Bubbles, Crashes and Risk*

William A. Branch
University of California, Irvine

George W. Evans
University of Oregon
University of St. Andrews

August 8, 2011

Abstract

In an asset-pricing model, risk-averse agents need to forecast the conditional variance of a stock's return. A near-rational restricted perceptions equilibrium exists in which agents believe prices follow a random walk with a conditional variance that is self-fulfilling. When agents estimate risk in real-time, recurrent bubbles and crashes can arise. These effect are stronger when agents allow for ARCH in excess returns.

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

1 Introduction

An open question in models of asset pricing is the role played by movements in risk premia in generating bubbles and crashes (Greenspan (2005)). Branch and Evans (2011), building on the adaptive learning literature in asset pricing (e.g. Timmermann (1993), Brock and Hommes (1998), Evans and Honkapohja (2001), Lansing (2010) and Adam, Marcet, and Nicolini (2010)), replace rational expectations with an econometric learning rule. A feature of Branch-Evans is that

---

agents estimate risk – the conditional variance of net stock returns – and when this is combined with estimating stock returns, adaptive learning can generate bubbles and crashes. A key mechanism is that adaptive learning introduces serial correlation that would not otherwise exist, which can lead agents' forecasting rule to track this correlation with an approximately self-fulfilling random-walk model.

This paper focuses in more detail on the role of risk in propagating bubbles and crashes, and in particular examines the role of ARCH effects that may be present. We begin by attributing to traders random-walk beliefs about stock prices. We then assume that traders estimate the risk of the stock by formulating one of two econometric models: a simple recursive algorithm and an ARCH model for the conditional variance of stock returns. It is not obvious, a priori, whether an econometric learning rule designed to account for ARCH effects (as in Engle (1982)) will weaken or strengthen the bubble effect. We find that when agents allow for ARCH effects the tendency for learning about risk to generate bubbles and crashes actually appears strengthened.

We proceed as follows. To begin, we demonstrate that there exists a unique, stable restricted perceptions equilibrium under random-walk beliefs. We then show that when agents use a constant gain (or "perpetual learning") algorithm, bubbles and crashes can emerge. ARCH effects can arise from agents' updating of their risk estimates, and these effects are stronger when agents estimate an ARCH model. Bubbles in this setting emerge as sequences of shocks push agents' estimates of risk down – in the case of the ARCH model agents explicitly forecast that this lower risk will persist into the near future. The lower estimates of risk leads to prices being bid up, buoyed by the feedback from random walk beliefs. This cycle persists until agents' estimates of risk eventually increases as the dynamics push estimates back towards the stable equilibrium value. The process as a whole generates recurrent bubbles and crashes.

2   Asset Pricing with Random-Walk Beliefs

We follow De Long, Shleifer, Summers, and Waldmann (1990) and adopt a mean-variance linear asset pricing model with one risky asset that yields dividends $\{y_t\}$ and trades at the price p_t, net of dividends, and a risk-free asset that pays the rate of return $R = \beta^{-1} > 1$. The demand for the risky asset is

$$z_{dt} = \frac{\hat{E}_t(p_{t+1} + y_{t+1}) - \beta^{-1}p_t}{a\sigma_t^2}$$

where $\hat{E}_t$ is the conditional expectations operator based on the agent's subjective probability distribution and σ_t^2 is the corresponding perceived conditional variance of excess returns $p_{t+1} + y_{t+1} - \beta^{-1} p_t$. 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

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

We assume that $y_t = y_0 + \varepsilon_t$ and that $z_{st} = \{\min(s_0, \Phi p_t)\} \cdot V_t$ where ε_t, V_t are uncorrelated *iid* shocks with $EV_t = 1, s_0 > 0$ and $\Phi = s_0/(\bar{p}\xi)$, where $\bar{p}$ is the mean stock price in a fundamentals based equilibrium and $0 < \xi < 1$.[1] Share supply is exogenous except when price falls well below its fundamental value. This assumption provides a flexible price floor in the event of a stock-price crash. We will assume throughout that agents know the true dividend process, so that $\hat{E}_t y_{t+1} = y_0$.

In Branch and Evans (2011), we studied the price dynamics under learning, about both expected future price p_{t+1} and the risk of the stock σ_t^2, and found that under constant-gain learning the model could generate recurring bubbles and crashes. A key to these results was that under learning agents might believe, often for a long stretch of time, that stock prices were following a random walk. Furthermore we found that these beliefs are nearly self-fulfilling. The current paper assumes that agents perceive stock prices to follow a random-walk process, and then use an econometric model to uncover the stock's riskiness in real time.

We also assume that under random-walk beliefs the conditional variance $\sigma_t^2 = \hat{E}_t \left(p_{t+1} - \hat{E}_t p_{t+1} + y_{t+1} - \hat{E}_t y_{t+1} \right)^2$ is believed to be given by

$$\sigma_t^2 = \hat{E}_t(\Delta p_{t+1} + \varepsilon_{t+1})^2.$$

Before turning to the dynamics under learning we first solve for the *restricted perceptions equilibrium* (RPE), in which agents view $\Delta p_{t+1} + \varepsilon_{t+1}$ as an *iid* process and treat σ_t^2 as a constant over time, $\sigma_t^2 = \sigma^2 = E(\Delta p_{t+1} + \varepsilon_{t+1})^2$. We then look for the the self-fulfilling value of σ^2.

When agents hold random-walk beliefs about stock prices, the actual process for stock prices is $p_t = \beta y_0 + \beta p_{t-1} - \beta a \sigma^2 z_{st}$. Here we have assumed that current price p_t is not part of the information set when expectations of p_{t+1} are formed, so that under random-walk beliefs $\hat{E}_t p_{t+1} = \hat{E}_t p_t = p_{t-1}$. The actual price process can be rewritten as

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

---

[1] In simulations, we set $\xi = 0.1$ which implies share supply is exogenous except when price falls below 10% of its mean value.

where v_t is the zero-mean shock defined by $V_t = 1 + v_t/s_0$.[2] Note that for $0 < \beta < 1$ close to $\beta = 1$ random-walk beliefs are self-fulfilling. The σ^2 implied by these beliefs is determined by

$$E(\Delta p_{t+1})^2 + \sigma_\varepsilon^2$$

for the price process (1). Noting that $\Delta p_t = \beta \Delta p_{t-1} - \beta a \sigma^2 \Delta v_t$, it is straightforward to compute the variance of Δp_{t+1}, which implies a mapping from perceived σ^2 to actual σ^2 of the form

$$\sigma^2 = \frac{2(a\beta\sigma^2)^2 \sigma_\nu^2}{1 + \beta} + \sigma_\varepsilon^2 \equiv T(\sigma^2),$$

where σ_ν^2 denotes the variance of v_t. An RPE is then a fixed point of the T-map, i.e. $\sigma^2 = T(\sigma^2)$. It follows that in a RPE

$$\sigma^2 = \frac{1 + \beta \pm \sqrt{(1+\beta)^2 - 8a^2\beta^2(1+\beta)\sigma_\varepsilon^2\sigma_\nu^2}}{4a^2\sigma_\nu^2\beta^2}$$

There are two positive roots. However, Proposition 1 below demonstrates that (only) the smaller root, σ_L^2 is stable under learning.

It is worth remarking that in a fundamentals-based rational expectations equilibrium (REE), the mean stock price is $\beta (y_0 - a\sigma^2 s_0) / (1 - \beta)$, which is identical to the mean stock price in a RPE. However, the equilibrium risk σ^2 is higher in a RPE than in the REE. In addition, as discussed below, learning about risk can give rise to additional stock-price dynamics that are qualitatively very different from the RE.

3   Two Learning Models for Risk

This section develops two theories of how agents might econometrically estimate risk and demonstrates that the RPE is stable under learning.

A simple recursive model for estimating the risk of a stock is given by[3]

$$\sigma_t^2 = \sigma_{t-1}^2 + \gamma_t \left[(p_t - p_{t-1} + \varepsilon_t)^2 - \sigma_{t-1}^2 \right] \tag{2}$$

For the stability results in this section, the gain γ_t is set to $\gamma_t = t^{-1}$ so that (2) corresponds to the recursive least squares estimator for a regression of the

---

[2]We are assuming that Φ and the support of the z_{st} are sufficiently small so that in the RPE we always have $z_{st} = s_0 V_t$.

[3]A similar learning rule was employed by Branch and Evans (2011) and LeBaron (2010).

conditional variance of returns on a constant. In the numerical simulations below we assume a constant gain $\gamma_t = \gamma$, where $0 < \gamma < 1$ is small, so that the recursive algorithm is a form of discounted least squares. Decreasing gains, such as $\gamma_t = t^{-1}$, allow for convergence to the RPE and is thus suitable for studying the local stability of an equilibrium. Constant gains are preferable in environments where agents might be concerned with structural change and also have the advantage of being a time-invariant, or perpetual, learning rule.

An alternative learning algorithm arises when agents perceive the conditional variance of returns to follow an autoregressive conditional heteroskedasticity (ARCH) process. Suppose that agents believe that risk follows the perceived law of motion

$$\sigma_t^2 = \alpha_0 + \alpha_1 \sigma_{t-1}^2 + \eta_t$$

where η_t is a perceived white noise error. Agents estimate their ARCH coefficients (α_0, α_1) by regressing squared excess returns on a constant and m of its own lagged values. Let $\theta' = (\alpha_0, \alpha_1)$, $z_t = p_t - p_{t-1} + \varepsilon_t$, $\bar{z}_t^2 = (1/m) \sum_{j=0}^{m-1} z_{t-j}^2$ and $X_t' = (1, \bar{z}_t^2)$. Then a recursive ARCH estimator is

$$
\begin{aligned}
\theta_t &= \theta_{t-1} + \gamma_t S_{t-1}^{-1} X_{t-1} \left(z_t^2 - \theta_{t-1}' X_{t-1} \right) \\
S_t &= S_{t-1} + \gamma_{t+1} \left(X_t X_t' - S_{t-1} \right)
\end{aligned}
\tag{3}
$$

Here S_t is an estimate of $E X_t X_t'$, the second moment matrix of the regressors. Combining (1) with learning algorithm (2) or (3) leads to a fully specified data-generating process under learning.[4] We have the following stability result for the risk estimator (2).

Proposition 1 *Under the adaptive learning algorithm (2) with exogenous share supply and gains $\gamma_t = t^{-1}$, the restricted perceptions equilibrium $\sigma^2 = \sigma_L^2$ is locally stable under learning.*

Analytic results are unavailable for algorithm (3), however, numerical analysis shows that $\sigma^2 = \sigma_L$ is stable under ARCH learning.

4   Bubbles, Crashes, and Risk

In this section, we consider a constant gain learning version of the model and use numerical simulations to demonstrate the theoretical possibility that stock

---

[4]For simplicity we have specified the ARCH as a parsimonious ARCH(m) model in which the m slope coefficients are constrained to be equal. The results are not greatly sensitive to this specification.

prices can exhibit nearly self-fulfilling ARCH effects and recurrent bubbles and crashes that result from the real-time updating of risk. We choose the following parameterization: $\beta = .98, y_0 = 1.5, s_0 = 1, a = 0.75, \sigma_\varepsilon^2 = 0.95, \sigma_\nu^2 = 0.45.$[5]

Figure 1 plots the results of a 10,000 period simulation for a small constant gain of $\gamma = .0001$. The top panel plots the stock price and the bottom panel plots the real time risk estimates. The right panels are for the ARCH learning model, while the left is the simple recursive algorithm. In a fundamentals REE stock price follows an iid process. With random walk beliefs, stock price is strongly serially correlated. Figure 1 demonstrates that with a small gain there is little variation in real time risk estimates and stock prices do not exhibit bubbles or crashes.

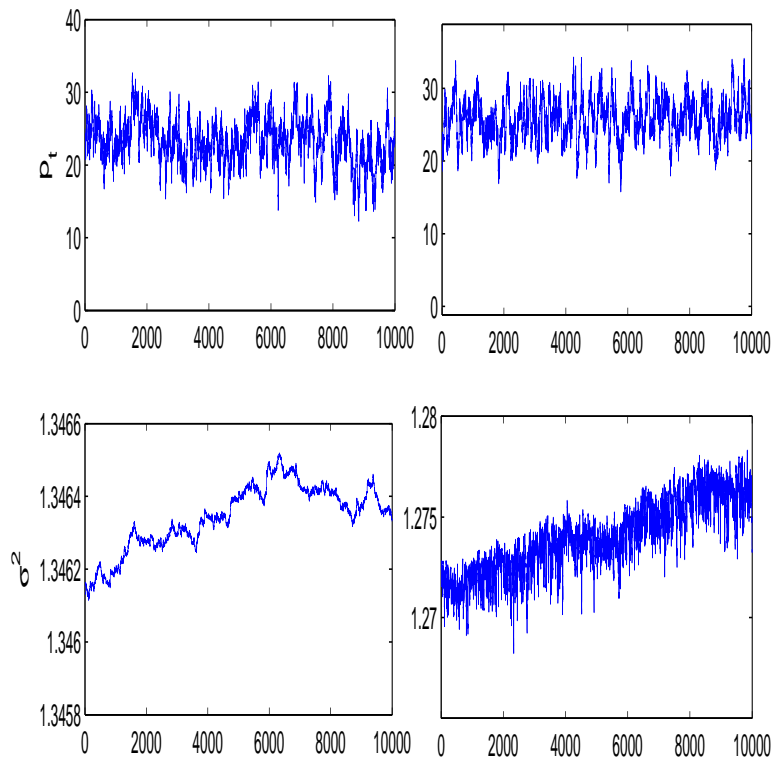Figure 1: Stock price dynamics with a small gain ($\gamma = .0001$).

Figure 2 presents the results from simulations with a larger gain $\gamma = 0.03$. There is significant movement in the estimated risk with periods of declining risk

---

[5]These parameter values are chosen for illustrative purposes; a serious calibration would require a more complicated model than the simple mean-variance framework employed here.

estimates and other periods of elevated risk. Clearly, movements in risk lead to bubbles and crashes in which stock price deviates significantly from its fundamentals value. A decline in perceived risk leads to a bubble: as traders perceive lower risk, their demand increases, leading to higher prices that persist because random walk beliefs interpret these innovations as permanent price increases. Recurring bubbles and crashes arise whether or not agents allow for ARCH effects in their estimation of risk.

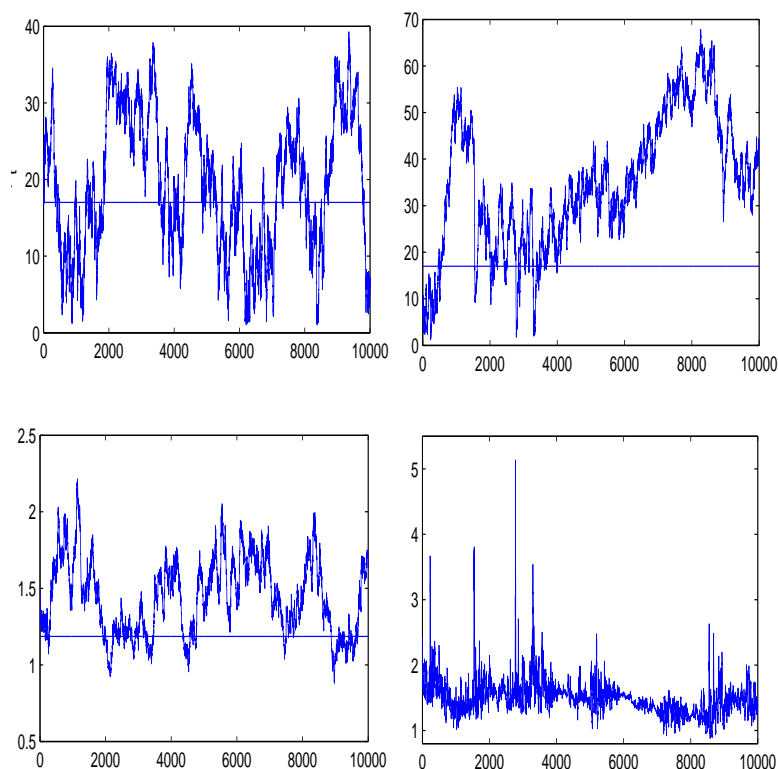Figure 2: Stock price dynamics with gain $\gamma = .03$.

Figure 2 also demonstrates that constant-gain learning about risk can *generate* ARCH effects. In Figure 1, with $\gamma = .0001$, real-time risk estimates were near their equilibrium value. In Figure 2, the estimated risk for the ARCH model (SE corner) exhibit ARCH effects as σ_t^2 fluctuates between periods of high and low volatility. Although perhaps not as noticeable for algorithm 1, it is possible to test for ARCH effects by constructing the test statistic in Engle (1982). We found that for $\gamma = .0001$ the test statistic fails to reject the null that squared excess returns are *iid*, while for $\gamma = .03$ we do reject the null of no ARCH. Learning about risk introduces ARCH effects and allowing for ARCH effects is partially self-fulfilling.

7

Furthermore, when agents allow for ARCH effects in their estimates, this reinforces the role played by risk in generating bubbles and crashes, as can be seen in the top panels of Figure 2.

Figure 3 focuses on a bubble episode from a long simulation assuming ARCH learning (3). The bottom panel also includes a plot of the real time least squares ARCH coefficient estimates $\alpha_{0,t}, \alpha_{1,t}$. The RPE value computed earlier corresponds to $\alpha_1 = 0$. Beginning at period 0 there is a sustained decrease in perceived risk. Moreover, the estimates for $\alpha_{1,t}$ attribute persistence to the decline in risk, which leads to further declines in perceived risk, and translates into a bubble as these price innovations feedback through the random walk beliefs of agents. This figure clearly demonstrates how real time learning about risk can generate bubbles in stock prices.

Figure 3: Bubble (ARCH learning) with gain $\gamma = .03$.
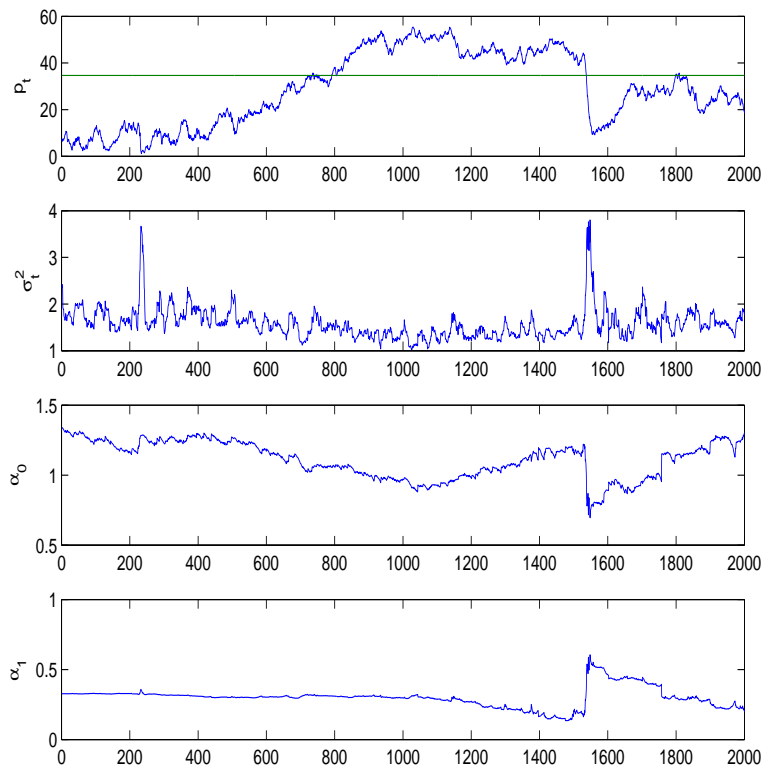
5   Conclusion

This paper adopted a least-squares learning environment to generate recurrent bubbles and crashes. The model consists of risk-averse traders each of whom believe that stock prices follow a random walk. Random walk beliefs were shown by Branch and Evans (2011) to be nearly self-confirming. Risk averse agents need to also forecast the riskiness of stocks – measured as the conditional variance of excess returns – and so they adopt an econometric forecasting model whose parameters are updated with a form of discounted least squares (constant gain learning). The results presented in this paper demonstrate that (1) there exists a unique restricted perceptions equilibrium that is stable under learning, (2) that when agents update their risk estimates in real time with constant gain least squares, recurrent bubbles and crashes can arise, and (3) ARCH effects arise endogenously from agents' learning. These ARCH effects are detectable, and when agents allow this in their algorithms, it strengthens the effect that risk has in generating bubbles and crashes.

Appendix

Proof of Proposition 1.

Let $z_t = p_t - p_{t-1} + \varepsilon_t$. Let $X_t = (1, z_t^2)'$. Then with exogenous share supply we have

$$\sigma_t^2 = \sigma_{t-1}^2 + t^{-1}\left(z_t^2 - \sigma_{t-1}^2\right)$$

for algorithm 1,

Define $\phi = \sigma^2$. Using standard techniques, e.g. Evans and Honkapohja (2001), the differential equation governing local stability of the learning dynamics is

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

where $h(\phi) = T(\sigma^2) - \sigma^2$. Locally stable RPE are associated with the stable rest points of this equation. It suffices to check $DT(\sigma_L^2) = \frac{4(a\beta)^2 \sigma_L^2 \sigma_\nu^2}{1+\beta} < 1$, whereas $DT > 1$ at the larger root for σ^2.

References

ADAM, K., A. MARCET, AND J. P. NICOLINI (2010): "Stock Market Volatility and Learning," Working paper.

BRANCH, W. A., AND G. W. EVANS (2011): "Learning about Risk and Return: A Simple Model of Bubbles and Crashes," *American Economic Journal: Macroeconomics*, 3(3), 159–191.

BROCK, W. A., AND C. H. HOMMES (1998): "Heterogeneous Beliefs and Routes to Chaos in a Simple Asset pricing Model," *Journal of Economic Dynamics and Control*, 22, 1235–1274.

DE LONG, B., A. SHLEIFER, L. SUMMERS, AND R. WALDMANN (1990): "Noise Trader Risk in Financial Markets," *Journal of Political Economy*, 98(4), 703–738.

ENGLE, R. F. (1982): "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation," *Econometrica*, 50, 987–1007.

EVANS, G. W., AND S. HONKAPOHJA (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, New Jersey.

GREENSPAN, A. (2005): "Reflections on Central Banking," speech delivered at the Federal Reserve Bank of Kansas City Jackson Hole Symposium August 26, 2005.

LANSING, K. (2010): "Rational and Near-Rational Bubbles without Drift," *Economic Journal*, pp. 1149–1174.

LEBARON, B. D. (2010): "Heterogeneous Gain Learning and Long Swings in Asset Prices," Working paper.

TIMMERMANN, A. G. (1993): "How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices," *Quarterly Journal of Economics*, 108, 1135–1145.