Risk Preference and Stability Under Learning ‡

Christophre Georges
Department of Economics
Hamilton College

Abstract

We consider a simple market environment in which traders with finite memory update forecasting rules at random intervals by OLS. In this context, changes in the perception of market risk can trigger volatility and bubbles. Consequently, higher degrees of risk response among traders can have a destabilizing effect on price dynamics. We consider the interaction of this effect with memory, the speed of learning, and the nature of the forecasting rules.

JEL Classification: D83, D84, E44.
Keywords: Learning, Expectations, Agent-based Modeling

‡ Contact: Christophre Georges, Department of Economics, Hamilton College, Clinton, NY 13323, USA. cgeorges@hamilton.edu, phone 315-859-4472, fax 315-859-4477, http://people.hamilton.edu/christophre-georges/.

1. Introduction

We report on the behavior of artificial agents who base their trading decisions on returns forecasting rules which they fit to recent data by OLS. The rules are overparameterized relative to the fundamentals of the market. Further, memory is limited, and agents face random wait times between OLS updates.

In Georges (2008a, 2008b), we found through simulation of an agent based model with these features that if traders use minimum state variable forecasting rules, price dynamics conform to a noisy version of the fundamental rational expectations equilibrium as predicted by Honkapohja and Mitra (2003). However, when traders fit overparameterized forecast rules to the data, the learning dynamics can become unstable. This instability tends to increase in the rate of learning (the frequency of OLS updating) and decrease in the memory of the traders. For nonlinear overparameterization, the instability can persist even with relatively large memories. Thus, the correspondence between the learning equilibrium and the fundamental rational expectations equilibrium is fragile.

In those papers, traders were assumed to be risk neutral, and respond only to expected returns. In the present paper, we introduce risk preference among the traders. Since changes in these agents' perceptions of risk will alter their trading behavior, their reactions to risk can either dampen or amplify volatility in the market. Indeed, we find in simulations that volatility, and particularly the incidence of speculative bubbles is increasing in the deviation of traders' preferences from risk neutrality, regardless of whether these agents tend to be risk averse or risk loving. Further, for any degree of risk preference, instability continues to be increasing in the rate of learning and the degree of over parameterization of returns forecasting rules, and decreasing in memory. Additionally, the feedback between the level and volatility of prices due to risk preference can lead to the formation of bubbles even when traders use minimum state variable (MSV) forecast rules for returns, which is not the case in the absence of risk preference.

2. Background

That risk preference can propagate volatility and set off bubbles has been noted elsewhere. Branch and Evans (2011) and LeBaron (2012) show that, under constant gain learning about return and volatility, volatility shocks can trigger both persistence in volatility and bubbles and crashes. LeBaron focusses on the heterogeneity of gains in forecast updating and is able to reproduce a wide variety of stylized facts on stock prices. Branch and Evans are able to generate some analytical results when the perceived model for returns is AR(1).

Here we focus on the dependence of unstable (bubble) dynamics on memory, the speed of learning, the degree of sophistication (overparameterization) of forecasting rules, and the strength of risk preference.

The traders considered below have bounded memory and forcast both returns and volatility using recent data on returns that are, themselves, generated collectively by their own trading behavior. The behavior of the model is studied by means of simulation. This exercise falls under the rubric of agent-based computational economics.

3. The Simple Market Environment

As in Georges (2008a, 2008b), there are two assets, a stock that pays a stochastic dividend d_t in each period t, and a bond with a fixed rate of return r in each period. Dividends are given by $d_t = \bar{d} + \varepsilon_t$, where $\bar{d}$ is constant and the ε_t are iid with zero mean and finite variance σ_ε^2. Denote the price of the stock in period t, P_t.

In each period t, each trader i constructs a forecast $F_t^i[P_{t+1} + d_{t+1}]$ of the price plus dividend of the stock in the following period as well as an estimate $F_t^i[\sigma_{t+1}^2]$ of the variance of her forecast error in that period. As in Georges (2008b) we will assume that the market price is determined by an average trader's behavior at any time. However, this behavior is now driven by that trader's forecasts of both return and volatility. Following the convention of assuming mean-variance or CRRA preferences for simplicity, we will assume the market clearing price of the stock in period t satisfies

$$(1) \qquad P_t = \frac{\bar{F}_t[P_{t+1} + d_{t+1}] - \bar{a} \cdot \bar{F}_t[\sigma_{t+1}^2]}{1 + r}$$

where $\bar{F}_t[\cdot]$ is the forecast of a representative agent (to be defined below), and $\bar{a}$ is the risk factor associated with the representative agent.

In Georges (2008a, 2008b), $\bar{a} = 0$, whereas here $\bar{a}$ may be $\lesseqgtr 0$.

4. Rational Expectations Benchmark

In this simple model, there is a unique stationary rational expectations equilibrium: $P_t = P^* \quad \forall t$, where $P^* \equiv \frac{\bar{d} - \bar{a}\sigma_\varepsilon^2}{r}$. Thus, given the dividend process assumed above, any volatility in price represents a deviation from the stationary REE.

It will be useful to define $x_t \equiv P_t + d_t$ and $x^* \equiv P^* + \bar{d} = \frac{(1+r)\cdot\bar{d} - \bar{a}\cdot\sigma_\varepsilon^2}{r}$, and note that the stationary REE can be expressed as $x_t = x^* + \varepsilon_t \quad \forall t$.

5. Forecast Rules

We suppose that all agents are technical traders who forecast returns using forecast rules with a common functional form

$$(2) \quad \mathrm{F}_t^i[x_{t+1}] = a_{0t}^i + a_{1t}^i x_t + a_{2t}^i \cdot x_{t-1} + a_{3t}^i \cdot x_{t-1}^2 + a_{4t}^i \cdot x_{t-2} + a_{5t}^i \cdot x_{t-2}^2 + a_{6t}^i \cdot x_{t-1} \cdot x_{t-2}$$

where a^i_{0t},...,a^i_{6t} are scalars that can vary across agents i and time t. [1] [2] Functional form (2) is arbitrary but simple and follows the spirit of Grandmont's (1998) "uncertainty principle" as well as Gigerenzer and Selten's (2001) "fast and frugal heuristics." It also nests the MSV forecast rule, which is consistent with the stationary REE, as well as simple trend chasing rules such as the AR(1) rule used by Branch and Evans (2011).

We take the average return forecast $\bar{F}[x_{t+1}]$ used in equilibrium condition (1) to be the forecast (2) using the algebraic averages ($\bar{a}_{0t}$, $\bar{a}_{1t}$, ...) of the traders' rule parameters. Note that the stationary REE forecast rule is given by ($\bar{a}_0$, $\bar{a}_1$, $\bar{a}_2$, $\bar{a}_3$, $\bar{a}_4$, $\bar{a}_5$, $\bar{a}_6$) = (x^*,0,0,0,0,0,0).

For volatility forecasts $F[\sigma^2_{t+1}]$, we assume that each trader simply calculates the mean squared forecast error from the previous M periods (where M stands for memory).

$$\text{F}^i_t[\sigma^2_{t+1}] = \sum_{k=1}^{M} \frac{1}{M} \cdot (x_{t-k} - F^i_{t-k-1}[x_{t-k}])^2 \tag{3}$$

6. Return Forecast Rule Updating

At the start of each period t, each agent updates her returns forecast rule (2) parameters (a^i_{jt}) with common probability *pupdate*. If agent i updates her rule in t, she chooses the rule that minimizes the sum of squared forecast errors over the preceding M periods. Thus, agents learn about returns using a finite memory asynchronous least squares learning algorithm. Note that, given *pupdate* < 1, not all agents will update their rules (by OLS) in a given period t. However, each agent who is selected to update her rule at t observes the same returns history and therefore will select the same new rule (i.e., the same parameter values a^i_{jt}).

In each period t, each agent (regardless of when she last updated her returns forecast rule by OLS) forms a new forecast of next period's return and volatility using her current rules (2-3). Note that volatility forecasts are common, whereas returns forecasts are heterogeneous due to the asynchrony of OLS updating.

[1] We condition forecasts on the level of x rather than on deviations of x from its stationary REE expected value as in the previous papers as the latter value now depends on the volatility of returns which tends to be systematically larger under learning than under rational expectations.

[2] Assuming that x_t is not in traders' information sets at time t, traders form iterated forecasts of x_{t+1} by first forecasting the current period's value x_t using the observed values from the preceeding three periods.

7. Simulations

For any given degree of risk preference, we continue to see the three broad patterns observed in Georges (2008a, 2008b). The incidence of explosive bubbles and crashes is generally increasing in the degree of overparameterization of returns forecast rules, decreasing in memory M, and increasing in the rate of OLS updating *pupdate* (i.e., in the rate of learning). We now also observe a fourth general pattern, which is that the incidence of explosive bubbles and crashes is increasing in the deviation of the risk preference parameter a from zero, $|a|$.

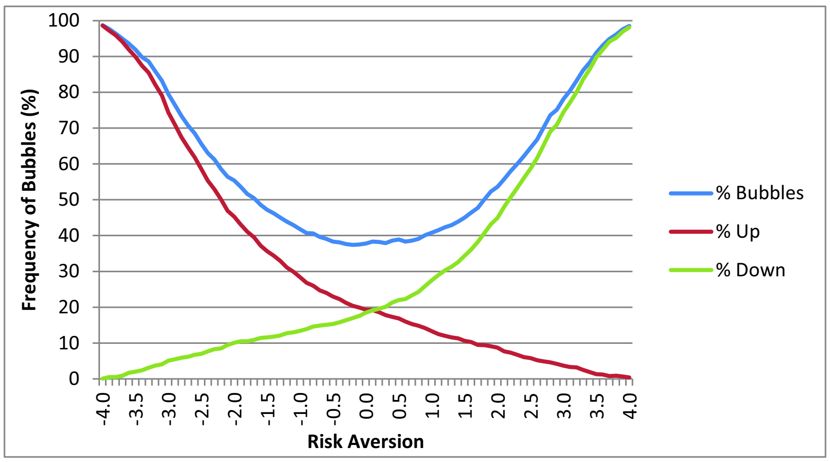

Figure 1: The frequency (in blue) with which bubbles emerge as the risk preference parameter a is varied from -4 to 4. For each of 81 values of a between -4 and 4 (increments of 0.1) we conduct 1000 runs of the model, each with 10,000 trading periods and identical parameter values except for the random seed. The blue curve plots, for each value of a, the percent of runs in which there is a bubble, defined as the price reaching an upper or lower threshold far from the stationary REE. Returns forecast rules are of form (2). Results are shown for memory $M = 200$ and rate of learning *pupdate* $= 0.5$. Also shown separately is the frequency of upward bubbles (in red) and downward bubbles (in green).

We see that the incidence of bubbles is increasing in $|a|$.[3] Further, positive values of a (risk aversion) are more likely to produce downward bubbles, and negative values of a are more likely to produce upward bubbles, while there is roughly the same incidence of upward and downward bubbles at $a = 0$. When agents detect an

[3] In these simulations $\bar{d} = 0.5$, $r = 0.05$, and ε is distributed uniformly on [-0.25,0.25] so that $\sigma_\varepsilon^2 \approx 0.02$. The upper and lower price thresholds for a bubble to be recorded are 20 and 0. These thresholds are far from both the stationary rational expectations equilibria and the normal (non-bubble) price dynamics of the model under simulation. Thus, the recorded bubbles reflect shifts from stable to explosive trajectories.

increase in the variance of returns, they will tend to sell as a result if risk averse, and buy if risk loving, driving price down or up in the two cases respectively. This may further increase perceived volatility and thus trigger further price changes in the same direction. This feedback is in addition to, and interacts with, the feedback through returns learning.

Figures 2 and 3 illustrate that previous results still hold for any given level of risk preference. The incidence of bubbles is decreasing in memory (Figure 2) and increasing in the rate of learning and the degree of overparameterization of the returns forecasting rule (Figure 3).

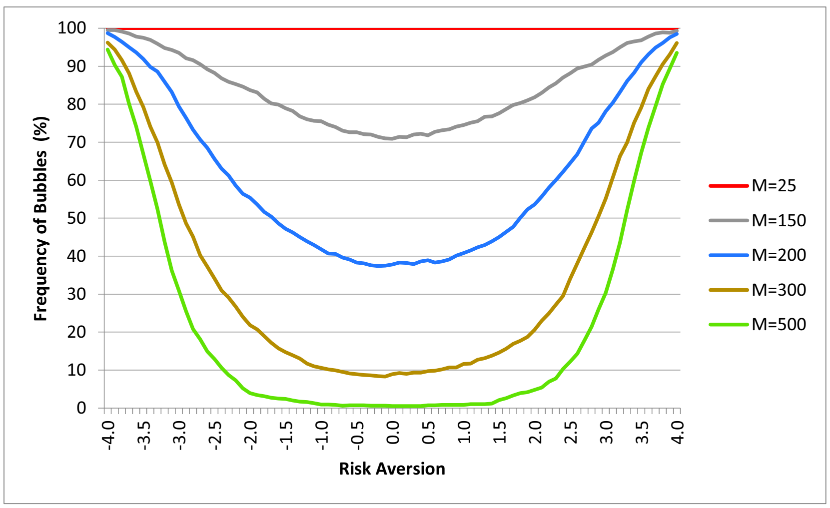

Figure 2: The frequency of bubbles as risk preference a and memory M are varied. $pupdate = 0.5$. The case $M = 200$ (blue) is the same as in Figure 1.

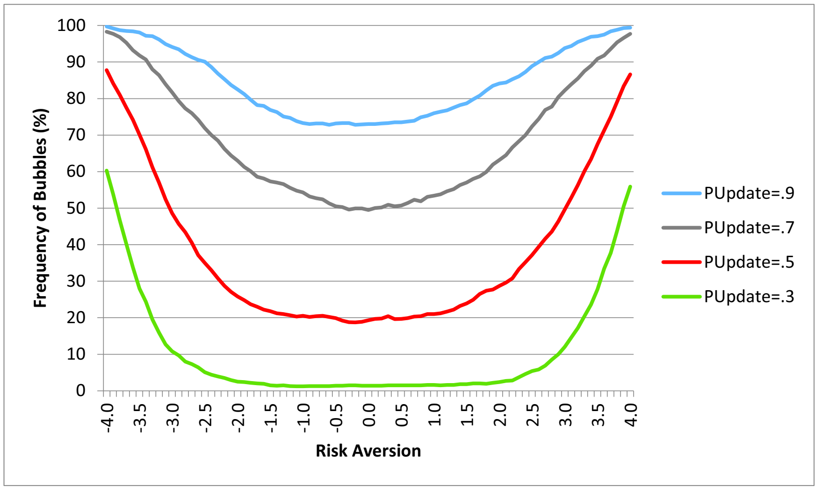

Figure 3: The frequency of bubbles as risk preference a and the rate of learning *pupdate* are varied. Additionally, here agents are using a restricted linear AR(2) version of rule (2) where the a^i_{jt} are set to zero for j=3,4,5, and 6. $M = 25$. With less overparameterization of the forecast rule, the incidence of bubbles is reduced, as is illustrated by the case $pupdate = 0.5$ (red) which corresponds to the case $M = 25$ (red) in Figure 2.

Increasing the degree of overparameterization and reducing M increases the variance of the OLS parameter estimates and thus promotes the incidence of extreme rules. Increasing *pupdate* increases the number of agents who adopt new rules in any period. When a new rule is extreme, it can lead to large price movements which in turn feedback onto the selection of rules, as well as the perception of risk, in subsequent periods.

A fifth further new result is that bubbles can now emerge for positive $|a|$ even in the case of MSV returns forecasting rules. This is not the case when $a = 0$. In the MSV case, not only does risk preference tend to increase the volatility of the price process, but it can also cause enough positive feedback to generate explosive bubbles.[4]

8. Conclusion

The response of traders to risk is shown to provide an additional channel for the propagation of bubbles in the simple framework of Georges (2008a, 2008b) while

[4] Indeed, for sufficiently large $|a|$, bubbles can emerge in the MSV case even when the rate of OLS updating is reduced to zero. Raising the OLS updating rate further amplifies this process and increases the incidence of these bubbles.

preserving the dependence of instability on memory, the rate of learning, and the degree of forecast rule overparameterization shown there.

Acknowledgements:

I thank Yinghan Ding, Rujun Han, James Kotari, Boriana Pratt, Philip Ewing and Kevin Xiao for research assistance.

References:

Branch, William A., and George W. Evans, 2011, Learning about Risk and Return: A Simple Model of Bubbles and Crashes. *American Economic Journal: Macroeconomics* 3. 159–191.

Georges, Christophre, 2008a, Bounded Memory, Overparameterized Forecast Rules, and Instability. *Economics Letters* 98. 129–135.

Georges, Christophre, 2008b, Staggered Updating in an Artificial Financial Market. *Journal of Economic Dynamics and Control* 32. 2809–2825.

Gigerenzer, Gerd and Reinhard Selten, Editors, **Bounded Rationality: The Adaptive Toolbox**, MIT Press, 2001.

Grandmont, Jean-Michel, 1998, Expectations Formation and Stability of Large Socioeconomic Systems, *Econometrica*, **66**:4, 741–781.

Honkapohja, Seppo, and Kaushik Mitra, 2003, Learning with Bounded Memory in Stochastic Models, *Journal of Economic Dynamics and Control*, **27:**8, 1437–57.

LeBaron, Blake, 2012, Heterogeneous Gain Learning and the Dynamics of Asset Prices. *Journal of Economic Behavior and Organization* 83. 424–445.