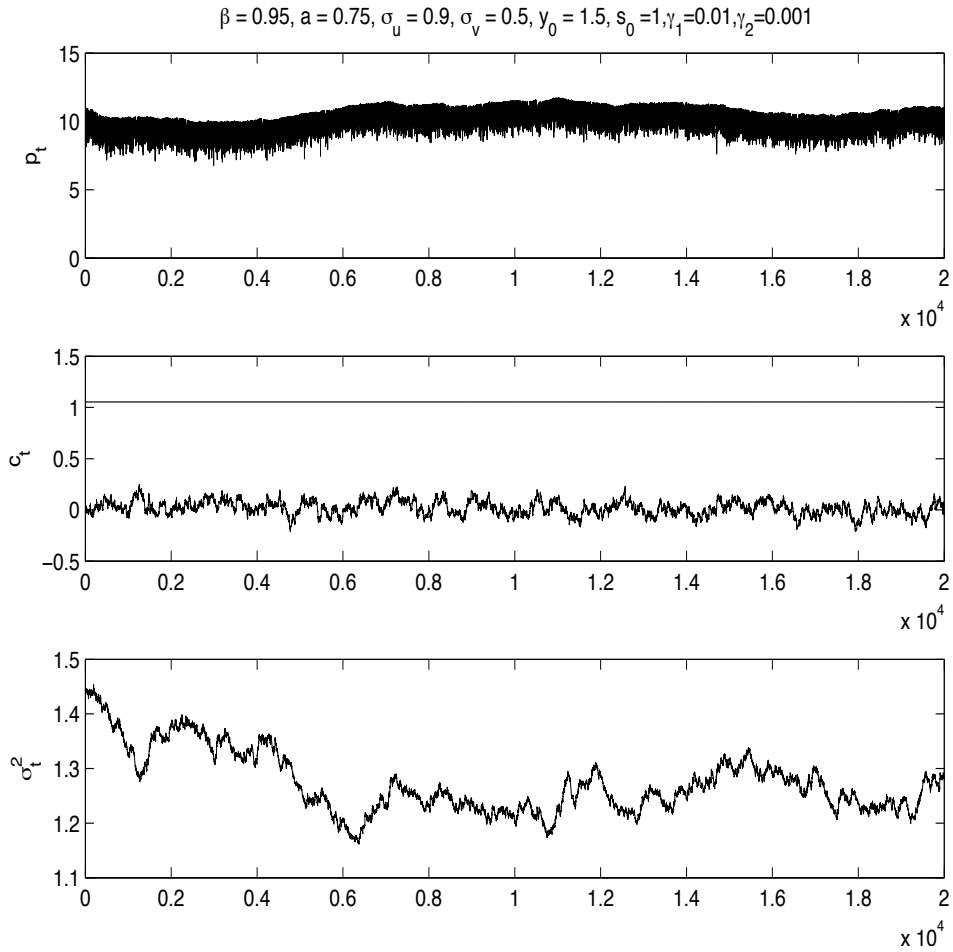
As mentioned above, with constant-gain learning some realizations of the exogenous shocks lead to dynamics that leave the basin of attraction for the fundamentals

---

[12]We also augment the model with a projection facility requiring $k(1 + c) \geq -y_0$. This restriction on beliefs rules out negative stock prices.

Figure 10: Constant gain learning with $\gamma_1 = \gamma_2 = .01$.
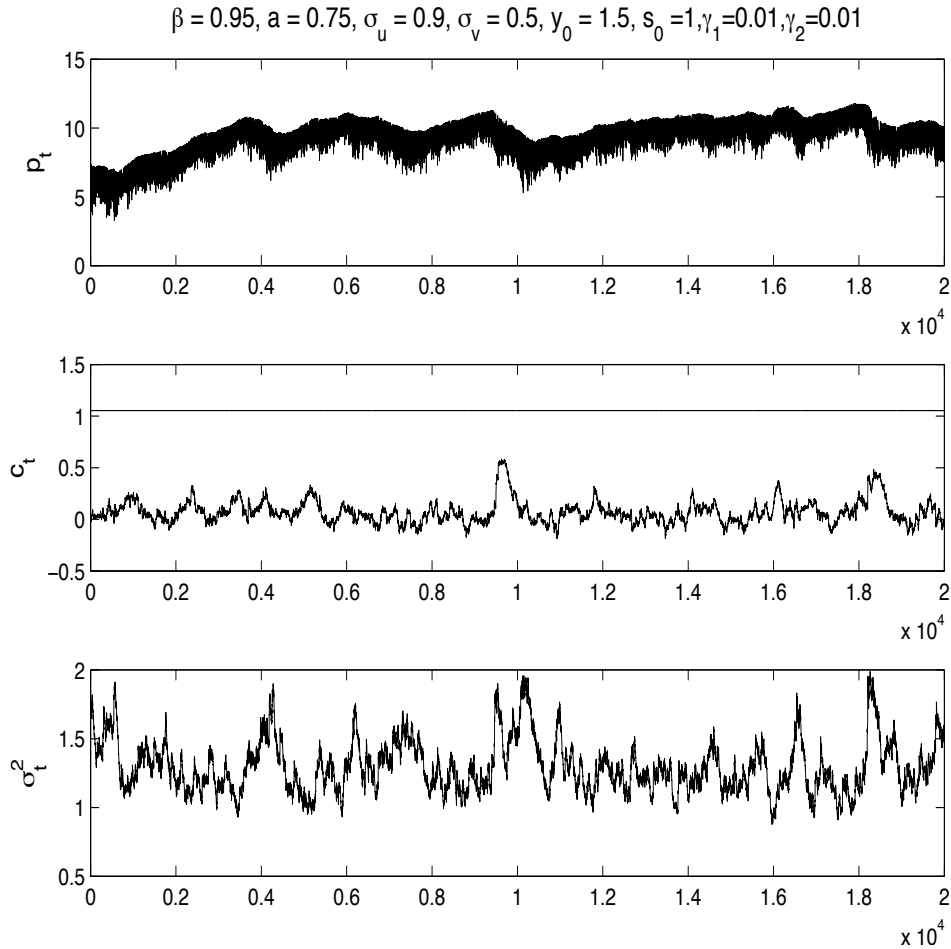
REE and escape to a random walk specification for beliefs. The mean dynamics, though, indicate that eventually the dynamics will return to a neighborhood of the fundamentals solution. For sufficiently small gains we would expect the dynamics to remain close to the fundamentals REE. Figure 9 confirms this intuition by plotting the results from a simulation with $\gamma_1 = .01, \gamma_2 = .001$. The top panel plots the stock price, while the bottom panels plot the estimated slope parameter c_t and risk σ_t^2, respectively. Notice that the belief parameters stay near their fundamentals REE value. As a result price is just a white noise process around its mean value.

Recall from the discussion in Section 4.2 that it is the drifting beliefs that can create self-fulfilling serial correlation. Figure 10 considers the effects on the price dynamics from an increase in the risk gain γ_2. With a much larger gain ($\gamma_2 = .01$) the estimated value of σ_t^2 displays much more volatility, though still not enough in response to the shocks to trigger an escape. The additional drift in beliefs induces

27

Figure 11: Constant gain learning with $\gamma_1 = .01, \gamma_2 = .02$.
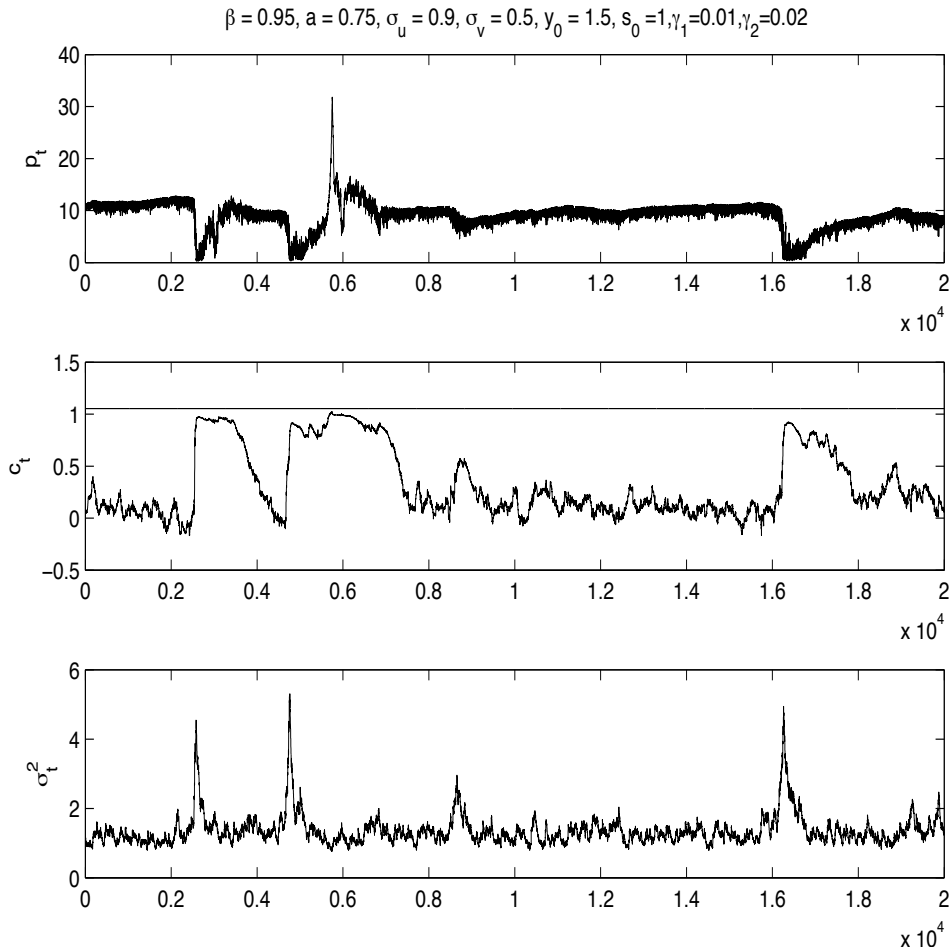
some serial correlation in the price process, but the model still does not exhibit bubbles or crashes.

As mentioned above σ^2 plays an essential role in the existence of bubbles and crashes. In Figure 10, there is not enough drift in the model – since the only shocks to p_t are the supply shocks v_t – to induce beliefs into the bubble region. With more volatility in the estimated risk parameter escapes become possible. Figure 11 begins to show the onset of bubbles and crashes. Figure 11 keeps the same gain for $\gamma_1 = .01$ and considers a slightly greater value for $\gamma_2 = .02$. Figure 11 demonstrates that for the first two thousand periods the dynamics look as they did in the previous figures. But then beginning about period 2200 there is a sudden qualitative change in the dynamics with three crashes and a bubble. Between the crashes and bubbles, the dynamics converge back to a neighborhood of the fundamentals REE. Notice how the beliefs for c, σ^2 track the mean dynamics.

Figures 9-11 suggest that it is joint learning about expected returns and σ^2 that is critical for bubbles and crashes. The remaining figures display simulation results for the gains $\gamma_1 = 0.01, \gamma_2 = .04$. These gain and parameter values were chosen for pedagogical reasons. Figure 12 demonstrates the remarkably different dynamics with these parameter values. For the first 1500 periods or so the dynamics look as they did in the first three figures, but then around period 1600 there is a dramatic qualitative and quantitative change in the nature of the price and belief dynamics. At about period 1500 there is a dramatic jump in both the RLS estimate c_t and the conditional variance estimate σ_t^2. Notice, in particular, as predicted by the mean dynamics, that beliefs hover near a random walk. This induces a sudden crash in the stock price, which is then followed by a series of bubbles and crashes in the sense of sustained deviations from the fundamentals price. Notice also that the price dynamics follow a path somewhat reminiscent of the detrended log S&P 500 index shown in Figure 2.

Figure 12: Constant gain learning with $\gamma_1 = .01, \gamma_2 = .04$.

$\beta = 0.95, a = 0.75, \sigma_u = 0.9, \sigma_v = 0.5, y_0 = 1.5, s_0 = 1, \gamma_1 = 0.01, \gamma_2 = 0.04$
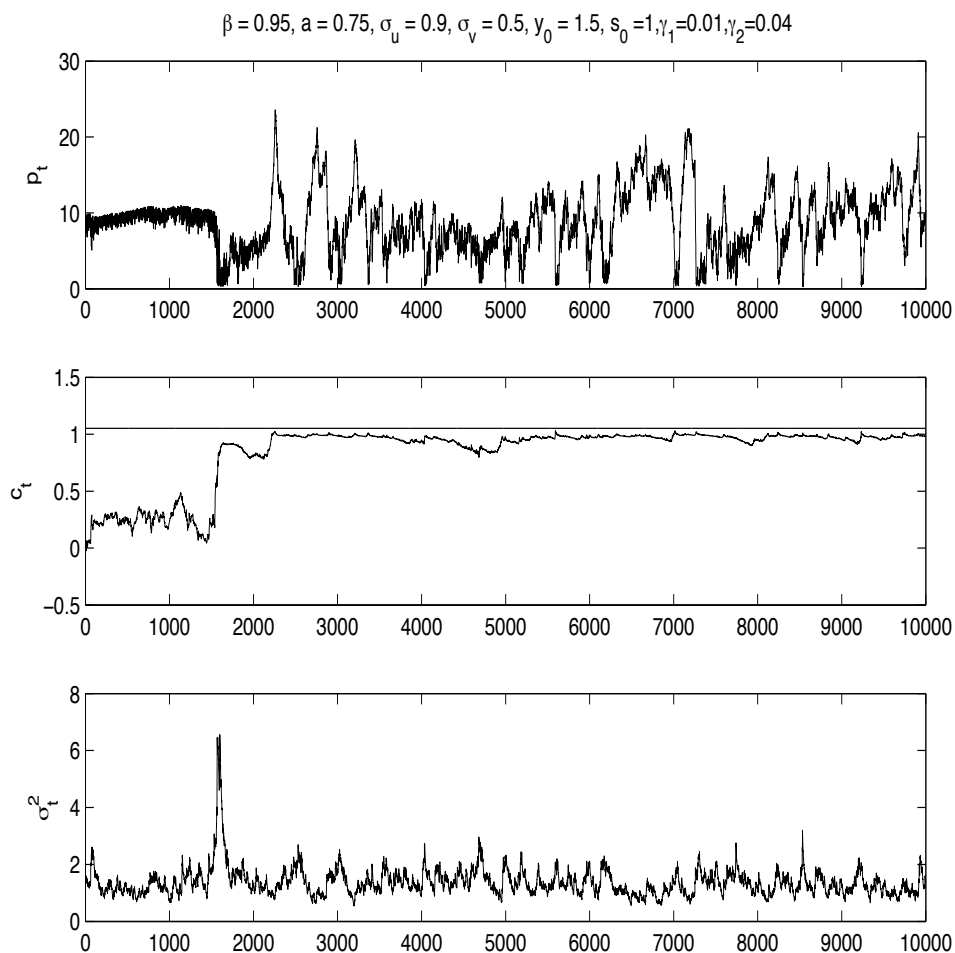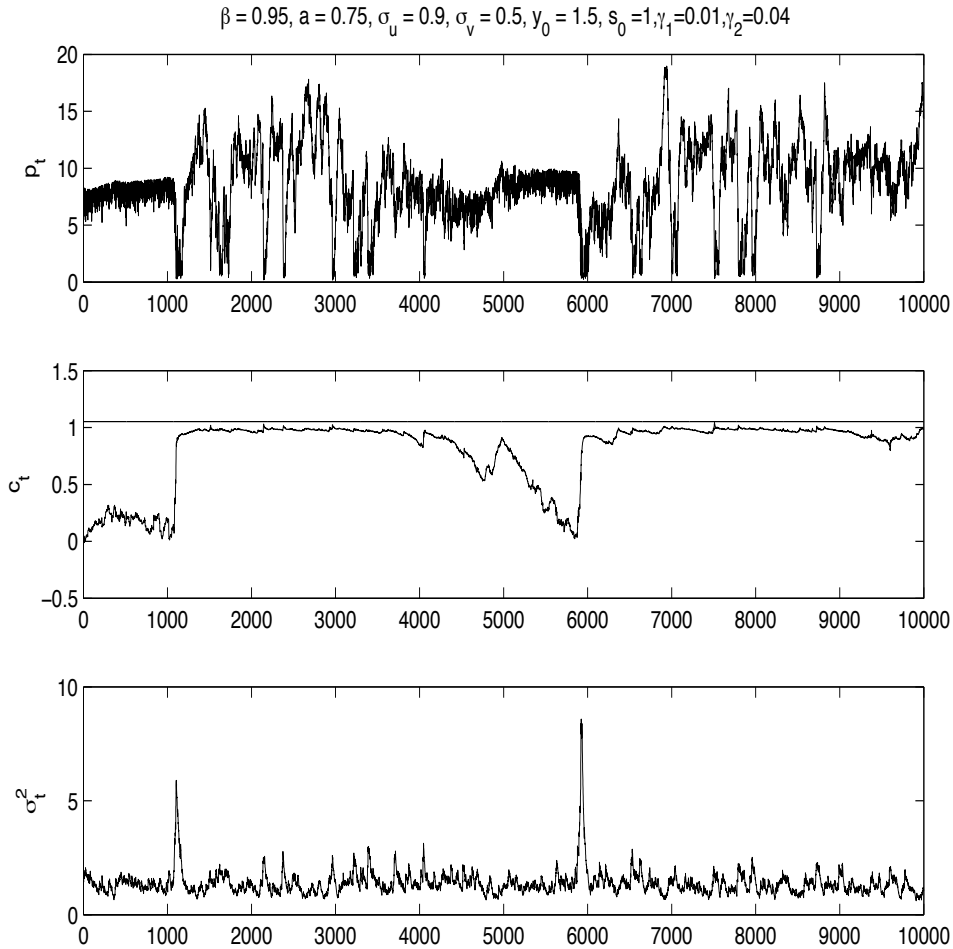
29

Figure 13: Constant gain learning with $\gamma_1 = .01, \gamma_2 = .04$.

The bottom two panels of Figure 12 illustrate how beliefs generate these recurrent bubbles and crashes. After the qualitative change in the dynamics around period 1500 there are frequent and sudden jumps in both σ^2 and c. This illustrates our central insight. Estimates of risk can propel a bubble: if agents suddenly expect lower risk then they are more willing to follow along a bubble path. However, the bubble path cannot persist forever: agents adapt their real-time estimates of risk; above a certain threshold the continued risk of staying on the bubble path is too high and this leads to a sudden crash of the stock price.[13] Figure 13 presents another simulation with the same parameter values that clearly illustrates this logic. In Figure 13, between periods 5000-6000 beliefs stay close to their fundamentals REE value, but there is

---

[13]In a sense, real-time estimation of risk σ^2 is acting like a projection facility in that it prevents estimates of c and k, and hence prices, from exploding. However, unlike a projection facility, the stabilizing role of σ^2 arises endogenously and has a natural economic interpretation.

then a sudden jump in c and σ^2 that leads to a crash and then a period of recurrent bubbles and crashes.

Interestingly, Figure 13 also demonstrates that the expected price path is strongly influenced by the mean dynamics. During periods 5000-6000 the dynamics stabilize and the belief parameters begin to drift back to their fundamentals REE levels, but again a sudden escape leads to recurrent bubbles and crashes.

One might wonder why beliefs do not appear to return to a neighborhood of the REE more often. Inspection of the eigenvalues of the Jacobian for the mean dynamics indicate that, for the specified parameter values, the speed of convergence is slow. In particular, there is an eigenvalue greater than $-.5$, which as Sargent (1999) notes, indicates convergence slower than the rate $t^{1/2}$.

Another issue that warrants briefly mentioning is the time scale and the frequency of bubbles. The current parameterization would suggest that bubbles occur about every 100 years or so, which is clearly not empirically realistic. By choosing values of β closer to 1, and selecting alternative gain parameters, it is possible to generate bubbles at a much higher frequency. However, the simulated stock prices become very noisy. Our parameter values $\beta = .95, \gamma_1 = .01$ and $\gamma_2 = .04$ were chosen because they generated figures that most clearly illustrate the mechanics of the model. A more carefully calibrated version of the model would require altering several modeling features as discussed below in Section 5.2.

5.1 Detecting Bubbles

Building on the insight that a bubbles REE can be decomposed into a fundamentals portion (e.g. dividends) plus a bubbles component, an extensive literature exists that tests for explosive asset price bubbles by testing for non-stationarity in the residuals from a regression of stock price on dividends. Diba and Grossman (1988) apply unit root tests to the residuals and reject the null hypothesis of bubbles in U.S. stock prices. Evans (1991), however, argues that, for a class of bubbles REE where the bubble periodically collapses, unit root tests, such as the Augmented Dickey Fuller (ADF) test, will fail to detect rational bubbles processes unless the frequency of collapse is very low. This is because periodically collapsing bubbles exhibit pseudo-stationarity, i.e. they mimic some key features of stationary data in finite samples. Similar results for a related class of rational bubbles have been found in Charemza and Deadman (1995). Building on this insight, Wu and Xiao (2002) propose an alternative Recursive Estimates Statistic (RES) which is better able to detect collapsible bubbles than conventional unit root tests. Their RES test indicates the presence of bubbles in Hong Kong stock market data and also provides some, but weaker, evidence of bubbles in U.S. stock prices. In order to illustrate the potential importance of our theoretical findings, we use Monte Carlo simulations of our simple model of bubbles to show

that, similar to these results, the RES, but not standard unit root tests, would detect the existence of bubbles. We find these results highly suggestive as they indicate that our model generates stochastic processes for asset prices that are consistent with key empirical features of the data.

The RES test is based on the residuals from a regression of the form

$$p_t = a + by_t + u_t,$$

which in our model is a constant plus noise. A test for the presence of bubbles is a test for non-stationarity in u_t. The RES constructs the following recursive statistic,

$$\max_{k=1,\ldots,n} \frac{k}{\sqrt{n}} \left| \frac{1}{k} \sum_{t=1}^{k} \hat{u}_t - \frac{1}{n} \sum_{t=1}^{n} \hat{u}_t \right| \tag{22}$$

where $\hat{u}_t$ is the estimated residual process and n is the length of the data sample. Wu and Xiao (2002) calculate the relevant critical values. Under an explosive bubble (22) will diverge to infinity, which gives the intuition for a simple test of the bubble.

To illustrate the potential empirical importance of our model we conducted the following experiment: we simulate the model for 15000 periods, storing the simulated price process. We next estimate the residuals by regressing price on a constant. We then conduct an ADF test of a unit root in the residuals[14] and an RES test for the presence of bubbles. The simulation was repeated 1000 times. For the ADF test, the alternative hypothesis is stationarity, so that bubbles are conventionally associated with a failure to reject the null. (As noted above, the Evans (1991) critique is that periodically collapsing bubbles will exhibit pseudo-stationarity). For the RES test, rejections are associated with bubbles. Table 1 presents the fraction of simulations in which the test detects bubbles based on 5% significance levels. Table 1 presents results for different values of the gain parameters as well as different sample sizes for the test (1000 or 100 periods).

| | Simulation Length: | | | |
| | 1000 | | 100 | |
Gain Parameters	RES	ADF	RES	ADF
$\gamma_1 = .001, \gamma_2 = .001$	0.024	0.000	0.0000	0.000
$\gamma_1 = .01, \gamma_2 = .01$	0.8760	0.0040	0.0220	0.000
$\gamma_1 = .01, \gamma_2 = .04$	0.980	0.0210	0.8340	.0640

Table 1. Bubbles Test Results for Simulated Data. For ADF, each cell entry reports the percentage of simulations consistent with a bubble (i.e.

---

[14]In the ADF test the lag length for each simulation is chosen using F-tests to ensure inclusion of all lags significant at the 5% level.

that fail to reject the null hypothesis of a bubble). For RES, each cell entry reports the percentage of simulations consistent with a bubble (i.e. that rejects the null hypothesis of no bubble). RES refers to the Recursive Estimates Statistic of Wu and Xiao (2002), ADF is an Augmented Dickey-Fuller test.

For very small gains both tests reject the presence of bubbles, as expected, since the simulated data remain close to the fundamentals rational expectations equilibrium. For gains that generate bubbles in our model, i.e. $\gamma_1 = .01, \gamma_2 = .04$, the results are also in line with our expectations: the RES test can detect bubbles even though the ADF test rejects a unit root in the residuals in favor of stationarity. The results for $\gamma_1 = .01, \gamma_2 = .04$ are consistent with the findings for Hong Kong and U.S. data in Wu and Xiao (2002) and Asian data in Chan and Woo (2007). In summary, our simulated test results combined with the above mentioned evidence suggest that an appropriately calibrated model could yield empirically plausible bubbles and crashes.

5.2 Further Discussion

We have developed a simple linear asset-pricing model capable of generating bubbles and crashes provided that agents use constant-gain learning to forecast expected returns and the conditional variance of stock returns. Risk aversion and uncertainty regarding the riskiness of the asset is a distinct and critical characteristic of our model. Risk generates the kind of drift and abrupt changes in beliefs – in response to fundamental shocks – that places the stock market onto a bubble path. Risk also bursts the bubble along an explosive price path when agents' risk estimates eventually increase. We believe this is a very reasonable and intuitive description of bubbles and crashes.

Our approach is closely related to others. Clearly, Sargent (1999, Chp. 6), in a different framework, anticipates some of the mechanisms that drive our asset pricing results. The learning dynamics are similar to the hyperinflation analysis of Marcet and Nicolini (2003) and Sargent, Williams, and Zha (2007) in that occasional shocks can trigger escapes in the learning dynamics. Timmermann (1994, 1996) examines learning in a present value model of asset pricing model and Carceles-Poveda and Giannitsarou (2006) study asset pricing with constant-gain learning in an RBC-type model, but these papers do not examine the implications of learning about asset price volatility. A distinguishing feature of our model is that risk plays a central role, both helping to trigger bubbles but also to reinforce the global stability of the fundamentals-based rational expectations equilibrium. Similar to our paper, Hong, Scheinkman, and Xiong (2005) assume that traders have mean-variance preferences and that there is asset float. In their paper, bubbles arise because insiders (those "floating" asset shares) and outsiders have different information about the underlying

asset. Outsiders overestimate the value, bidding up the price, and then when the lock-ups expire insiders sell their shares and prices crash. In our paper, asset float is a necessary component for the environment to provide agents an incentive to estimate the variance of returns. Again, however, real-time estimation of risk by private agents is the driving factor of our model.

An issue that should be addressed in future research is the choice of the time interval. There are three separate questions: the length of private agents' planning horizon; the frequency with which they update their recursive models; and, the frequency with which they update their information sets. In the present paper, for theoretical convenience these are all chosen to be the same unit. In work in progress, we construct a model with planning horizons that are longer than the estimation and information gathering horizons. This introduces additional complexity to the model that nonetheless would be important for a serious empirical exercise.

6 Conclusion

This paper generates bubbles and crashes in a simple linear asset pricing model with adaptive learning. The existence of recurrent bubbles in a model with adaptive learning has been an open question in macroeconomics. Our central insight is that in an environment in which traders are risk averse and boundedly rational, in the sense that they know the reduced form of the actual law of motion governing prices but not the parameters, then they must forecast both the conditional mean and the conditional variance of stock returns. We show that, when agents adopt constant-gain econometric learning, the qualitative nature of the dynamics can generate frequent deviations from the fundamentals solution taking the form of bubbles and crashes.

We identify two roles for real-time learning of risk. First, occasional shocks can lead agents to revise their estimates of risk in dramatic fashion. A sudden decrease in the estimated risk of a stock can propel the system away from the efficient-markets fundamentals equilibrium and into a bubble. Second, along a bubble path, risk estimates will increase until eventually the perceived risk is so high that asset demand will collapse and stock prices will crash. Thus, risk in an adaptive learning setting plays a central role in triggering and collapsing asset price bubbles. These results are intuitive and provide insights into the role adaptive learning and bounded rationality play in large swings in asset prices.

Appendix

Proof of Proposition 2

We sketch the proof to this proposition by making use of Proposition 7.8 of Evans and Honkapohja, itself a re-statement of Benveniste, Metivier, and Priouret (1990, Theorem 7, Chp. 4.4.3, Part II). The proposition in the text is based on the proposition stated below. Let D be an open set containing the fundamentals REE parameters $\theta^*, S^*, \sigma^{2*}$. In the case of exogenous share supply, the actual law of motion followed by price is

$$p_t = T(k_{t-1}, c_{t-1})X_{t-1} - \beta a \sigma_{t-1}^2 v_t.$$

It is clearly the case that the state dynamics are conditionally linear and can be written as

$$\bar{X}_t \equiv \begin{bmatrix} X_t \\ X_{t-1} \\ u_t \\ v_t \end{bmatrix} = \begin{bmatrix} A(\phi_{t-1}) & 0 & 0 & 0 \\ I & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \bar{X}_{t-1} + \begin{bmatrix} B(\phi_{t-1}) & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} W_t$$

where $I, 0$ are conformable matrices, and

$$X_t = A(\phi_{t-1})X_{t-1} + B(\phi_{t-1})W_t$$

with $X_t' = (1, p_t)', W_t' = (1, u_t, v_t)'$. The validity of the proposition depends on the following properties as established in Evans and Honkapohja (2001).

P1 W_t is iid with finite absolute moments.

P2 For any compact $Q \subset D$, $\sup_{\phi \in Q} |B(\phi)| \leq M$ and $\sup_{\phi \in Q} |A(\phi)| \leq \rho < 1$, and $|\cdot|$ is an appropriately defined matrix norm.

P3 For any compact $Q \subset D$, $\exists C, q$ s.t. $\forall \phi \in Q$ and for all t $|\mathcal{H}(\phi, x)| \leq C(1 + |x|^q)$.

P4 For any compact $Q \subset D$, $\mathcal{H}(\phi, x)$ is twice continuously differentiable with bounded second derivatives.

P5 $h(\phi)$ has continuous first and second derivatives on D.

The conditional linearity simplifies verification of these conditions. In particular, Proposition 7.5 of Evans and Honkapohja (2001) show that conditions M1-M5 of their Proposition 7.8 are implied by P1-P2. For their assumption $A3'$ we also make use of the remark on p. 155, which shows that P4 is sufficient.

For given ϕ let $p_t(\phi) = T(k, c)X_{t-1} - \beta a \sigma^2 v_t$ and let $X_t(\phi)' = (1, p_t(\phi))'$. Then $X_t(\phi)$ is stationary for ϕ sufficiently close to the fundamentals REE. Therefore, fix D to be an open set around $(\theta^*, S^*, \sigma^{2*})$ such that $\forall (\theta, S, \sigma^2) \in D$

1. $(\theta^*, S^*, \sigma^{2*})$ are such that σ^{2*} is the unique solution in D to the quadratic $\sigma_u^2 + (a\beta\sigma^2)^2 - \sigma^2 = 0$, θ^* is the unique fixed point of $T(\theta)$ on D with $\sigma^2 = \sigma^{2*}$, $S^* = EX_{t-1}(\phi^*)X_{t-1}(\phi^*)'$,

2. for some $\tilde{\varepsilon} > 0$, $\det(S) \geq \tilde{\varepsilon} > 0$,

3. and $k(1 + c) \geq y_0$ and $-1 < c < \bar{c} < \beta^{-1/2}$.

Write $\bar{X}_t = \bar{A}(\phi_{t-1})\bar{X}_{t-1} + \bar{B}(\phi_{t-1})W_t$, where the expressions for $\bar{A}, \bar{B}$ are given above. Clearly the eigenvalues of $\bar{A}$ consist of zero and the eigenvalues of A. The set D is defined so that the roots of $A(\phi)$ are inside the unit circle implying $\bar{A}(\phi)$ will also have roots with modulus less than one. It is straightforward to verify that assumptions P1-P5 hold.

We employ the following result from Evans and Honkapohja (2001):

Proposition 3 (EH(2001), Proposition 7.8) *Assume P1-P5. Consider the normalized random variables* $U^\gamma(t) = \gamma^{-1/2}\left[\phi^\gamma(t) - \tilde{\phi}(t, \phi_0)\right]$. *As* $\gamma \to 0$, *the process* $U^\gamma(t)$, $0 \leq t \leq T$, *converges weakly to the solution* $U(t)$ *of the stochastic differential equation*

$$dU(t) = D_\phi h(\tilde{\phi}(t, \phi_0))U(t)dt + \mathcal{R}^{1/2}(\tilde{\phi}(t, \phi_0))dW(t)$$

with initial condition $U(0) = 0$, *where* $W(t)$ *is a standard vector Wiener process, and* $\mathcal{R}$ *is a covariance matrix whose* i, jth *elements are*

$$\mathcal{R}^{ij}(\phi) = \sum_{k=-\infty}^{\infty} Cov\left[\mathcal{H}^i(\phi, \bar{X}_k^\phi), \mathcal{H}^j(\phi, \bar{X}_0^\phi)\right]$$

Finally Proposition 2 can be established by noting that the solution to the stochastic differential equation $U(t)$ has the following properties

$$
\begin{aligned}
EU(t) &= 0 \\
\frac{dVar(U(t))}{dt} &= D_\phi h(\tilde{\phi}(t, \phi_0))V_u(t) + V_u D_\phi h(\tilde{\phi}(t, \phi_0))' + \mathcal{R}(\tilde{\phi}(t, \phi_0)),
\end{aligned}
$$

where $V_u = Var(U(t))$.

Details on Approximating the Mean Dynamics With Endogenous Share Supply. The recursive algorithm (16)-(18) is for the case of exogenous share supply. When share supply may become endogenous requires additional care in constructing the mean dynamics. The condition for exogenous supply, $s_0 \leq \Phi p_t$, is satisfied if and only if

$$s_0 \leq \Phi\frac{\beta\left(k(1+c) + y_0\right)}{1 + \beta a\sigma^2\Phi(1 + v_t)} + \Phi\frac{\beta c^2}{1 + \beta a\sigma^2\Phi(1 + v_t)}p_{t-1}, \text{ or}$$

$$s_0\Phi^{-1} + s_0\beta a\sigma^2(1 + v_t) \le \beta\left(k(1 + c) + y_0\right) + \beta c^2 p_{t-1}. \tag{23}$$

Given $\theta = (k, c.\sigma^2)$, equations (3), (4) and (23) specify $p_t = F(p_{t-1}, v_t; \theta)$. For computing mean dynamics the complication is that whether (23) is satisfied, and thus whether (3) or (4) applies, depends on v_t.

Mean dynamics are computed by fixing θ and R and computing the ODE, where the expectation is taken over v_t and $p_t(\theta)$, the p_t process for fixed θ. In general this must be done using the process given by (3), (4) and (23), and for any given θ one must take account of the possibility that either regime will occur, depending on v_t. However, at least for "small" v_t, a reasonable approximation would be to split the θ space into two regions: in one region the probability is high that (for the given θ) the $p_t(\theta)$ process will be given by (3), and in the other region the probability is high that the $p_t(\theta)$ process will be given by (4).

For the (3) region $p_t(\theta)$ converges to a stationary AR(1) with mean

$$Ep_t(\theta) = \frac{\beta\left(k(1 + c) + y_0 - a\sigma^2 s_0\right)}{1 - \beta c^2} \equiv \bar{p}_H,$$

provided $\beta c^2 < 1$. If $\beta c^2 > 1$ the condition $s_0 \le \Phi p_t$ is satisfied (for $\lim_{t\to\infty} Ep_t(\theta)$). For $\beta c^2 < 1$ the condition is satisfied, using the above expression for $Ep_t(\theta)$ provided

$$s_0\Phi^{-1} + s_0\beta a\sigma^2 \le \beta\left(k(1 + c) + y_0\right) + \beta c^2\bar{p}_H.$$

Here we have set $v_t = 0$, and replaced p_{t-1} by its mean under (3). The condition can be rewritten as

$$\sigma^2 \le \bar{\sigma}_H^2(c, k), \text{ where}$$

$$\bar{\sigma}_H^2(c, k) = (s_0\beta a)^{-1}\left\{\beta\left(k(1 + c) + y_0\right) - s_0\Phi^{-1} + \beta c^2\bar{p}_H\right\}.$$

For the (4) region the linear approximation of the $p_t(\theta)$ process is of the form

$$p_t = \frac{\beta\left(k(1 + c) + y_0\right)}{1 + \beta a\sigma^2\Phi} + \frac{\beta c^2}{1 + \beta a\sigma^2\Phi}p_{t-1} - \delta v_t, \tag{24}$$

which has mean

$$Ep_t = \bar{p}_L \equiv \frac{\beta\left(k(1 + c) + y_0\right)}{1 - \beta c^2 + \beta a\sigma^2\Phi}.$$

Here

$$\delta = \frac{\beta^2 a\sigma^2\Phi\left(k(1 + c) + y_0 + \beta c^2\bar{p}_L\right)}{(1 + \beta a\sigma^2\Phi)^2}$$

Based on this mean, the condition $s_0 > \Phi p_t$ for (4) (with approximation (24)) will be satisfied when

$$\sigma^2 > \bar{\sigma}_L^2(c, k), \text{ where}$$

$$\bar{\sigma}_L^2(c, k) = (s_0\beta a)^{-1}\left\{\beta\left(k(1 + c) + y_0\right) - s_0\Phi^{-1} + \beta c^2\bar{p}_L\right\},$$

where we again set $v_t = 0$ and where we set p_{t-1} at its mean under (24). Since $\bar{p}_L < \bar{p}_H$ we have $\bar{\sigma}_L^2(c, k) < \bar{\sigma}_H^2(c, k)$. Thus when $\sigma^2 > \bar{\sigma}_H^2(c, k)$ and the distribution of v_t has small enough support, it is very likely that the (approximate) dynamics (24) will be followed.

In the main text we present numerical results for the mean dynamics based on the above approximation. Thus, for $\sigma^2 \leq \bar{\sigma}_H^2(c, k)$, we assume the mean dynamics are based on exogenous supply. For $\sigma^2 > \bar{\sigma}_H^2(c, k)$ the mean dynamics are instead assumed to be given by the alternative mean dynamics based on (24). Note for (24) the corresponding mapping from PLM to ALM has k, c components

$$(k, c) \rightarrow \left(\frac{\beta \left(k(1 + c) + y_0 \right)}{1 + \beta a \sigma^2 \Phi}, \frac{\beta c^2}{1 + \beta a \sigma^2 \Phi} \right).$$

and there is a corresponding expression for σ^2 component of the ODE:

$$h_{\sigma^2} = (T(\theta) - \theta) M(\theta, S, \sigma^2) (T(\theta) - \theta)' + \sigma_u^2 + \delta^2 \sigma_v^2$$

It is worth remarking that this procedure ignores the chance that the process will have endogenous supply when $\sigma^2 \leq \bar{\sigma}_H^2(c, k)$ and it ignores the chance that it will have exogenous supply when $\sigma^2 > \bar{\sigma}_H^2(c, k)$. Within and near the region $\bar{\sigma}_L^2(c, k) < \sigma^2 < \bar{\sigma}_H^2(c, k)$ the approximation will be at its worst, since both regimes will have a significant chance of arising. But in order to provide intuition for the real time learning results, this approximation suffices.

Procedure for Computing the Confidence Ellipses

Here we outline how we computed the confidence ellipsoids. Details on the general procedure are given in Evans and Honkapohja (2001, Chp. 14, p. 348-356). The confidence ellipsoids assume that the parameter estimates k_t, c_t will be distributed asymptotically normal. Under similar assumptions to those for Proposition 2 this property can be established formally.

In Evans and Honkapohja (2001) it is shown that $\theta_t \sim N(\theta^*, \gamma V)$ for small γ and large t, where $\theta' = (k, c)'$ and V solves the matrix Riccati equation

$$D_\theta h(\bar{\phi}) V + V (D_\theta h(\bar{\phi}))' = -\mathcal{R}_\theta(\bar{\phi})$$

where $\mathcal{R} = E \mathcal{H}(\phi) \mathcal{H}(\phi)'$ is as given in the proof to Proposition 2. Notice that the way this Riccati equation is expressed omits the $D_S h(\bar{\phi})$ and $D_{\sigma^2} h(\bar{\phi})$ terms. This is because R is a block diagonal matrix:

$$\mathcal{R} = E \mathcal{H}(\bar{\phi}) \mathcal{H}(\bar{\phi})' = \begin{bmatrix} (a\beta)^2 (\bar{\sigma}^2)^2 \sigma_v^2 M^{-1} & 0 & 0 \\ 0 & Evec \mathcal{H}_R vec \mathcal{H}_R' & 0 \\ 0 & 0 & \sigma_u^2 + (a\beta)^2 (\bar{\sigma}^2)^2 \sigma_v^2 - \bar{\sigma}^2 \end{bmatrix}$$

where $M = E X_{t-1} X_{t-1}'$. The text solves V numerically, sets $\gamma = .05$, and plots the 50% and 90% concentration ellipses.

References

[1] Adam, Klaus, Albert Marcet, and Juan Pablo Nicolini, 2007, "Learning and Stock Market Volatility," mimeo.

[2] Allen, Franklin, Stephen Morris, and Andrew Postlewaite, 1993, "Finite Bubbles with Short Sale Constraints and Asymmetric Information," *Journal of Economic Theory*, 61, 206-229.

[3] Benveniste, A., M. Metivier, and P. Priouret, 1990, *Adaptive Algorithms and Stochastic Approximations*, Springer-Verlag, Berlin.

[4] Bernanke, Ben S., 2002, "Asset Price Bubbles and Monetary Policy," remarks by Governor Ben S. Bernanke before the New York Chapter of the National Association for Business Economics, October 15, 2002, NY, NY.

[5] Blanchard, Olivier J., 1979, "Speculative Bubbles, Crashes, and Rational Expectations," *Economics Letters*, 3, 4, 387-389.

[6] Blanchard, Olivier J., and Mark W. Watson, 1982, "Bubbles, Rational Expectations, and Financial Markets," in *Crisis in the Economic and Financial Structure: Bubbles, Bursts, and Shocks,* Paul Wachtel (editor), Lexington Books.

[7] Bohm, V., and Chiarella, C., 2005, "Mean-Variance Preferences, Expectations Formation and the Dynamics of Random Asset Prices", Mathematical Finance, 15, 1, p. 61-97.

[8] Branch, William A., and George W. Evans, 2006, "A Simple Recursive Forecasting Model," *Economics Letters*, 91, 158-166.

[9] Branch, William A., and George W. Evans, 2006, "Asset Return Dynamics and Learning," mimeo.

[10] Brock, William A., and Cars H. Hommes, "Heterogeneous Beliefs and Routes to Chaos in a Simple Asset Pricing Model", *Journal of Economics Dynamics and Control*, 22 (1998), 1235-74.

[11] Carceles-Poveda, Eva and Chryssi Giannitsarou, 2006, "Asset Pricing with Adaptive Learning," mimeo.

[12] Chan, Hing Lin, and Kai-Yin Woo, 2007, "Testing Stochastic Explosive Root Bubbles in Asian Emerging Stock Markets," forthcoming *Economics Letters.*

[13] Charemza, W.W., and D.F. Deadman, 1995, "Speculative Bubbles with Stochastic Explosive Roots," *Journal of Empirical Finance*, 2, 153-163.

[14] Cho, In Koo, Noah Williams, and Thomas J. Sargent, 2003, "Escaping Nash Inflation," *Review of Economic Studies*, 69, 1, 1-40.

[15] Cochrane, John H., 2001, "Asset Pricing," Princeton University Press.

[16] Cochrane, John H., 2005, "Asset Pricing Program Review: Liquidity, Trading and Asset Prices," mimeo.

[17] DeLong, J. Bradford, Andrei Schleifer, Lawrence H. Summers, and Robert J. Waldmann, 1990, "Noise Trader Risk in Financial Markets," *Journal of Political Economy*, 98, 703-738.

[18] Diba, Behzad T., and Herschel I. Grossman, "The Theory of Rational Bubbles in Stock Prices," *Economic Journal*, 98, 392, 746-754.

[19] Evans, George W., 1991, "Pitfalls in Testing for Explosive Asset Prices," *American Economic Review*, 81, 4, 922-930.

[20] Evans, George W., and Seppo Honkapohja, 1998, "Economic Dynamics with Learning: New Stability Results," *Review of Economic Studies*, 65, 23-44.

[21] Evans, George W., and Seppo Honkapohja, 2001, "Learning and Expectations in Macroeconomics", Princeton University Press, Princeton, NJ.

[22] Froot, Kenneth A., and Maurice Obstfeld, 1991, "Intrinsic Bubbles: The Case of Stock Prices," *American Economic Review*, 81, 1189-1214.

[23] Grandmont, Jean-Michel, 1998, "Expectations Formation and Stability of Large Socioeconomic Systems," *Econometrica*, 66, 741-781.

[24] Greenspan, Alan, 2005, "Reflections on Central Banking," speech delivered at the Federal Reserve Bank of Kansas City Jackson Hole Symposium August 26, 2005.

[25] Hong, Harrison, Jose Scheinkman, and Wei Xiong, 2005, "Asset Float and Speculative Bubbles," forthcoming *Journal of Finance*.

[26] Hong, Harrison, Jose Scheinkman, and Wei Xiong, 2006, "Advisors and Asset Prices: A Model of the Origins of Bubbles," mimeo.

[27] Kindleberger, Charles P., 1977, *Manias, Panics, and Crashes*, Basic Books, New York.

[28] LeRoy, Stephen, and Richard Porter, 1981, "The Present Value Relation: Tests Based on Variance Bounds," *Econometrica*, 49, 555-557.

[29] Lewellen, Jonathan, and Jay Shanken, 2002, "Learning, Asset-Pricing Tests, and Market Efficiency," *Journal of Finance*, 57, 3, 1113-1145.

[30] Marcet, Albert, and Juan-Pablo Nicolini, 2003, "Recurrent Hyperinflation and Learning," *American Economic Review*, 93, 5, 1476-1495.

[31] Marcet, Albert, and Thomas J. Sargent, 1989, "Convergence of Least Squares Mechanisms in Self-Referential Stochastic Models," *Journal of Economic Theory*, 48, 337-368.

[32] McGough, Bruce, 2006, "Shocking Escapes," *Economic Journal*, 116, 507-528.

[33] Milani, Fabio, 2007, "Expectations, Learning and Macroeconomic Persistence," *Journal of Monetary Economics*, 54, 7, 2065-2082.

[34] Muth, John F., "Optimal Properties of Exponentially Weighted Forecasts," *Journal of the American Statistical Association*, 69, 618-626.

[35] Ofek, Eli, and Matthew Richardson, 2003, "Dotcom Mania: The Rise and Fall of Internet Stock Prices," *Journal of Finance*, 58, 1113-1137.

[36] Orphanides, Athanasios, and John C. Williams, 2005a, "The Decline of Activist Stabilization Policy: Natural Rate Misperceptions, Learning, and Expectations," *Journal of Economic Dynamics and Control*, 29, 1927-1950.

[37] Orphanides, Athanasios, and John C. Williams, 2005b, "Imperfect Knowledge and Monetary Policy," in *The Inflation Targeting Debate*, edited by Ben S. Bernanke and Michael Woodford, NBER and University of Chicago Press, Chicago, 201-234.

[38] Santos, Manuel, and Michael Woodford, 1997, "Rational Asset Pricing Bubbles," *Econometrica*, 65, 29-57.

[39] Sargent, Thomas J., 1999, *Conquest of American Inflation*, Princeton University Press.

[40] Sargent, Thomas J., Noah Williams, and Tao Zha, 2006a, "Shocks and Government Beliefs: The Rise and Fall of American Inflation," forthcoming *American Economic Review*.

[41] Sargent, Thomas J., Noah Williams, and Tao Zha, 2007, " The Conquest of South American Inflation," mimeo.

[42] Shiller, Robert J., "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends," *American Economic Review*, 71, 421-436.

[43] Timmermann, Allan, 1994, "Can Agents Learn to Form Rational Expectations? Some Results on Convergence and Stability of Learning in the UK Stock Market," *Economic Journal*, 104, 777-798.

[44] Timmermann, Allan, 1996, "Excess Volatility and Predictability of Stock Returns in Autoregressive Dividend Models with Learning," *Review of Economic Studies*, 523-557.

[45] West, Kenneth D., 1987, "A Specification Test for Speculative Bubbles," *Quarterly Journal of Economics*, 102, 553-580.

[46] Williams, Noah, 2004, "Escape Dynamics in Learning Models," mimeo.

[47] Wu, Guojun, and Zhijie Xiao, 2002, "Are There Speculative Bubbles in Stock Markets? Evidence From an Alternative Approach," mimeo.