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
University of St. Andrews

Revised: August 23, 2010

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 need to make forecasts of the conditional variance of a stock's return. Recursive updating of both the conditional variance and the expected return implies several mechanisms through which learning impacts stock prices. Extended periods of excess volatility, bubbles and crashes arise with a frequency that depends on the extent to which past data is discounted. A central role is played by changes over time in agents' estimates of risk.

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

1

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   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)). 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.

In this paper, we consider 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 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 estimated and 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 surprising and interesting ways 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 an econometric forecasting rule. For the illustrative model parameterization adopted in subsequent sections, it turns out that in a fundamentals based rational expectations equilibrium the mean stock price is about 8.7, and along an equilibrium path price is simply a constant plus white noise, with a standard deviation of about 0.7. Under learning, the dynamics can undergo an abrupt change leading to the recurrent bubbles and crashes illustrated in Figure 1.[1] Rational expectations prices lie between 7.3 and 10.1 ($\pm$ two standard deviations) about 95% of the time and between 6.6 and 10.8 ($\pm$ three standard deviations) over 99.7% of the time, whereas there are frequent departures from these ranges in the learning environment we analyze: in the Figure 1 simulation, prices are in the 6.6 to 10.8 range less than 40% of the time.

Figure 1 about here

To establish our results, we consider a simple asset pricing model in which the

---

[1]The paramerization is that used for Figure 12.

stock price today depends on expected *cum* dividend price next period and negatively on share supply, meant to proxy asset float. Share supply and dividends are both assumed to follow exogenous iid processes. We assume agents are risk-averse so that they seek to forecast both the expected net return tomorrow and the conditional variance of excess returns. It turns out that by requiring agents to also estimate the conditional variance, the global learning dynamics of our model are dramatically different.

The analysis in this paper identifies several channels through which agents' adaptive learning about risk and return affects stock prices. Occasional shocks to fundamentals can lead agents to adjust their estimates for risk and expected return; combined, these two forces cause stock prices to deviate from their fundamental values. For example a sustained period of small shocks to prices can lead to a downward revision in risk estimates that raises stock prices. More generally, various specific sequences of shocks, reinforced by the feedback from adaptive beliefs, introduce serial correlation that would not otherwise exist, and can lead agents' forecasting rule to track this serial correlation via a random walk forecasting model. Random walk beliefs can be approximately self-fulfilling and various scenarios for stock prices are possible, including bubbles and crashes. Changing estimates of risk are also useful in explaining how explosive bubbles can crash suddenly. If stock prices follow a bubble path, estimates of risk will increase along a bubble path. Eventually, the increased risk estimates can lead to decreased demand for the risky asset and falls in the stock price, at which point demand collapses and price crashes well below its fundamental value.

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. Section 6 discusses our results in the context of the literature, and Section 7 concludes.

2 Model

We employ a simple mean-variance linear asset pricing model, similar to DeLong, Shleifer, Summers, and Waldmann (1990).[2] 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 $R = \beta^{-1} > 1$, where β is the discount

---

[2]See the Appendix for details of the set-up.

factor. In this environment, demand for the risky asset is

$$z_{dt} = \frac{E_t^*(p_{t+1} + y_{t+1}) - \beta^{-1}p_t}{a\sigma_t^2},$$

where $E_t^*(p_{t+1} + y_{t+1})$ denotes the conditional expectation of $p_{t+1} + y_{t+1}$ based on the agent's subjective probability distribution and σ_t^2 denotes the agents conditional variance of excess returns $p_{t+1} + y_{t+1} - Rp_t$. The equilibrium price p_t is given by $z_{dt} = z_{st}$, where z_{st} is the (random) supply of the risky asset at time t.

It follows that

$$p_t = \beta E_t^* (p_{t+1} + y_{t+1}) - \beta a\sigma_t^2 z_{st} \tag{1}$$

For $a = 0$, equation (1) reduces to the standard risk-neutral asset-pricing formula, which can also 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. If agents are risk-neutral and financial markets complete then $\beta R = 1$, where β is the discount rate.[3]

When $a > 0$, equation (1) can be derived from an overlapping generations model along the lines of DeLong et al. (1990). As we note in the Appendix, this leads to mean-variance preferences when agents have constant absolute risk aversion (CARA) utility and believe that returns are normally distributed. Mean-variance preferences are a frequently employed approach to tractably modeling limited risk tolerance (downward sloping asset demand) and gives rise to a mean-variance maximizing setting in which agents optimize their portfolio by maximizing risk-adjusted expected wealth. See, for example, Bohm and Chiarella (2005) 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 conditional variance of returns, σ_t^2. Agents' concern with risk makes σ_t^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_t^2 will sometimes arrest explosive bubbles and can lead to crashes.

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 $\sigma_t^2 = \text{Var}_t^* (p_{t+1} + y_{t+1} - Rp_t) = \text{Var}_t^* (p_{t+1} + y_{t+1})$. Having price depend explicitly on z_{st} implies that price depends

---

[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.

on agents' perceived risk. In subsequent sections, we demonstrate that the nature of stock price dynamics depends, in part, on how perceived risk adjusts to the fundamental shocks to dividends and outside share supply. A stochastic process for share supply is meant to proxy for variations in the availability of publicly tradeable shares (i.e. asset float). Variations in asset float 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 (2009)) as many IPO's have a large share of their offerings "locked up" by insiders and, therefore, not publicly tradeable. Here we motivate the presence of asset share supply by appealing to this literature and noting that short sales constraints in markets (such as IPO lock up restrictions) can give rise to an important role to supply variation in asset pricing. The Appendix presents a stylized overlapping generations model where variation in per capita share supply can arise from stochastic variation in the population of young agents. Equation (1) makes it clear that with risk-neutral agents share supply does not have an effect on price.

We assume that the exogenous process for dividends is

$$y_t = y_0 + u_t,$$

with $y_0 > 0$ and $Eu_t = 0$. Share supply is assumed to follow a multiplicative process of the form

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

where u_t, V_t are uncorrelated iid shocks with $EV_t = 1$, and $s_0 > 0$. Here $\Phi = s_0/(\bar{p}\xi)$, where $\bar{p}$ is the mean stock price in a fundamentals based REE and $0 < \xi < 1$ is a fairly small proportion. In our numerical illustrations we set $\xi = 0.1$, which implies that share supply is exogenous except when price falls below 10% of its fundamentals value. The endogeneity of share supply at low prices is meant to capture asset float drying up in financial markets that perform poorly. This ensures that price remains non-negative, thereby providing a price floor in the event of a crash in stock prices.

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 $E^* = E$. The fundamentals-based REE p_t^f can be found by assuming $\sigma_t^2 = \sigma^2$ and iterating (1) forward to give $p_t = \sum_{j=1}^{\infty} \beta^j E_t y_{t+j} - \beta \sum_{j=0}^{\infty} \beta^j a\sigma^2 E_t z_{st+j}$. There is additionally a class of bubbles REE, which are given by adding to the fundamentals solution a "rational bubble" term $\beta^{-t}\eta_t$, where η_t is an arbitrary martingale, i.e. $E_t\eta_{t+1} = \eta_t$. For $0 < \beta < 1$ the bubbles REE is explosive. To generate empirically plausible 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)).

6

Provided the support of the supply z_{st} is not too large, the model will have REE in which share supply is always exogenous. Letting $V_t = 1 + v_t/s_0$, i.e. $s_0 V_t = s_0 + v_t$, where v_t is an *iid* mean-zero disturbance, and restricting attention to solutions with $\sigma_t^2 = \sigma^2$, the model becomes

$$p_t = \beta E_t^* p_{t+1} + \beta(y_0 - a\sigma^2 s_0) - \beta a\sigma^2 v_t. \tag{2}$$

The fundamentals solution takes the form

$$p_t = \beta(1 - \beta)^{-1}(y_0 - a\sigma^2 s_0) - \beta a\sigma^2 v_t,$$

and it can be shown that the bubbles solutions have the alternative representation

$$p_t = a\sigma^2 s_0 - y_0 + \beta^{-1} p_{t-1} - a\sigma^2 v_{t-1} + \xi_t$$

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

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 assume that agents attempt to learn, in real-time, about the underlying stochastic process followed by the stock price, in particular about the conditional mean and variance of the excess rate of return. Because the model is self-referential, agents' learning can produce, as endogenous reactions to the intrinsic fundamental shocks, periodic bubbles and crashes.

To this end we depart from rational expectations by following, e.g., Marcet and Sargent (1989a) and Evans and Honkapohja (2001), and assume that agents form expectations based on a perceived law of motion, i.e. a forecasting model that allows for both the fundamentals and bubbles REE. With a forecasting rule in hand, agents forecast next period's stock price and expected returns, and use these expectations to form their demands for the risky asset. The equilibrium stock price will be the one that clears the market. Specifically, we assume that at the beginning of period t agents forecast p_{t+1} based on the following perceived law of motion:

$$p_t = k + cp_{t-1} + \varepsilon_t, \tag{3}$$

where ε_t is an unobserved white-noise disturbance. Under learning k and c are parameters estimated from past data. With the perceived law of motion (3), subjective conditional expectations are

$$E_t^* p_{t+1} = k(1 + c) + c^2 p_{t-1}. \tag{4}$$

For convenience we are adopting the timing assumption that no contemporaneous variables, including z_{st}, are observable to agents at t when they make their forecasts.[4] To ensure stock prices remain non-negative, we also impose that $k(1 + c) \geq -y_0$.

---

[4]Thus $E_t^* p_{t+1} = E_t^* (k + c(k + cp_{t-1} + \varepsilon_t) + \varepsilon_{t+1}) = k(1 + c) + c^2 p_{t-1}$. The stability properties of the REE solutions under learning does not hinge on the details of the information assumption.

The demand for the risky asset also depends on σ^2, the subjective conditional variance of excess returns, which is also estimated from past data. Inserting these beliefs into (1), using again $s_0 V_t = s_0 + v_t$ and allowing for the possibility of endogenous share supply, yields the actual law(s) of motion

$$p_t = \beta(y_0 + k(1+c) - a\sigma^2 s_0) + \beta c^2 p_{t-1} - \beta a\sigma^2 v_t, \text{ if } s_0 \leq \Phi p_t \tag{5}$$

$$p_t = \frac{\beta\left(k(1+c) + y_0\right)}{1 + \beta a\sigma^2\Phi(1+v_t)} + \frac{\beta c^2}{1 + \beta a\sigma^2\Phi(1+v_t)} p_{t-1}, \text{ if } s_0 > \Phi p_t, \tag{6}$$

Thus asset prices p_t are determined by (5)-(6), given agents' estimates of (c, k, σ^2) based on past data. Under real-time learning the parameter estimates are updated each period, using the most recent data point, as described in Section 3.2 below. Apart from the details of the adaptive learning rules, this completes the description of how the path of asset prices under learning is determined. The first issue we study is whether under learning there can be convergence to REE over time.

3 Stability under Learning

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

3.1 Expectational stability

To examine local stability of the REE under learning we consider the mapping from the perceived law of motion (3) to the implied actual law of motion (5)-(6). In this section we treat perceived risk σ^2 as a given parameter. If beliefs are sufficiently close to an REE, and provided p_{t-1} is not too low, then asset share supply will be exogenous and the actual law of motion can be re-written simply as

$$p_t = T(k, c)(1, p_{t-1})' - \beta a\sigma^2 v_t, \tag{7}$$

where

$$T(k, c) = \left(\beta(y_0 + k(1+c) - a\sigma^2 s_0), \beta c^2\right) \tag{8}$$

defines a map from the perceived to the actual law of motion. There are two fixed points of the T-map (8), $(\beta(y_0 - a\sigma^2 s_0)/(1 - \beta), 0)$ and $(a\sigma^2 s_0 - y_0, \beta^{-1})$, which correspond to the fundamentals and bubbles REE.

The T-map (8) can be interpreted as follows. If agents had perceived law of motion given by (3), with parameters (k, c) fixed at possibly non-RE values, and if agents maintained these fixed parameters over time, then their forecast rule would be given by (4). With exogenous share supply the stochastic process followed by p_t would take the same functional form as (3), but with parameters $T(k, c)$ instead of (k, c). Not surprisingly, REE are fixed points of the T-map, since an REE occurs when perceptions and outcomes align.

Under real-time learning the parameters (k, c) are not fixed, and instead are adjusted gradually over time using least-squares to update their values in response to the evolving data. The econometric learning literature, e.g. Evans and Honkapohja (2001), has shown that the T-map can be used to compute a simple stability condition, known as E-stability, which governs whether or not RE parameters are locally stable under learning, and that the differential equation, used to defined E-stability, also provides information on the global dynamics under learning.[5]

Formally, the E-stability principle states that locally stable rest points of the ordinary differential equation

$$\frac{d(k, c)'}{d\tau} = (T(k, c) - (k, c))' \tag{9}$$

will be attainable under least squares and closely related learning algorithms.[6] That the E-stability principle governs stability under learning is intuitive, because under (9) the parameters (k, c) are gradually adjusted in the direction of the actual law of motion parameters generating the data given these perceptions. Local stability of (9) at an REE thus answers the question of whether under such adjustment a small displacement of (k, c) would return to the REE. It is straightforward to apply the E-stability principle to assess the local stability under learning of the REE in our set-up. In fact, Evans and Honkapohja (2001) show, in a closely related model,[7] that with $0 < \beta < 1$ (i) the fundamentals REE $(\beta(y_0 - a\sigma^2 s_0)/(1 - \beta), 0)$ is E-stable, and hence stable under learning, and (ii) 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 the

---

[5]When analyzing the global dynamics below we allow for endogenous supply.

[6]Here τ denotes notional time, which however can be linked to real time t as discussed below.

[7]For the case with fixed and known σ^2, the formal structure of the model under learning is also similar to the hyperinflation model of Marcet and Sargent (1989b).

process under learning to diverge from the bubbles REE, observing such equilibria seems unlikely.

3.2 Stability with learning about risk

In the previous section the stability under learning was examined, while taking as given the agents perception of risk σ^2. In an REE σ^2 is an equilibrium object, and it is also natural, and we would argue crucial, to extend the analysis of learning to include learning about the degree of risk.

Recalling that $\sigma^2 = Var_t(p_{t+1} + y_{t+1} - \beta^{-1}p_t)$, it follows that in an REE

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

In the case of the fundamentals REE,

$$\sigma^2 = E_t \left(-a\beta\sigma^2 v_{t+1} + u_{t+1}\right)^2 \tag{10}$$

The right-hand side of this equation can be viewed as giving, for any specified perceived value of σ^2, the implied actual value of σ^2, and solutions to (10) deliver the REE values for the fundamentals REE:

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

There are two positive solutions, but we will see that it is the smaller root σ_L^2 that is stable under learning. For the bubbles REE straightforward calculations show that $\sigma^2 = \sigma_u^2 + \sigma_\xi^2$. We remark that in the fundamentals REE, p_t is affected directly by the supply shock, but not the dividend shock. However the variances of both shocks affect the distribution of p_t via σ^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 simple 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. Assuming that at time t agents use parameters estimated using data through time $t - 1$, and that their forecasts are conditioned 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 is

$$p_t = \beta \left(y_0 + k_{t-1}(1 + c_{t-1})\right) + \beta c_{t-1}^2 p_{t-1} - \beta a\sigma_{t-1}^2 z_{st}, \tag{11}$$

where $z_{st} = s_0 + v_t$ if supply is exogenous. Allowing for endogenous supply (11) can be rewritten as (5)-(6) with k, c, σ^2 replaced by $k_{t-1}, c_{t-1}, \sigma_{t-1}^2$.

Letting $X_t = (1, p_t)'$, the real-time learning algorithm can be written as

$$
\begin{align}
\theta_t &= \theta_{t-1} + \gamma_{1,t} S_{t-1}^{-1} X_{t-1} \left(p_t - \theta_{t-1}' X_{t-1} \right) \tag{12}\\
S_t &= S_{t-1} + \gamma_{1,t+1} \left(X_t X_t' - S_{t-1} \right) \tag{13}\\
\sigma_t^2 &= \sigma_{t-1}^2 + \gamma_{2,t} \left(\left(p_t - \theta_{t-1}' X_{t-1} + u_t \right)^2 - \sigma_{t-1}^2 \right). \tag{14}
\end{align}
$$

The first two equations in (12)-(14) are the updating equations for recursive least squares. Here S_t is an estimate of $EX_t X_t'$, the second moment matrix of the regressors, which is needed for least-squares updating. Equation (14) is a recursive algorithm for estimating the conditional variance of net returns.[8]

For the stability results in this section we assume the "gains" $\gamma_{1,t}, \gamma_{2,t}$ are set to $\gamma_{1,t} = \gamma_{2,t} = t^{-1}$ as in standard least squares. For the results on mean dynamics, given in the next section, and for the numerical simulations, we instead assume constant gains and allow $\gamma_{1,t} = \gamma_1 \neq \gamma_2 = \gamma_{2,t}$. With constant-gain learning, the recursive algorithm becomes a form of discounted least squares. Decreasing gains allow for the possibility of full convergence to REE, and are thus convenient for studying local stability questions. Constant gains are appropriate if agents want to allow for the possibility of structural change of an unknown form, and also have the advantage that the system is time-invariant with an ergodic distribution that can be studied. Under decreasing gains equation (14) in effect estimates the conditional variance by the sample mean of squared excess returns, as would be appropriate in the fundamentals REE, while with a constant gain $0 < \gamma_2 < 1$ the algorithm can track drifting volatilities of an unspecified form. More general formulations could be considered in which the conditional variance was assumed by agents to depend on observables, in which case the algorithm for learning about second moments would look more like the algorithm for learning about first moments.

The first, and most basic, stability question is whether the E-stability results for the fundamentals and bubbles REE, given in the preceding section, carry over to the current setting in which estimates of risk, as well as the coefficients of the price process, are updated in real time. Evans and Honkapohja (2001) provide conditions that ensure convergence of recursive systems like (12)-(14). These conditions draw on convergence theorems for stochastic recursive algorithms. For now, we assume that initial beliefs lie in the region in which share supply is exogenous, in which case the

---

[8]Agents' period t estimate of the conditional variance of excess returns is now denoted σ_{t-1}^2 because it is based on data through period $t - 1$.

price process under learning is

$$p_t = \beta \left(y_0 - a s_0 + k_{t-1}(1 + c_{t-1}) \right) + \beta c_{t-1}^2 p_{t-1} - \beta a \sigma_{t-1}^2 v_t. \qquad (15)$$

Later we illustrate how weak convergence results are impacted by (possibly) endogenous supply.

To study the stability under learning of an REE, for the case of decreasing gain, the approach is to use a continuous time approximation to (12)-(15), the fit of which improves as time gets large. Defining $\phi_t = (\theta_t, \text{vec}(S_t), \sigma_t^2)'$, one can write (12)-(14) as

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

where $\bar{X}_t' = (1, p_t, p_{t-1}, u_t, v_t)'$. Results from stochastic approximation theory show that asymptotically the dynamics of (12)-(15) are governed by the associated ODE (ordinary differential equation)

$$\begin{aligned} \frac{d\phi}{d\tau} &= h(\phi), \text{ where} \qquad\qquad (16) \\ h(\phi) &= \lim_{t \to \infty} E\mathcal{H}(t, \phi, \bar{X}_t(\phi)). \end{aligned}$$

Here $\phi = (\theta, \text{vec}(S), \sigma^2)'$ and τ is "notional" time. The explicit computation of $h(\phi)$ is given in the Appendix, and details of the technique are described in Marcet and Sargent (1989) and Evans and Honkapohja (2001). Local stability of this ODE governs the local stability of the REE under (12)-(15). In the Web Appendix we show:

Proposition 1 *Consider the model (1) with exogenous share supply. Under the adaptive learning algorithm (12)-(15) with gains $\gamma_{1,t} = \gamma_{2,t} = t^{-1}$:*

1. *The fundamentals REE with $\sigma^2 = \sigma_L^2$ is locally stable under learning.*

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 (1989a) 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. As we will now see, with constant-gain learning, bubble-like global dynamics can periodically emerge

as temporary large deviations from the fundamentals REE. Furthermore the increases in perceived risk along these bubble paths eventually acts 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. Thus, the onset of recurring bubbles and crashes will arise from the global dynamic properties of the model under learning. As shown below in Proposition 2, the ODE (16) also provides insight on global dynamics under both decreasing and constant gain learning. In particular, the ODE (16) describes the mean dynamics of the parameters under learning from given initial values.

This section provides numerical examples of the global learning dynamics in order to illustrate the possible theoretical outcomes. Numerical parameter choices are made in order to best demonstrate the possible dynamics rather than to provide an empirically meaningful calibration. This section is organized as follows. First, in section 4.1 we show that, even if expectations start at the fundamentals REE, under constant-gain learning belief parameters will vary randomly around it according to a distribution that is scaled by the gain parameter, i.e. to the sensitivity of agents' beliefs to recent data. Second, we will show that if belief parameters wander far enough away from the fundamentals REE, the path back to the REE, governed by the mean dynamics, may include a period in which stock prices are believed to follow a random walk. In section 4.2 we show that such periods are likely to be sustained because random walk beliefs are approximately self-confirming. Finally, section 4.3 considers more specifically which "escape paths" are most likely to drive the system large distances away from the fundamentals and to generate random walk beliefs.

Figure 2 illustrates one key part of the intuition that can be understood in terms of the E-stability results of Section 3.1. Consider either REE, and fix σ^2 and S at the REE values. Examination of the form of the ODE (16), given in the Appendix, shows that the dynamics of $\theta' = (k, c)$ are then given simply by the E-stability differential equation (9). In addition, from (8) it can be seen that the c component of the T-map is $T_c = \beta c^2$, so that c evolves independently of k. These observations imply that with σ^2, S fixed at REE values, the mean dynamics of c are autonomous and given by $dc/d\tau = T_c(c) - c$. The figure demonstrates how learning dynamics adjust perceptions about the slope parameter c (horizontal axis) in response to outcomes (vertical axis).

In an REE T_c intersects the 45° line. There are clearly two REE, the fundamentals at $c = 0$ and the bubbles at $c = 1/\beta$. The arrows in the figure show the direction of adaptation of c. For initial values $c > 1/\beta$ we have $T_c(c) - c > 0$, so that c tends to be revised upward, and the ensuing estimates c will tend to explode. On the other hand, for initial $c < 1/\beta$ we have $T_c(c) - c < 0$, so that c is revised downward, with convergence to the fundamentals REE value $c = 0$.

Figure 2 about here

Although trajectories originating in $[0, 1/\beta)$ will eventually settle down at the fundamentals REE, the global dynamics along a convergent path could still be interesting. In particular, away from the fundamentals REE, the dynamics introduce serial correlation into p_t. This serial correlation may be self-reinforcing leading agents to (temporarily) believe that $c > 0$, and in some cases paths will arise in which the agents believe that the price is close to random walk behavior. We will later see that such paths are associated with bubbles and crashes.

Additional insight can be obtained by studying the global dynamics of the ODE (16), the solutions to which can be shown to provide the "mean dynamics" to real-time learning under both decreasing and constant gain. Anticipating the real-time simulations of Section 5, we use constant gains $\gamma_{1,t} = \gamma_1$ and $\gamma_{2,t} = \gamma_2$, allowing also for $\gamma_1 \neq \gamma_2$. As noted above, constant gains are better able to track the stochastic process generating the data when there is structural change taking an unknown form. Constant gain learning has been widely used in the learning literature as discussed further in the Section 6. One implication is that instead of converging asymptotically to an REE, estimates can converge to a stationary process around a stable REE.

We begin with formal results, for the case of exogenous supply, on the ODE approximation for the case of small constant gains. Details of the case where share supply may become endogenous are in the Web Appendix. Fixing the ratio γ_2/γ_1 at some value, we use ϕ_t^γ to denote the value of $\phi_t = (\theta_t, vec(S_t), \sigma_t^2)'$ under the process (12)-(14) when γ_1 is set at some (small) value $\gamma_1 = \gamma$. In order to make a comparison between solutions to the continuous time ODE and to the discrete time recursive algorithm, we need to define a corresponding continuous time sequence for ϕ_t^γ, which we denote as $\phi^\gamma(t)$. To construct $\phi^\gamma(\tau)$, we set $\tau_t^\gamma = t\gamma$, and define $\phi^\gamma(\tau) = \phi_t^\gamma$ if $\tau_t^\gamma \leq \tau < \tau_{t+1}^\gamma$. The following proposition establishes the mean dynamics result in a neighborhood of the fundamentals REE, and also provides information on the stochastic distribution.

Proposition 2 *Assume exogenous share supply. For any ϕ_0 within a suitable neighborhood of the fundamentals REE, define $\tilde{\phi}(\tau, \phi_0)$ as the solution to the ODE $d\phi/d\tau =$*

14

$h(\phi)$, with initial condition ϕ_0. Consider the random variable, indexed by the constant gain γ, $U^\gamma(\tau) = \gamma^{-1/2}\left(\phi^\gamma(\tau) - \tilde{\phi}(\tau, \phi_0)\right)$. As $\gamma \to 0$, $U^\gamma(\tau), 0 \leq \tau \leq T$, converges weakly to a zero mean random variable.

The proof is contained in the Web Appendix. We remark that the "neighborhood" of validity of this proposition need not be small and, as shown in the Web Appendix, can include a wide range of values for ϕ. (The neighborhood must also include the trajectory $\tilde{\phi}(\tau, \phi_0)$ for $0 \leq \tau \leq T$).[9] This result establishes that, over finite periods of time, the constant-gain learning dynamics will converge weakly to the solution of the ODE $d\phi/d\tau = h(\phi)$, where $\tau \approx \gamma t$. Thus the "mean dynamics" approximate the expected path, under real-time learning with a small constant gain, from given initial conditions. 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.[10]

Section 4.1 demonstrates that, if beliefs are displaced away from the REE, the transitional path may include agents temporarily believing stock prices follow a random walk. Section 4.2 further shows that such random-walk beliefs are almost self-fulfilling. One way to try to understand how beliefs of this form might arise is the approach of Cho, Williams and Sargent (2003) (CWS), who use the notion of "escape dynamics" as a way of characterizing the "most likely unlikely" shock process that will lead a model away from a rational expectations equilibrium. CWS show that an ODE, similar to the mean dynamics ODE, governs the path for beliefs that move *away* from a neighborhood of the REE. This ODE takes the form $d\phi/d\tau = h(\phi) + \dot{v}$ where $\dot{v} = v(\phi)$. CWS interpret $\dot{v}$ as arising from a continuous time approximation to the constant gain learning algorithm under a "most likely unlikely" distribution for the shock process. The intuition behind their analysis is to look for a sequence of shocks that moves the system out of a neighborhood of the REE via the shortest, or least costly, route. To give insight into the types of escape dynamics possible in this model, in Section 4.3 we simulate the model under various "unlikely" sequences of shocks.

4.1 Mean dynamics

Constant gain learning allows estimates to be more alert to structural change, but it also makes agents' beliefs more responsive to shocks. Consequently random divi-

---

[9]Similar results can be expected to hold in the case where share supply may become endogenous, but verification of the technical conditions in this case are difficult. Instead, we use an approximation and then present numerical results.

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

dend and supply shocks continue to displace the system from the fundamentals RE solution. The resulting displacements trigger mean dynamics that can sometimes temporarily move further away from the fundamentals REE. Section 4.3 studies what types of shocks might provide the trigger to move beliefs far away from their REE values. 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 are more likely to arise.

To illustrate this reasoning Figure 3 plots the 95% and 50% confidence ellipses for (k, c) around the fundamentals REE assuming relative constant gains $\gamma_2/\gamma_1 = 2$, for reasons which will become apparent below. To compute this figure we follow Evans and Honkapohja (2001, Chp.7) who show that asymptotically, under constant-gain learning, the parameter estimates are approximately normally distributed around the REE value, with variance proportional to the gain.[11] This figure was generated by assuming the following baseline parameterization: $\beta = 0.95, a = 0.75, \sigma_u^2 = 0.9, \sigma_\nu^2 = 0.5, y_0 = 1.5, s_0 = 1$. Figure 3 illustrates that the confidence ellipses around the fundamentals REE have a decreasing principal axis, suggesting 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 larger c associated with smaller k along the principal axis. The relative size of these ellipses depends on the sizes of the constant gain. Figure 3 is our first indication that, under constant-gain learning, estimates of agents will occasionally evolve toward random-walk beliefs, with a frequency that is higher for larger gains.

Figure 3 about here

For the key parameters (k, c, σ^2), the confidence ellipsoid consists of the (k, c) ellipse in Figure 3 and a confidence region for the risk aversion parameter σ^2, which is a small interval around the fundamentals REE value.[12] One can think of constant-gain learning dynamics as re-initializing the mean dynamics. Figure 4 illustrates representative mean dynamics for an initial value of $c > 0$, with a corresponding k on the principal axis, and with initial σ^2 somewhat below its stable REE value.[13] Again we set $\gamma_2/\gamma_1 = 2$. Setting σ^2 below and c above their REE values corresponds to a decrease in perceived risk and to an increase in perceived serial correlation in price, so that initially mean prices are above the fundamental REE value. The figure plots

---

[11] See the Web Appendix for further details.

[12] In the asymptotic distribution, σ^2 is uncorrelated with k and c.

[13] The working paper version of this paper contains plots of the mean dynamics for other starting points.

the mean price along the learning path, the belief parameters c, k, and the perceived risk estimate σ^2.

<div align="center">Figure 4 about here</div>

The fundamentals REE is seen to be a stable rest point for the mean dynamics, in line with Proposition 1. However, in addition, the transition path for the mean dynamics, in particular the behavior of the autoregressive parameter c, is very interesting. At first the estimate of c moves toward the fundamentals REE, but then it reverses course and increases to a value of $c = 1$, where it remains for a period before eventually converging to $c = 0$. This evolution is accompanied by an increase in σ^2, including a sharp spike, before returning to its fundamental value. Note that $k \approx 0$ during the period during which $c \approx 1$. Thus the mean dynamics show agents coming to believe that stock prices approximately follow a random walk. Along the path, the mean price implied by k, c, σ^2 begins well above the fundamentals with $c > 0$ and σ^2 below its REE value, but then collapses along with the temporary increase in σ^2.

Through numerical explorations, we found that greater sensitivity in updating estimates of σ^2 was more likely to trigger random walk beliefs in the mean dynamics. For this reason, in the real-time dynamics below, we choose values of $\gamma_2 > \gamma_1$. Is this choice empirically realistic? We believe so. Merton (1980) argues that under appropriate assumptions the instantaneous conditional variance of the excess return, in a continuous time framework, can be estimated much more precisely than can the conditional mean. However, what is relevant within our model is the conditional variance over the investment horizon of the representative agent. The strong volatility clustering observed empirically in excess returns suggests the need for a relatively large gain γ_2 to track the time variations in conditional variance.

Revisions of risk estimates together with random-walk beliefs play a key role in the learning dynamics. In a perceived low-risk environment traders will act on these beliefs and asset prices will be driven up. Similarly high risk estimates tend to drive asset prices down. In either case the resulting price dynamics push estimates of the price process towards a random walk. In essence, agents come to believe that recent changes in price are permanent shifts and not mean-reverting fluctuations. The random-walk beliefs are nearly self-fulfilling, as we will see next. Furthermore, random-walk beliefs, when combined with variation over time in risk estimates, tend to generate bubbles and crashes, as we will see in the simulations in Section 5. Besides generating departures from the fundamentals departures REE, another role played by risk is to crash bubbles: along an explosive price path, risk estimates can increase and eventually cause price to collapse. The relative gain γ_2/γ_1 is important for ensuring that bubbles will crash.

<div align="center">17</div>

4.2 Random-walk beliefs

The mean dynamics illustrated in the preceding section suggests that if agents' estimates evolve toward random-walk beliefs, they can stay close to these beliefs for an extended period of time before finally returning to the fundamentals REE values. To understand this, we follow Sargent (1999, Ch. 6), in adapting an insight from Muth (1960), and show that if agents hold random walk beliefs then the resulting stochastic process can be almost self-fulfilling in the sense that the deviation from rational expectations of the random walk approximation may be almost indetectable. The basic idea is that a random walk model approximates well a model with time-varying means.

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 learning model (3) provided $c = 1, k = 0$. Assuming also exogenous supply and $\sigma_t^2 = \sigma^2$, and plugging these beliefs into the equation for stock price (2), these beliefs imply the price process

$$p_t = \beta(y_0 - a\sigma^2 s_0) + \beta L p_t - \beta a\sigma^2 v_t \tag{17}$$

where L is the lag operator. In terms of $\mathrm{MA}(\infty)$ processes, under random walk beliefs,

$$p_t = (1 - L)^{-1}\varepsilon_t \tag{18}$$

while under these beliefs the actual price process is

$$p_t = \mu + f(L)v_t \tag{19}$$

where $\mu = (y_0 - a\sigma^2 s_0)\beta/(1 - \beta)$ and $f(L) = -\beta a\sigma^2/(1 - \beta L)$. From (17) one can see that σ^2, the one-step ahead conditional variance of returns is given by (10), the same expression as for the fundamentals solution. Thus, with the same value for σ^2, the mean price p_t under random walk beliefs is the same as for the fundamentals solution.

Sargent (1999) emphasizes two features of misspecified random-walk beliefs that are evident in (18) and (19). First, random walk beliefs introduce serial correlation into a model that is not serially correlated under rational expectations (in the fundamentals solution). The moving average processes (18) and (19) demonstrate that the perceived serial correlation becomes almost self-fulfilling. Second, random walk beliefs can track constants well. In (18) there is no constant but in (19) there is. The random walk uses higher-order moments to track low frequency movements (i.e. the mean) in the price process.

For the problem at hand a key point is that random-walk beliefs induce a price process that is almost self-fulfilling. 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. The result, given in Figure 5, shows that there is indeed a good match between spectral densities. We conclude that if a sequence of random shocks leads agents to have random-walk beliefs concerning asset prices, these beliefs may last for a substantial period of time. Intuitively, because random-walk beliefs are close to self-fulfilling, it is difficult to detect the misspecification except using long stretches of data. The mean dynamics do eventually return the system to the fundamentals REE, but transitional dynamics with random-walk beliefs will be important. Furthermore, with constant-gain learning, there can be periodic returns to random-walk beliefs and, as we shall see, the episodes of random-walk beliefs can, for some parameter settings, be dominant.

Figure 5 about here

In addition to being nearly self-fulfilling, random-walk beliefs also lead to a substantial amount of excess volatility, in the sense that the unconditional variance of p_t is much higher than in the fundamentals solution. This is a result of the strong serial correlation in prices under (19). For the same innovation standard deviation $\beta a \sigma^2 \sigma_v$, the unconditional variance of p_t under random-walk beliefs, compared to the fundamentals solution, is higher by a factor of $(1 - \beta^2)^{-1}$. For the same reason, changes in the estimate of the conditional variance of returns σ^2 will have a magnified effect on unconditional price volatility. The process (19) remains linked to the fundamentals solution, but the near-random walk behavior, with much larger volatility around the mean, implies that under these beliefs the price process is almost "detached" from the fundamentals solution. For $0 < \beta < 1$ with β near one, random-walk beliefs are close to the rational bubble beliefs, in which the autoregressive parameter is β^{-1}. The price process resulting from random walk beliefs might thus be viewed as a bubble regime of the model.

4.3   Escape paths

In section 4.1 we showed first that, even if expectations start at the fundamentals REE, under constant-gain learning belief parameters will vary randomly around it according to a distribution that is scaled by the gain parameter, i.e. to the sensitivity of agents' beliefs to recent data. Secondly we showed that if belief parameters wander far enough away from the fundamentals REE, the path back to the REE, governed by the mean dynamics, may include a period in which stock prices are believed to

follow a random walk. In section 4.2 we further showed that such periods are likely to be sustained because random walk beliefs are approximately self-confirming.

In the current section we consider more specifically which "escape paths" are most likely to drive the system large distances away from the fundamentals and to generate random walk beliefs. The literature on "large deviations" and "escape paths" studies this question by looking for the "most likely unlikely sequences" of shocks that will drive the system a given distance away from the equilibrium. The central idea is that over long stretches of time there will, with high probability, be shock sequences that lead to large deviations. One can look for which "unlikely sequences" of shocks are "most likely" to occur over a given stretch of time, and study the features of the resulting paths. Given γ_2/γ_1, in the small gain limit $\gamma_1 \to 0$ there is a dominant escape path.

To study this question we use a version of the technique employed in Cho, Williams, and Sargent (2003). We consider different sequences of trinomial shocks for u and v that take on fixed (over time) values in $\{-\sigma_u, 0, \sigma_u\}$ and $\{-\sigma_v, 0, \sigma_v\}$, respectively, and then simulate the model under learning starting from the fundamentals REE. We identify the various possible paths, with fixed values (over time) of the shocks drawn from these trinomial sets, that drive stock price away from the REE, and measure the relative likelihood of the alternative escape paths in terms of the speed with which the escapes occur. As in the study of the mean dynamics, we proceed numerically, using the baseline parameterization above with $\gamma_1 = .01, \gamma_2 = .02$. We initialize c, k, σ^2 and the sample second moment matrix S at the fundamental REE, and then simulate the model under specified non-random sequences of shocks. Once stock price rises to 1.5 times its fundamentals price (bubble) or 0.5 times the fundamentals price (crash) we say that an "escape" has occurred. We use trinomial shocks because, in our framework in which estimated risk plays a central role, we need zero shocks to consider the impact of "unlikely" sequences of very small shocks.

We first illustrate how a bubble might arise by looking at a simulation with (u, v) fixed at their mean value $(0, 0)$ each period. Figure 6 shows the path for p_t, the estimated AR(1) coefficient c_t and the perceived risk σ_t^2. The repeated shocks $(u, v) = (0, 0)$ induce decreases in σ_t^2. In line with the lower perceived risk, the demand for the risky asset increases and so does p_t. Under learning with small constant gains, price increases take place gradually. The bottom panel presents a scatter plot of p_t and p_{t-1}, showing that at the end of period 80 the data will lead the econometric model to fit a random walk for stock prices, i.e. a zero intercept and a slope coefficient equal to one. The upward trend in stock prices that leads to the least squares estimate for c_t to increase will, in turn, amplify the increase in p_t. Continuing in this manner, eventually, $c > \beta^{-1}$ and p_t explodes. The price dynamics from that point forward

depend critically on the perceived risk estimates as Figure 7 illustrates.

<center>Figures 6, 7 about here</center>

Figure 7 considers the sequence of "zero" shocks for two alternative gains on the risk estimates: $\gamma_2 = .0001$ and $\gamma_2 = .01$. We make agents' estimates (c, k) evolve slowly by setting $\gamma_1 = .005$. The left panels plot the simulations with $\gamma_2 = .0001$. A positive bubble arises as price increases gradually, leading to increasing estimates for c, which feeds back into further price increases.[14] As above, eventually, $c > \beta^{-1}$ and price explodes along a bubble path. Because the risk estimates adjust very slowly the explosive trend in price leads c to increase faster than σ^2 and there is no limit to the bubble. The right panels illustrate simulations for the same sequence of shocks but with the higher gain $\gamma_2 = .01$. The greater responsiveness of risk accelerates the upward trend in prices and the feedback through expectations. However, as $c > \beta^{-1}$ and price begins to explode, perceived risk increases sharply and arrests the explosive trajectory of prices.[15] Once prices begin to fall, a self-fulfilling crash occurs: expected price falls sharply and there are further sharp increases in σ^2. This feedback loop continues until price crashes to the floor. Thus, learning about risk can play two important roles in a bubble and its eventual crash: it can reinforce the feedback effect of expected price and, if sufficiently responsive, can eventually arrest an explosive upward movement in stock prices.

The simulations in Section 5 below show that the first "escape" path away from the fundamentals REE is often a crash rather than a positive bubble. To see how a crash might arise starting from the fundamentals REE Figure 8 simulates the model, starting from an REE, for the "unlikely" sequence of fixed shocks $(u, v) = (0, .2)$. Intuitively, we expect such an unlikely sequence to trigger a crash because positive share supply shocks will decrease stock prices, with crashes developing because of self-fulfilling expectations about stock prices.

<center>Figure 8 about here</center>

Figure 8 plots two separate simulations. The dotted line is for a small gain $\gamma_1 = .002$ while the solid lines are for $\gamma_1 = .01$. In each case (see the third panel), the risk

---

[14]The bottom panels plot early periods of the simulation in order to keep the scale illustrative of the emerging trend in price.

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

<center>21</center>

estimates are here assumed to be held fixed at their REE value (by setting $\gamma_2 = 0$) so as to focus on the expected price effects. Since we are considering a sequence of repeated fixed shocks $(u, v) = (0, .2)$, this eventually leads agents to expect prices to return not to its original fundamental value, but rather to a new "pseudo steady-state" $\bar{p}' \approx 4.7$. In the case of the small gain p_t moves monotonically to the new steady-state. With the larger gain p_t converges to the new steady-state price but there is over-shooting arising from price expectations with the price falling to the floor before returning to $\bar{p}'$. Thus, crashes arise during this sequence of shocks because of an overshooting in expectations via learning. When σ^2 adjusts in real-time as well, the crash is magnified by increases in σ^2 brought on initially by the sequence of shocks.

Table 1 about here

These sequences of shocks illustrate representative cases of the possible routes to bubbles or crashes. Table 1 presents the results from various "unlikely" sequences of fixed shocks. For each non-random sequence of trinomial shock, Table 1 gives the type of price dynamics, the time to reach a point outside of a neighborhood of the fundamental equilibrium, and the results from a long simulation. In each case we set $\gamma_1 = .01, \gamma_2 = .02$. Cho, Williams, and Sargent (2003) identify the "most likely unlikely" sequence of shocks as the one that will move a system to a point away from the equilibrium in the shortest number of periods. According to Table 1, the "most likely unlikely" sequence of shocks is in the third row with $(u, v) = (-\sigma_u, \sigma_v)$. These shocks are the most likely route away from the steady-state because they decrease price both directly, through increased share supply, and indirectly by increasing perceived risk σ^2. As we saw above, these direct and indirect effects on price produce dynamics well approximated by a random walk model and price eventually crashes to the floor. Table 1 indicates that fixed shocks $(u, v) = (-\sigma_u, \sigma_v)$ lead stock price away from the fundamentals in the shortest amount of time.[16] In most cases listed in Table 1 the resulting paths include both a crash and a bubble.

The escape path generated by the "most likely unlikely" sequence of shocks is the one that dominates in the small gain limit, but for finite gains a variety of escape paths will arise, and larger gains can exhibit a wide variety of escape dynamics. In addition to bubbles arising from sequences of zero (or near zero) shocks $(u, v) = (0, 0)$, Table 1 shows that bubbles will also arise in response to sequences of negative share supply shocks, which put prices on a gradual upward trajectory. In summary, depending on the shock realizations, the model with learning about risk and return is able to generate a rich set of theoretical possibilities for stock price dynamics.

---

[16]However, note that with normally distributed random shocks (u, v) values near $(0, 0)$ are more likely than values near one standard deviation.

22

5 Recurrent Bubbles and Crashes

The results in Sections 3 and 4 have indicated that while the fundamentals REE is locally stable under real-time learning, displaced estimates of risk and returns, sufficiently away from the fundamentals equilibrium value, can induce learning dynamics that send beliefs for a sustained period of time into a random-walk region that is nearly self-fulfilling and that exhibits a much higher level of price volatility. Changing estimates of risk play key roles in these dynamics by pushing prices away from the fundamentals equilibrium, leading to bubbles or crashes or bubbles followed by crashes, and in contributing to price volatility in the random-walk regime.

Under constant-gain real-time learning, in which agents discount past data, we anticipate the possibility of seeing a regime of bubbles and crashes periodically emerge from the fundamentals solution, before subsiding and returning to the fundamentals for a period of time before eventually again emerging. We would expect the frequency with which the regime of bubbles and crashes appears to be controlled by gain parameters. To study the issue further requires stochastic simulations for the system (11)-(14), with $\gamma_{1,t} = \gamma_1$ and $\gamma_{2,t} = \gamma_2$. Figures 9-12 present the numerical results. We choose the same parameter values as above, and in order to focus on the effects of the gain for risk estimates, γ_2, we fix $\gamma_1 = 0.01$ and look at the impact of varying γ_2. Larger gains γ_2 correspond to greater discounting of past data and hence a greater sensitivity to recent data. In each figure we report the results from a typical simulation of length 10,000, which follows a 5,000 length transient period.

Figures 9-12 about here

Figure 9 gives results from a simulation with $\gamma_1 = .01$ and a very small gain for risk estimates, $\gamma_2 = .001$. The top panel plots the stock price p_t, while the bottom panels plot the estimated autoregressive parameter c_t and the risk estimate σ_t^2, respectively. The belief parameters stay near their fundamentals REE value, and as a result p_t is close to the fundamentals REE, a constant plus a white noise process. If γ_2 is increased to $\gamma_2 = .01$ the plots (not shown) for p_t now exhibit a smooth low frequency process superimposed on the fundamentals REE, and the estimated value of σ_t^2 displays more volatility. However, the model still does not exhibit bubbles or crashes.

In Figure 10 the gain is increased to $\gamma_2 = .02$. Initially the dynamics look as they did in the previous figures, but beginning around 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

23

REE. Notice how the beliefs for c, σ^2 correspond to the mean dynamics pattern seen in Figure 4. Figure 11 plots a "zoomed in" portion of a typical simulation, plotting together the stock price p_t and the inverse risk-measure $1/\sigma_t^2$. The figure shows movements in the risk-estimate preceding large qualitative changes in the stock price. In particular, before a bubble episode there is a significant decrease in the perceived risk, while before the crashes the risk estimate increases. This is in line with the analysis in Section 4.3.

Figures 9-10 illustrate how learning about risk and return interact to create bubbles and crashes. In Figure 12 we increase γ_2 further to $\gamma_2 = .04$, which leads to a further qualitative shift. Again, initially the dynamics are not far from the fundamentals REE, but then around period 1600 there is a dramatic change in the nature of the price and belief dynamics, starting with jumps in both c_t and σ_t^2. These induce a 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. The price dynamics follow a path somewhat reminiscent of the detrended log S&P 500 index.

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 1600 there are frequent jumps in σ^2, and c spends considerable time near $c = 1$. For this parameter setting the random-walk beliefs regime becomes almost permanent. In this regime, prices remain centered at the fundamentals value, with positive and negative deviations about equally likely. With the larger gain $\gamma_2 = 0.04$, the endogenous shifts in volatility create sustained movements in prices that are well tracked by a random walk and largely offset the pressure from the mean dynamics to return to the fundamentals REE.

In Figures 9-10 and 12 the horizontal line in the middle panel, showing c_t, is set equal to $c = \beta^{-1}$, the rational bubble value, which is slightly in excess of $c = 1$. In Section 3 we saw that the rational bubble solutions were not locally stable under learning. However, in Section 4 we saw that, following some plausible random displacements from the fundamentals equilibrium, mean dynamics paths often visited for substantial periods of time the random-walk beliefs that are prominent in Figure 12. Because random-walk beliefs are close to rational bubble beliefs, they are almost self-fulfilling. Moreover, since random walk beliefs are nearly detached from the fundamentals value, and generate substantial excess volatility, it is natural to describe this as a bubble regime. In contrast to the rational bubbles literature, a central role in our model is played by agents' estimates of risk. Furthermore, in our model estimates of both returns and risk are driven by fundamental shocks. Revisions in risk are associated with escapes from the fundamentals solution and with sustaining the regime of bubbles and crashes.

A final issue that warrants comment 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 $\gamma_2 = .02, .04$, 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.

6 Further Discussion and Literature Review

We have developed a simple linear asset-pricing model capable of generating bubbles and crashes when agents use constant-gain learning to forecast expected returns and the conditional variance of stock returns. The approach here has been informed by an influential literature on periodically collapsing rational bubbles. 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.[17] A controversial issue for rational bubbles is that the trigger for the bubble collapse is modelled by an exogenous sunspot process. Our model instead predicts bubbles and crashes as self-fulfilling responses to fundamental shocks, arising from the adaptive learning of agents.

Our approach is also related to other strands of the literature. The learning dynamics are similar to the hyperinflation analysis of Marcet and Nicolini (2003) and Sargent, Williams, and Zha (2009) in that occasional shocks can trigger, via the learning dynamics, sudden departures from a rational expectations equilibrium. Adam, Marcet, and Nicolini (2009) 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. Timmermann (1994, 1996) examines learning in a present value model of asset pricing and Carceles-Poveda and Giannitsarou (2006) study asset pricing with constant-gain learning in an RBC-type model. Timmermann (1993, 1994, 1996), as in our model, uses adaptive learning to generate excess volatility in asset returns. We

---

[17]There is a wide literature that catalogs theoretical objections to bubbles. For instance, Diba and Grossman (1988) show that, since free disposal implies price can never be negative, if a bubble collapses to zero then a rational bubble can never again arise.

generate distinct results since Timmermann (1993, 1994, 1996) emphasizes traders forecasting the exogenous dividend process whereas our model of learning is self-referential, and since neither Timmermann, Carceles-Poveda-Giannitsarou, or Adam, et al. examine the implications of learning about asset price volatility. Finally, that learning can generate large stock returns has been pointed out by Geweke (2001) and Weitzman (2007) in a Bayesian learning context. These papers demonstrate that, with CRRA utility, Bayesian learning implies an infinite stochastic discount factor, a property that is not needed in our framework.

A distinguishing feature of our model is that risk plays a central role. 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, and it is the real-time estimation of risk by private agents that is a driving factor of our model.

The onset of bubbles and crashes, as illustrated in Figure 10, is reminiscent of the escape dynamics identified by Sargent (1999), Cho, Williams, and Sargent (2002), Williams (2004), and Cho and Kasa (2008). We showed that certain "unlikely" sequences of shocks, reinforced by the feedback from 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. This "escape" from a serially uncorrelated process to a serially correlated time-series, well approximated by a random walk, arises endogenously, and this shift in beliefs leads to recurrent bubbles and crashes.

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 longer than the estimation and information gathering windows. This introduces additional complexity to the model that would be important for a serious empirical exercise.

7 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 as well as boundedly rational, in the sense that they do not know the true law of motion governing prices, changes to their forecasts of both the conditional mean and the conditional variance of stock returns play a central role in asset price dynamics. In particular we show that when agents use constant-gain econometric learning, which to some extent discounts past data, learning dynamics can generate frequent deviations from the fundamentals solution taking the form of bubbles and crashes.

We identify several 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 or increase in the estimated risk of stocks can propel the system away from the fundamentals equilibrium and into a bubble or crash. Second, along an explosive bubble path, increased risk estimates tend to increase and can become high enough to lead asset demand to collapse and stock prices to crash. Third, under learning, estimates for stock returns will occasionally escape to random-walk beliefs that can be viewed as a bubble regime in which stock prices exhibit substantial excess volatility. In this regime revisions of risk estimates play an important role in generating the movements of prices that sustain the random-walk beliefs. In summary, risk in an adaptive learning setting plays a key role in triggering asset price bubbles and crashes. These intuitive and plausible results provide insights into the mechanisms by which expectations, learning and bounded rationality generate large swings in asset prices.

Appendix

Overlapping Generations Framework. We here describe a simple overlapping generations model, based on DeLong et al., which delivers the pricing equation (1). Agents live two periods. The number n_t of young agents is an identically and independently distributed random process with an inverse mean of one. There is a single consumption good. When young, each agent receives an endowment of ω units of the good. Agents consume only when old, with CARA utility, as described below. All of the endowment is saved, using one of two assets. Using a riskless storage technology agents receive $R = \beta^{-1} > 1$ units when old for every unit saved when young. Alternatively agents can purchase a risky asset, which is in fixed supply s_0. Because n_t is random, the per capita supply of the risky asset z_{st} is random, and we write $z_{st} = s_0 V_t$, where $V_t = 1/n_t$. The risky asset pays a random dividend paid the following period, $y_{t+1} = y_0 + u_{t+1}$, where u_t is white noise. The price of the risky asset

at time t is p_t and when old the agent, after receiving the dividend, sells the asset at price p_{t+1}.

Preferences take the CARA form

$$U(c_{t+1}) = -\exp\{-ac_{t+1}\},$$

where $a > 0$ is the coefficient of absolute risk aversion, and young agents choose their portfolio to maximize the conditional expectation of $U(c_{t+1})$. Agents assume that $p_{t+1} + y_{t+1}$ and hence c_{t+1} is normally distributed, and thus it is equivalent for them to maximize

$$E_t^* U(c_{t+1}) = -\exp\{-aE_t^* c_{t+1} + (a^2/2)\text{Var}_t^* c_{t+1}\}.$$

Here E_t^* denotes the conditional expectation and Var_t^* the conditional variance of a random variable, based on the subjective probability distribution of the agents. Letting z_{dt} denote the number of "shares" or units of the risky asset chosen by the young agents, their budget constraint is given by

$$c_{t+1} = (\omega - p_t z_{dt})\beta^{-1} + z_{dt}(p_{t+1} + y_{t+1}).$$

Thus

$$
\begin{aligned}
E_t^* c_{t+1} &= (\omega - p_t z_{dt})\beta^{-1} + z_{dt} E_t(p_{t+1} + y_{t+1}) \\
\text{Var}_t^* c_{t+1} &= z_{dt}^2 \text{Var}_t^*(p_t + y_{t+1}) \equiv \sigma_t^2,
\end{aligned}
$$

The optimal choice of z_{dt} must satisfy the first-order condition

$$-p_t \beta^{-1} + E_t^*(p_{t+1} + y_{t+1}) - az_{dt}\sigma_t^2 = 0, \text{ or}$$

$$z_{dt} = \frac{E_t^*(p_{t+1} + y_{t+1}) - \beta^{-1}p_t}{a\sigma_t^2}.$$

The equilibrium price p_t is determined by $z_{dt} = z_{st}$. Under the assumptions given above, per capita supply $z_{st} = s_0 V_t$ is exogenous, where z_{st} is iid with $Ez_{st} = s_0$. Under rational expectations, $E_t^*(p_{t+1} + y_{t+1}) = E_t(p_{t+1} + y_{t+1})$, the true conditional expectation under the objective probability distribution, and $\sigma_t^2 = \text{Var}_t(p_t + y_{t+1})$, the true conditional variance. For u_t and V_t independent normally distributed processes, the "fundamentals" solution is given in the text, and it can be shown that $p_{t+1} + y_{t+1}$ is normally distributed with $\text{Var}_t(p_t + y_{t+1})$ constant over time.

In the version of the model with endogenous supply at low prices, it is assumed that net supply is reduced when p_t falls sufficiently far. This might arise, for example, if there is another class of agents – e.g. long-lived agents with an alternative use of the

28

asset that becomes profitable at low prices – with a demand for the asset proportional to price when p_t falls below a specified threshold. This leads to a net supply of assets available to young agents that takes the form $z_{st} = \{\min(s_0, \Phi p_t)\} \cdot V_t$.

Stochastic Approximation ODE. Set $\gamma_{1,t} = \gamma_{2,t} = t^{-1}$ and let $z_t = p_t - \theta'_{t-1}X_{t-1} + u_t = \left(T(\theta_{t-1}; \sigma^2_{t-1}) - \theta_{t-1}\right)X_{t-1} - a\beta\sigma^2_{t-1}v_t + u_t$. Then (12)-(15), with exogenous supply, is

$$\theta_t = \theta_{t-1} + t^{-1}S^{-1}_{t-1}X_{t-1}\left(X'_{t-1}(T(\theta_{t-1}; \sigma^2_{t-1}) - \theta_{t-1})' - a\beta\sigma^2_{t-1}v_t\right) \qquad (20)$$

$$S_t = S_{t-1} + t^{-1}\left(\frac{t}{t+1}\left(X_tX'_t - S_{t-1}\right)\right) \qquad (21)$$

$$\sigma^2_t = \sigma^2_{t-1} + t^{-1}\left(z_tz'_t - \sigma^2_{t-1}\right), \qquad (22)$$

where (15) has been used to substitute for p_t under learning. Defining $\phi_t = (\theta_t, \text{vec}(S_t), \sigma^2_t)'$ this specifies the updating equation $\mathcal{H}(t, \phi, \bar{X}_t(\phi))$. Using the framework of Evans and Honkapohja (2001), the ODE (ordinary differential equation) associated with the asymptotic behavior of the learning algorithm is given by (16), with $h(\phi)$ components given by

$$h_\theta = S^{-1}M(\theta, \sigma^2)(T(\theta; \sigma^2) - \theta)'$$
$$h_S = M(\theta, \sigma^2) - S$$
$$h_{\sigma^2} = (T(\theta; \sigma^2) - \theta)M(\theta, \sigma^2)(T(\theta; \sigma^2) - \theta)' + \sigma^2_u + (a\beta\sigma^2)^2\sigma^2_v - \sigma^2,$$

where $M(\theta, \sigma^2) = EX_t(\theta, \sigma^2)X_t(\theta, \sigma^2)'$. Note that at an REE $S = M(\theta, \sigma^2)$ and the h_θ components reduce to $T(\theta; \sigma^2) - \theta$. Here $\theta = (k, c)'$ and $T(\theta; \sigma^2)$ is given by (8).

References

[1] Adam, Klaus, Albert Marcet, and Juan Pablo Nicolini, 2009, "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] 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.

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

[5] 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.

[6] 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.

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

[8] Branch, William A., and George W. Evans, 2009, "Asset Return Dynamics and Learning," forthcoming *Review of Financial Studies.*

[9] Carceles-Poveda, Eva and Chryssi Giannitsarou, 2006, "Asset Pricing with Adaptive Learning," forthcoming *Review of Economic Dynamics.*

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

[11] Cho, In Koo, and Kenneth Kasa, 2008, "Learning Dynamics and Endogenous Currency Crises," *Macroeconomic Dynamics,* 257-285, 12.

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

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

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

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

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

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

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

[19] Geweke, John, 2001, "A Note on Some Limitations of CRRA Utility," *Economics Letters*, 71, 3, 341-355.

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

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

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

[23] Hong, Harrison, Jose Scheinkman, and Wei Xiong, 2006, "Advisors and Asset Prices: A Model of the Origins of Bubbles," forthcoming *Journal of Financial Economics*.

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

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

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

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

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

[29] Marcet, Albert, and Thomas J. Sargent, 1989b, "Convergence of Least-Squares Learning and the Dynamic of Hyperinflation." In *Economic Complexity: Chaos, Sunspots, Bubbles, and Nonlinreaity*, eds. W. Barnett, J. Geweke, and K. Shell, Cambridge University Press, Cambridge, 119-37.

[30] Merton, Robert C., 1980, "On Estimating the Expected Return on the Market: An Exploratory Investigation," *Journal of Financial Economics*, 8, 323-361.

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

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

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

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

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

[36] Sargent, Thomas J., Noah Williams, and Tao Zha, 2009, " The Conquest of South American Inflation," *Journal of Political Economy*, 117, 211-256.

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

[38] Timmermann, Allan, 1993, "How Learning in Financial Markets Generates Excess Volatility and Predictability of Excess Returns," *Quarterly Journal of Economics*, 108, 1135-1145.

[39] 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.

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

[41] Weitzman, Martin L., 2007, "Subjective Expectations and Asset-Return Puzzles," *American Economic Review*, 97, 4, 1102-1130.

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

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

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

Web Appendix to "Learning about Risk and Return: A Simple Model of Bubbles and Crashes" by William A. Branch and George W. Evans

Proof of Proposition 1. We use the stochastic approximation results described in Evans and Honkapohja (1998, 2001) and Marcet and Sargent (1989). Set $\gamma_{1,t} = \gamma_{2,t} = t^{-1}$ and define $z_t = p_t - \theta'_{t-1}X_{t-1} + u_t = \left(T(\theta_{t-1}; \sigma^2_{t-1}) - \theta_{t-1}\right)X_{t-1} - a\beta\sigma^2_{t-1}v_t + u_t$. Then (12)-(15) in the main text, for the case of exogenous supply, can be re-written as the equations (20)-(22) in the print Appendix to the paper, reproduced here:

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

$$S_t = S_{t-1} + t^{-1}\left(\frac{t}{t+1}\left(X_tX'_t - S_{t-1}\right)\right)$$

$$\sigma^2_t = \sigma^2_{t-1} + t^{-1}\left(z_tz'_t - \sigma^2_{t-1}\right).$$

Defining $\phi_t = (\theta_t, \text{vec}(S_t), \sigma^2_t)'$, 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 learning algorithm (16), i.e.

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

is given by

$$h_\theta = S^{-1}M(\theta, \sigma^2)(T(\theta; \sigma^2) - \theta)'$$

$$h_S = M(\theta, \sigma^2) - S$$

$$h_{\sigma^2} = (T(\theta; \sigma^2) - \theta)M(\theta, \sigma^2)(T(\theta; \sigma^2) - \theta)' + \sigma^2_u + (a\beta\sigma^2)^2\sigma^2_v - \sigma^2,$$

and where $M(\theta, \sigma^2) = EX_t(\theta, \sigma^2)X_t(\theta, \sigma^2)'$. Locally stable REE under (12)-(15) are associated with stable rest points of the ODE. The Jacobian matrix of this ODE, evaluated at the REE, provides the relevant stability conditions:

$$\begin{pmatrix} \beta(1+c)-1 & \beta k & 0 & 0 & 0 & 0 & -\beta a s_0 \\ 0 & 2\beta c - 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 & 0 & 0 \\ \frac{\partial M(1,2)}{\partial k} & \frac{\partial M(1,2)}{\partial c} & 0 & -1 & 0 & 0 & 0 \\ \frac{\partial M(1,2)}{\partial k} & \frac{\partial M(1,2)}{\partial c} & 0 & 0 & -1 & 0 & 0 \\ \frac{\partial M(2,2)}{\partial k} & \frac{\partial M(2,2)}{\partial c} & 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 2a^2\beta^2\sigma^2_v\sigma^2 - 1 \end{pmatrix}.$$

Local stability requires all eigenvalues to have negative real parts. The Jacobian matrix has eigenvalues $-1 + 2c\beta$, $-1 + \beta + c\beta$, $-1 + 2a^2\beta^2\sigma^2_v\sigma^2$, and repeated values

of -1. The root $-1 + 2a^2\beta^2\sigma_v^2\sigma^2$ corresponds to the derivative of the quadratic $\sigma_u^2 + \sigma_v^2(a\beta\sigma^2)^2 - \sigma^2$ and it is easily verified that this is negative at the lower root $\sigma^2 = \sigma_L^2$. At the fundamentals solution $c = 0$, the other nonzero roots are -1 and $-1 + \beta$. Since $0 < \beta < 1$ all roots of the Jacobian matrix are negative, which implies E-stability (and thus stability under learning). At the RE bubble solution $c = \beta^{-1}$, there is one root equal to one, which implies E-instability.

Proof of Proposition 2 We proceed by first noting that under constant-gain learning $\gamma_{1,t} = \gamma_1 > 0$, $\gamma_{2,t} = \gamma_2 > 0$, it is possible to rewrite the real-time learning algorithms (12)-(14) in the form

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

where $\bar{X}_t' = (1, p_t, p_{t-1}, u_t, v_t)'$. The components of $\mathcal{H}$ are implicitly defined by (12)-(14), with a fixed multiplicative term γ_2/γ_1 incorporated into (14). The superscript γ has been added to the parameter estimates ϕ^γ to emphasize their dependence on the gain $\gamma = \gamma_1$. 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, denoted $\phi^\gamma(\tau)$, given by $\phi^\gamma(\tau) = \phi_t^\gamma$ if $\tau_t^\gamma \leq \tau < \tau_{t+1}^\gamma$, where $\tau_t^\gamma = t\gamma$.

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

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

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).

2

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.

Here $h(\phi)$ is as defined earlier except that the σ^2 component of $h(\phi)$ is multiplied by the fixed ratio γ_2/γ_1. The conditional linearity simplifies verification of these conditions. Proposition 7.5 of Evans and Honkapohja (2001) shows 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 $p_t(\phi) = T(k, c; \sigma^2)X_{t-1} - \beta a \sigma^2 v_t$ 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 $\forall(\theta, S, \sigma^2) \in D$, we have: (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 $T(\theta; \sigma^2)$ on D with $\sigma^2 = \sigma^{2*}$, $S^* = EX_t(\phi^*)X_t(\phi^*)'$, (2) for some $\tilde{\varepsilon} > 0$, $\det(S) \geq \tilde{\varepsilon} > 0$, and (3) $k(1 + c) \geq -y_0$ 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 $\bar{A}, \bar{B}$ 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 use the following result from Evans and Honkapohja (2001):

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

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

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

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

3

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

$$
\begin{aligned}
EU(\tau) &= 0 \\
\frac{dVar(U(\tau))}{d\tau} &= D_\phi h(\tilde{\phi}(\tau, \phi_0))V_u(\tau) + V_u D_\phi h(\tilde{\phi}(\tau, \phi_0))' + \mathcal{R}(\tilde{\phi}(\tau, \phi_0)),
\end{aligned}
$$

where $V_u = Var(U(\tau))$.

Details on Approximating the Mean Dynamics With Endogenous Share Supply. Under learning we continue to have

$$
p_t = \beta \left(y_0 + k_{t-1}(1 + c_{t-1})\right) + \beta c_{t-1}^2 p_{t-1} - \beta a \sigma_{t-1}^2 z_{st},
$$

but when share supply may become endogenous additional care is required to construct the mean dynamics. The condition for exogenous supply, $s_0 \le \Phi p_t$, is satisfied if and only if

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

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

Given $\tilde{\theta} = (k, c; \sigma^2)$, equations (5), (6) and (23) specify $p_t = F(p_{t-1}, v_t; \tilde{\theta})$. For computing mean dynamics the complication is that whether (23) is satisfied, and thus whether (5) or (6) applies, depends on v_t.

Mean dynamics are computed by fixing $\tilde{\theta}$ and $\tilde{S}$ and computing the ODE, where the expectation is taken over v_t and $p_t(\tilde{\theta})$, the p_t process for fixed $\tilde{\theta}$. In general this must be done using the process given by (5), (6) and (23), and for any given $\tilde{\theta}$ 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 $\tilde{\theta}$ space into two regions: in one region the probability is high that (for the given $\tilde{\theta}$) the $p_t(\tilde{\theta})$ process will be given by (5), and in the other region the probability is high that the $p_t(\tilde{\theta})$ process will be given by (6).

For the region defuned by equation (5), $p_t(\tilde{\theta})$ converges to a stationary AR(1) with mean

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

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

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

4

Here we have set $v_t = 0$, and replaced p_{t-1} by its mean under (5). The condition can be rewritten as

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

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

For the region defined by equation (6) the linear approximation of the $p_t(\tilde{\theta})$ 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

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

Here

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

Based on this mean, the condition $s_0 > \Phi p_t$ for (6) (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 perceived law of motion to the actual law of motion 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 the σ^2 component of the ODE:

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

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 the procedure. 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 $M = EX_{t-1}X_{t-1}'$. The text solves V numerically, sets $\gamma_2/\gamma_1 = 2$, and plots the 50% and 95% concentration ellipses.

References

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

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

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

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

Table 1: Results from unlikely sequences

u_0	v_0	Description	Time to $p_t = \begin{cases} 1.5\bar{p} \\ 0.5\bar{p} \end{cases}$	Long simulation
σ_u	σ_v	crash	716	crash $\rightarrow$ steady-state
σ_u	$-\sigma_v$	bubble	146	bubble $\rightarrow$ crash $\rightarrow$steady-state
$-\sigma_u$	σ_v	crash	32	crash $\rightarrow$bubble $\rightarrow$steady-state
$-\sigma_u$	$-\sigma_v$	bubble	92	bubble $\rightarrow$crash $\rightarrow$ steady-state
0	0	bubble	132	bubble $\rightarrow$crash $\rightarrow$ steady-state

Figure 1: Simulated stock price dynamics.

1
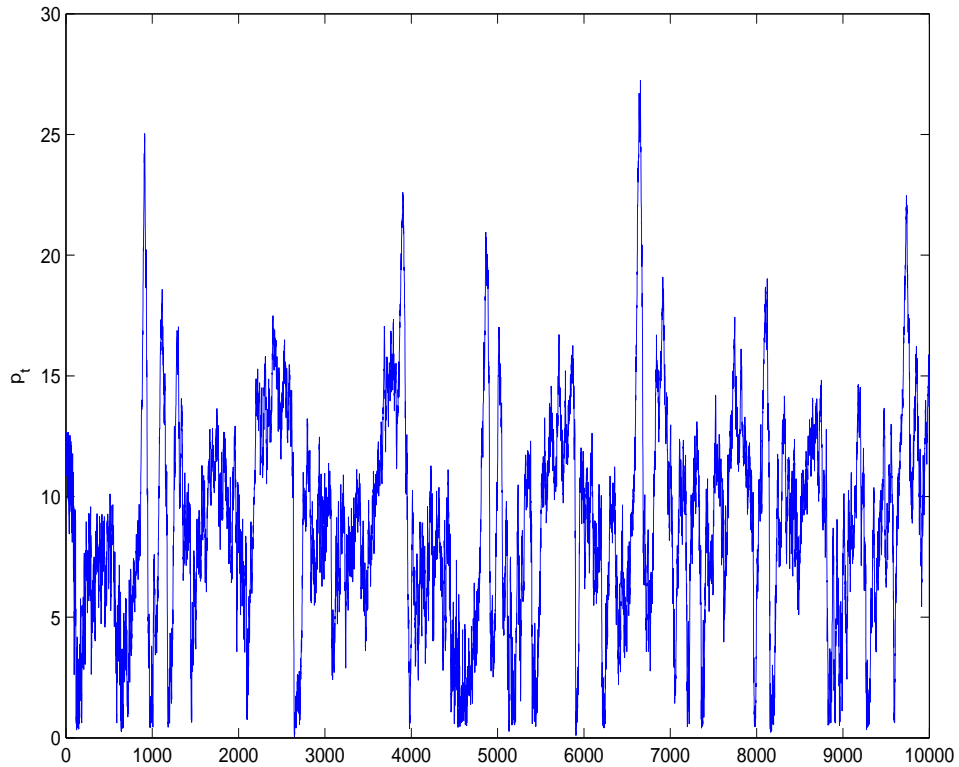
Figure 2: T-map (T_c component).
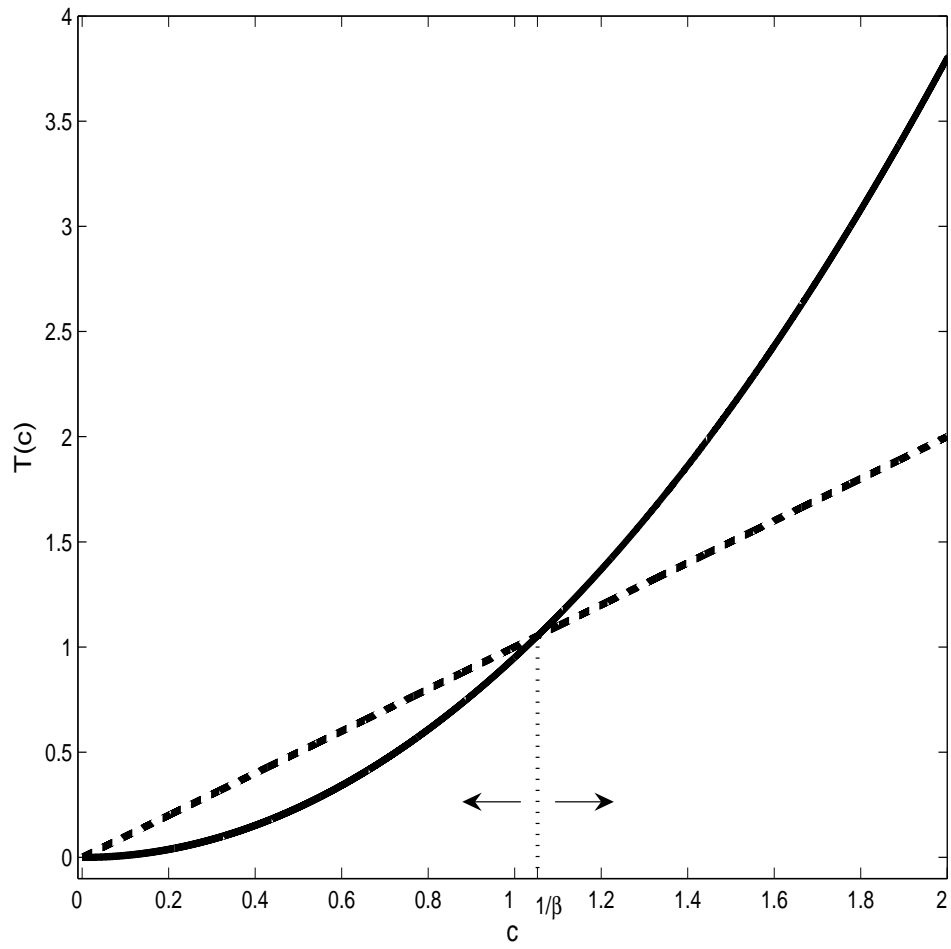
Figure 3: Confidence ellipses around fundamentals REE for constant gain learning version of the model.
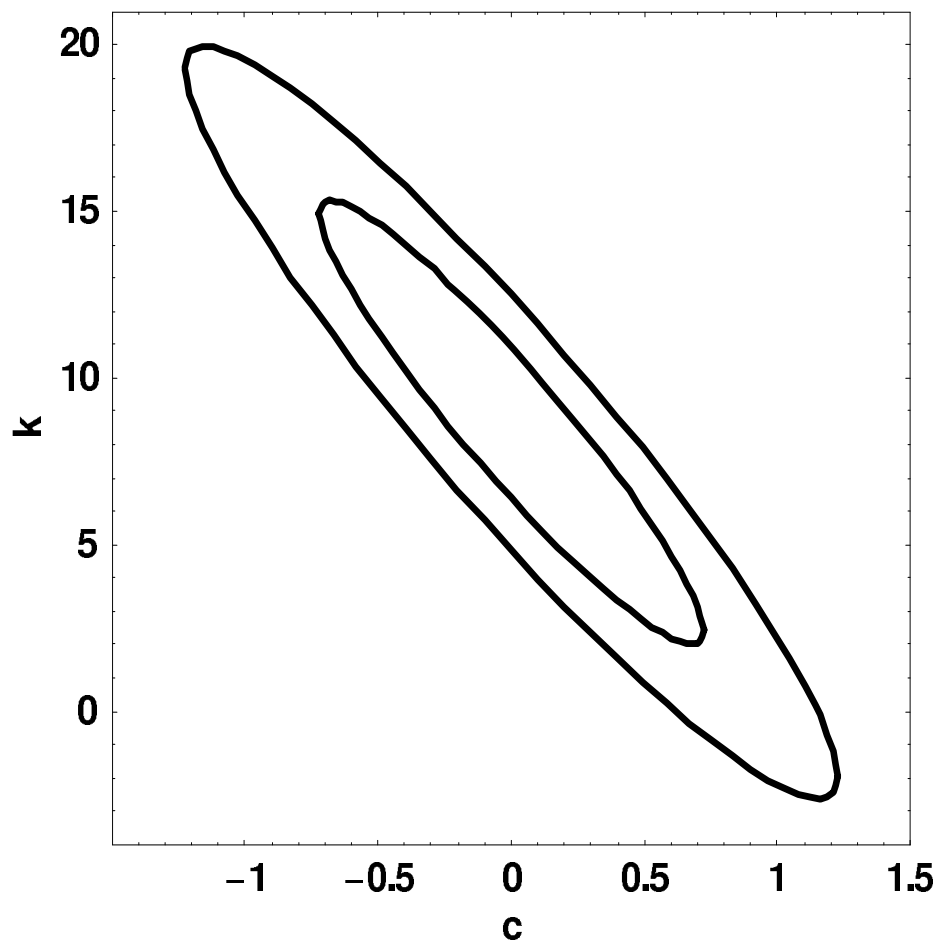
Figure 4: Mean Dynamics, initial values for slope parameter c, k and perceived risk σ^2, drawn from the confidence ellipsoid. Large gain on perceived risk (γ_2).

Mean Dynamics, $\beta = 0.95$, $a = 0.75$, $\sigma_u^2 = 0.9$, $\sigma_v^2 = 0.5$, $y_0 = 1.5$, $s_0 = 1$
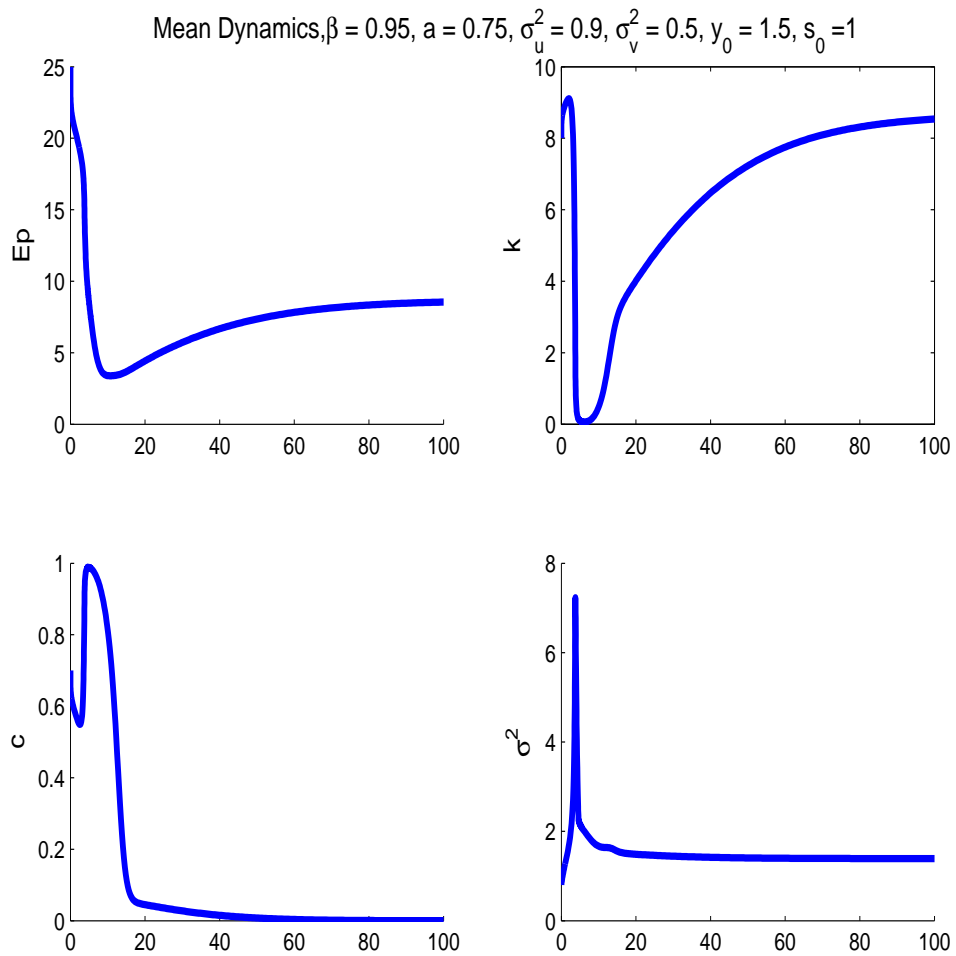
Figure 5: Spectral densities (in logs) for random walk beliefs and the associated actual law of motion.

Figure 6: Path under $(u, v) = (0, 0)$.

v = 0, u = 0 sequence generates bubble
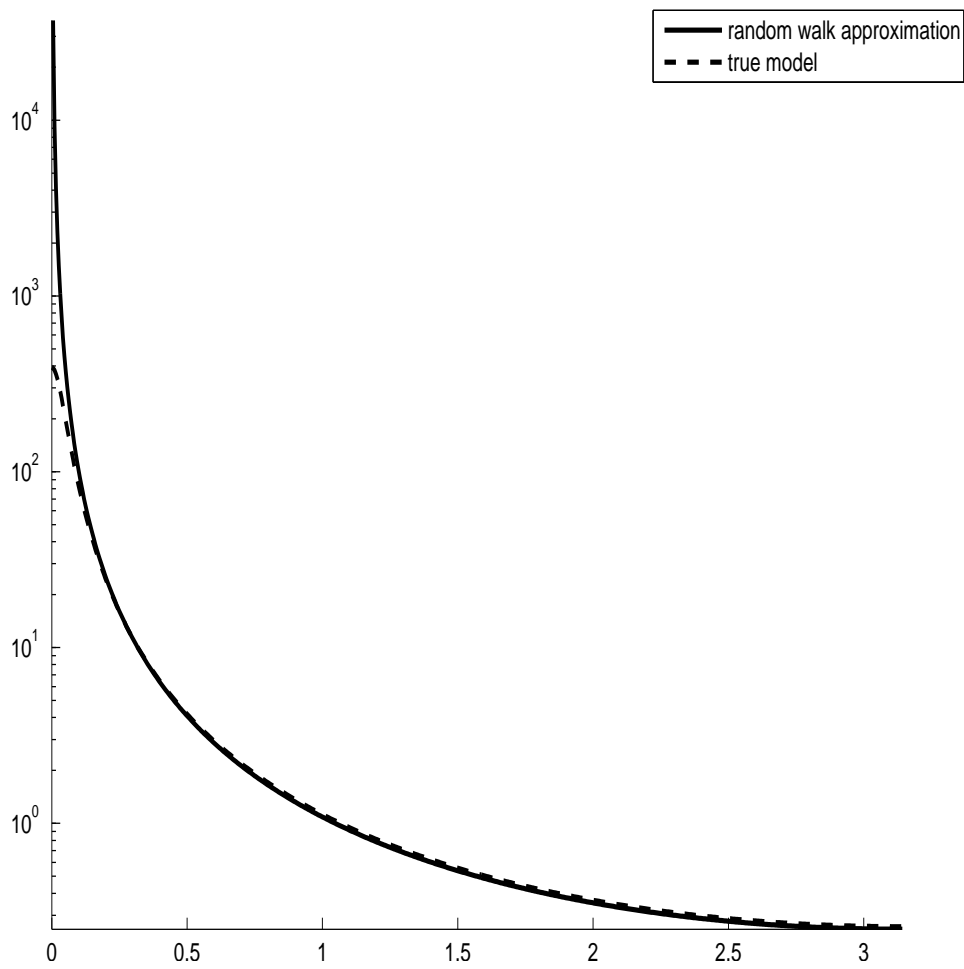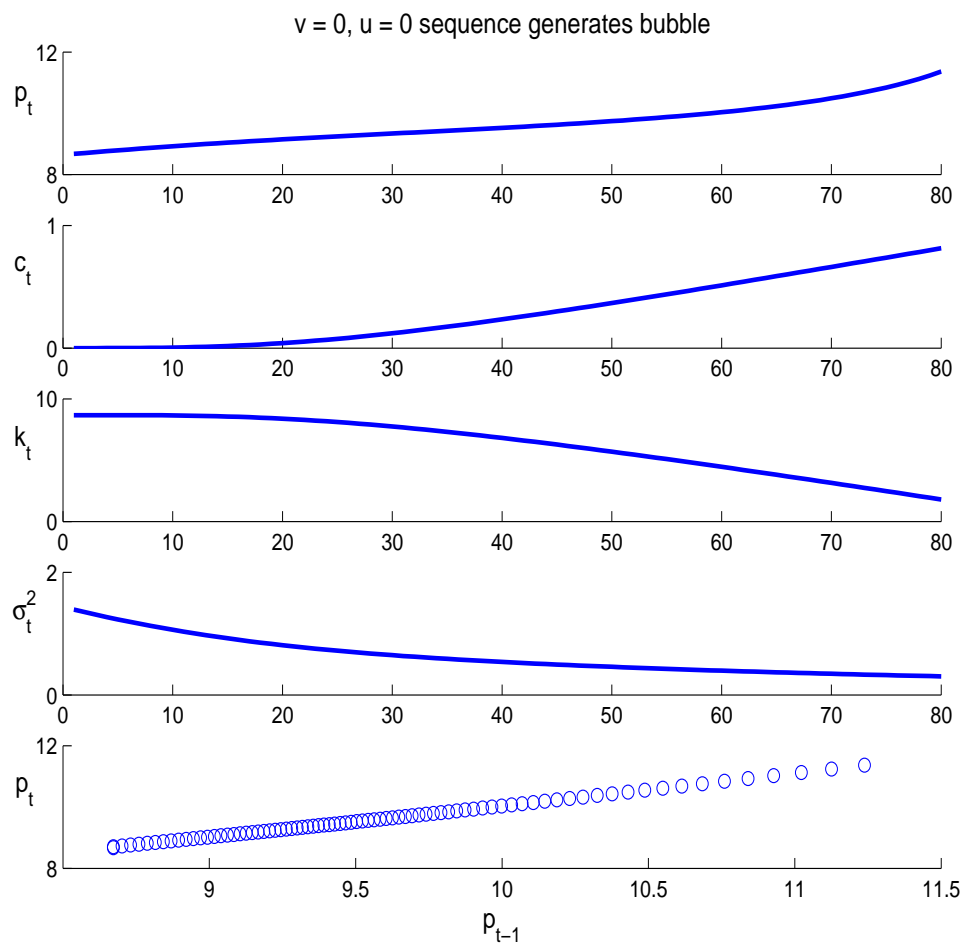
Figure 7: The role of risk in bubbles.
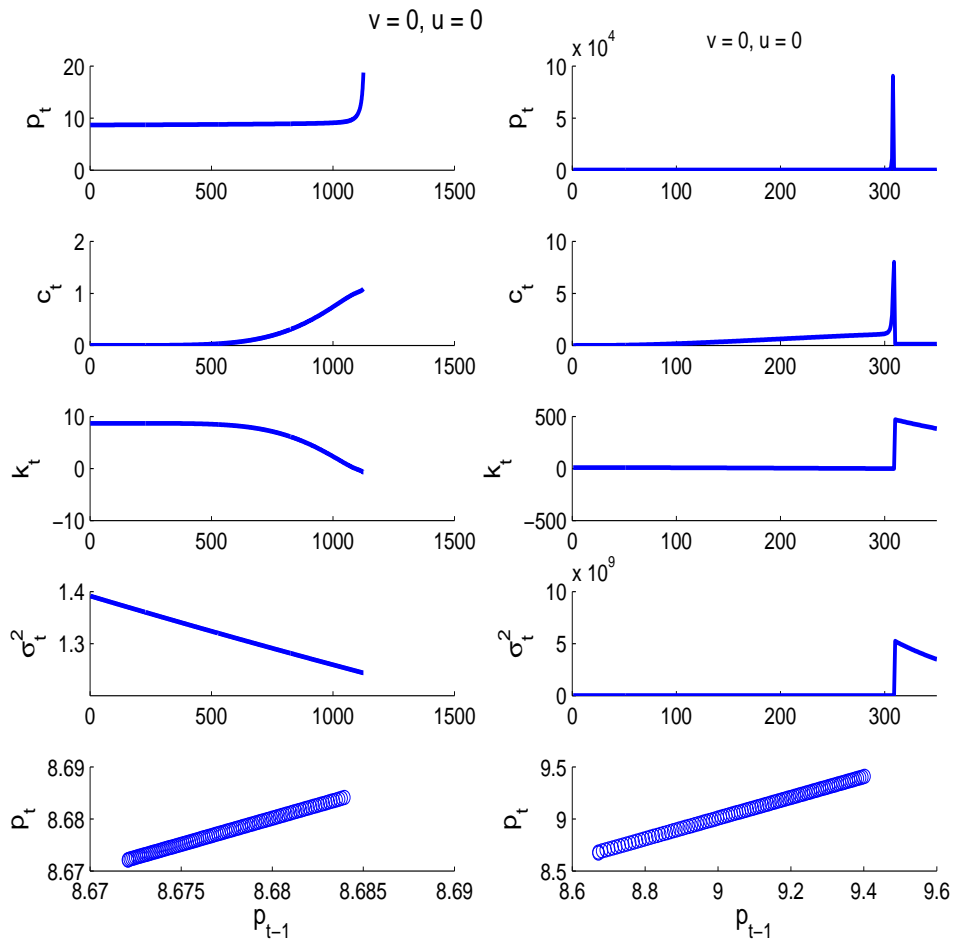
Figure 8: Crash.
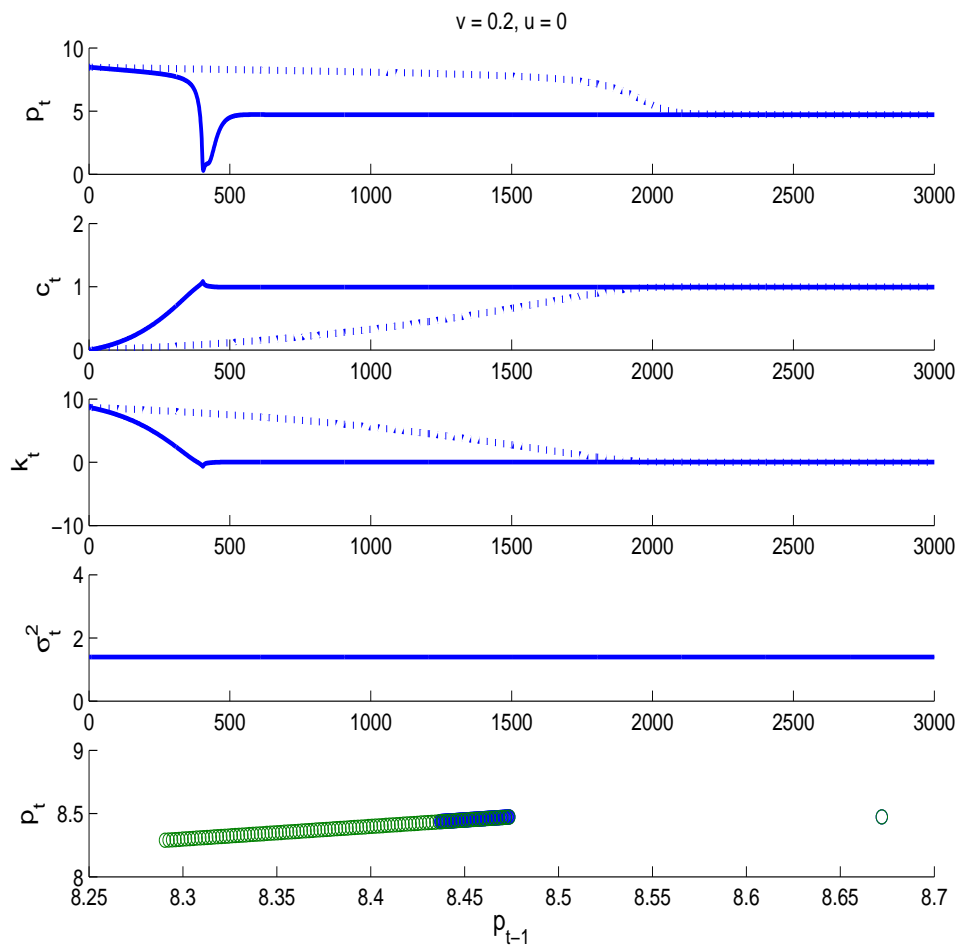
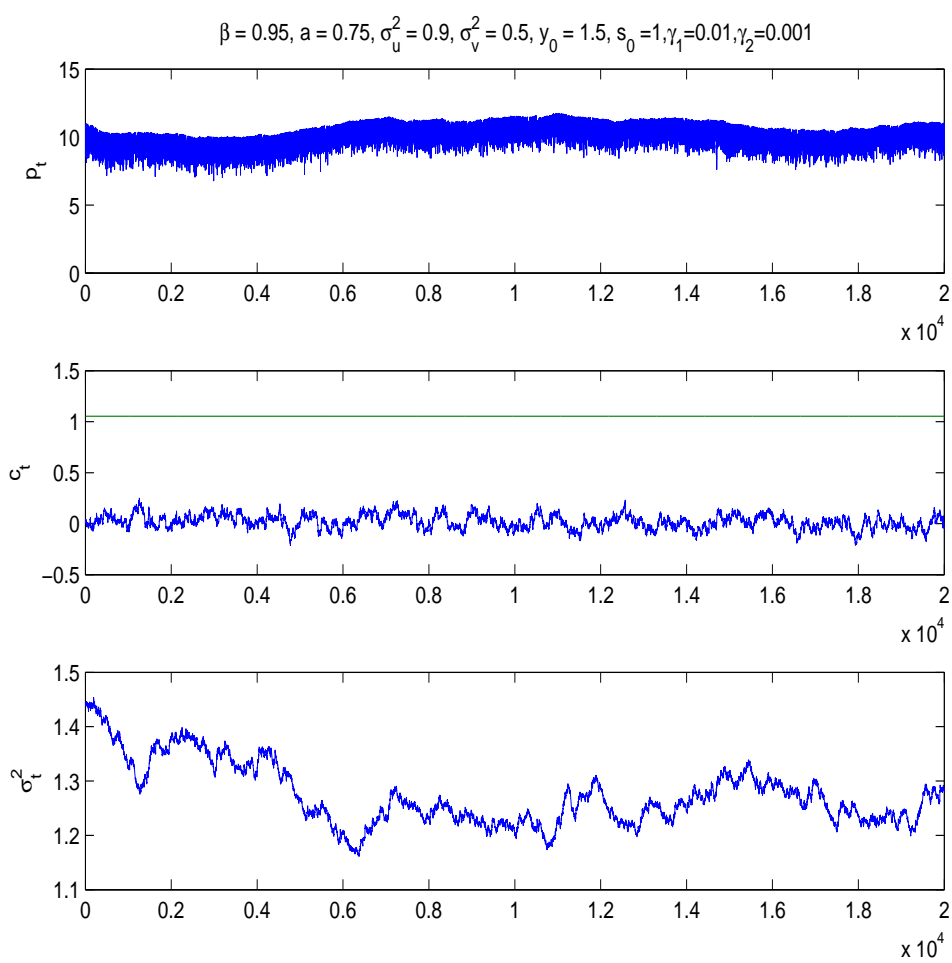Figure 9: Constant gain learning with $\gamma_1 = .01, \gamma_2 = .001$.

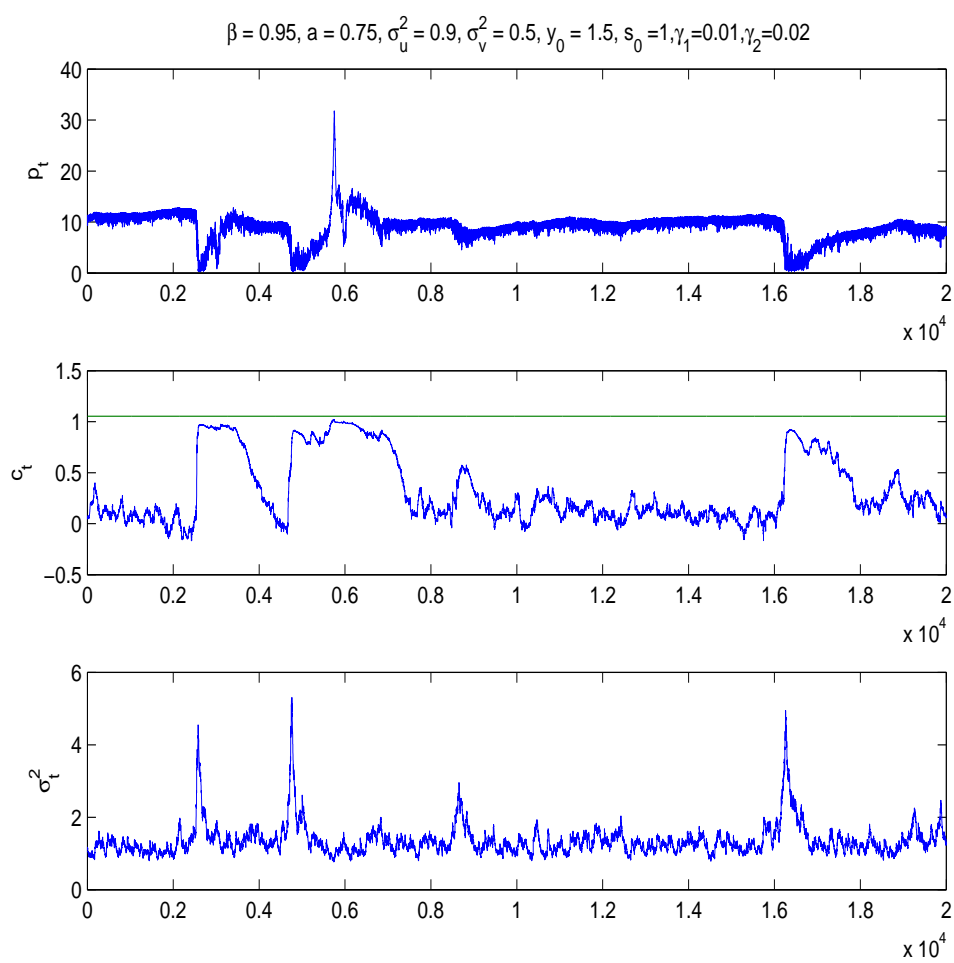Figure 10: Constant gain learning with $\gamma_1 = .01, \gamma_2 = .02$.

Figure 11: Comparison of price and risk dynamics. Constant gain learning with $\gamma_1 = .01, \gamma_2 = .02$.

$\beta = 0.95$, $a = 0.7$, $\sigma_u^2 = 0.9$, $\sigma_v^2 = 0.5$, $y_0 = 1.5$, $s_0 = 1$, $\gamma_1 = 0.01$, $\gamma_2 = 0.02$
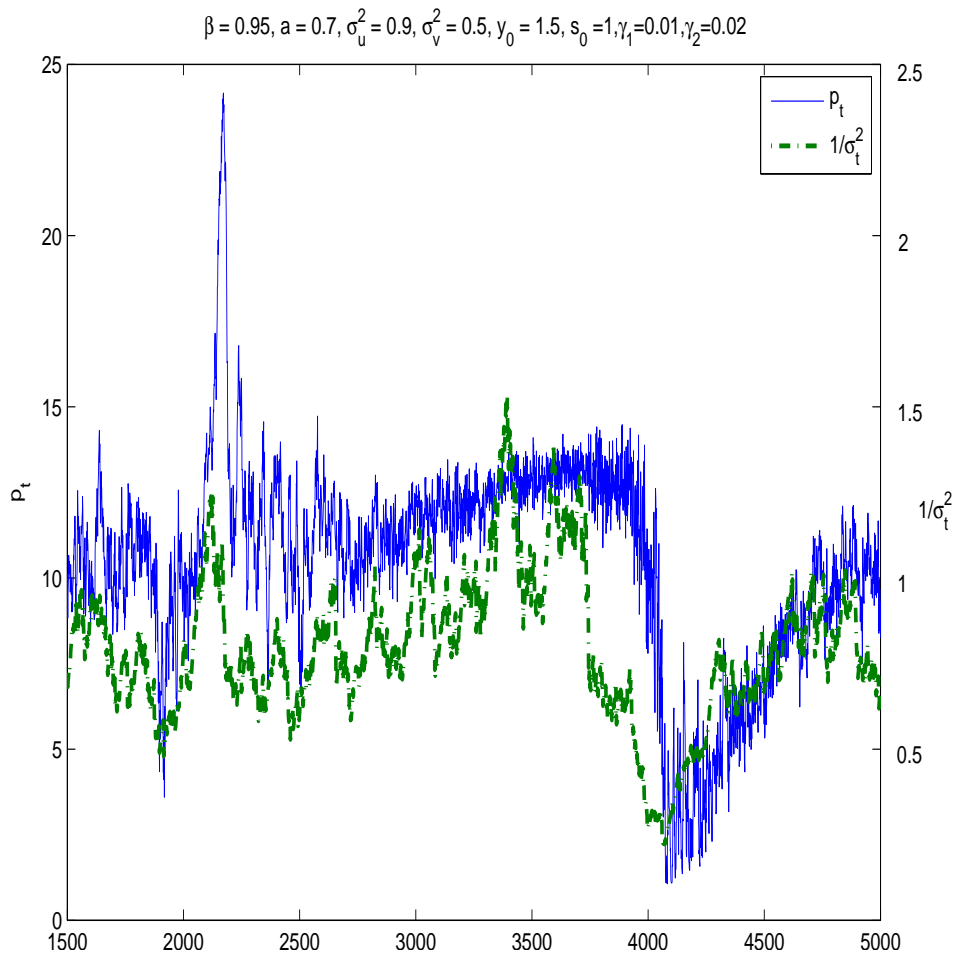
Figure 12: Constant gain learning with $\gamma_1 = .01, \gamma_2 = .04$.