



ECONOMIC RESEARCH
FEDERAL RESERVE BANK OF ST. LOUIS
WORKING PAPER SERIES

Learning and Excess Volatility

Authors	James Bullard, and John Duffy
Working Paper Number	1998-016A
Citable Link	https://doi.org/10.20955/wp.1998.016
Suggested Citation	Bullard, J., Duffy, J.; Learning and Excess Volatility, Federal Reserve Bank of St. Louis Working Paper 1998-016. URL https://doi.org/10.20955/wp.1998.016

Published In	Macroeconomic Dynamics
---------------------	------------------------

Federal Reserve Bank of St. Louis, Research Division, P.O. Box 442, St. Louis, MO 63166

The views expressed in this paper are those of the author(s) and do not necessarily reflect the views of the Federal Reserve System, the Board of Governors, or the regional Federal Reserve Banks. Federal Reserve Bank of St. Louis Working Papers are preliminary materials circulated to stimulate discussion and critical comment.

LEARNING AND EXCESS VOLATILITY

September 1998

Abstract

We introduce adaptive learning behavior into a general equilibrium lifecycle economy with capital accumulation. Agents form forecasts of the rate of return to capital assets using least squares autoregressions on past data. We show that, in contrast to the perfect foresight dynamics, the dynamical system under learning possesses equilibria characterized by persistent excess volatility in returns to capital. We explore a quantitative case for these learning equilibria. We use an evolutionary search algorithm to calibrate a version of the system under learning and show that this system can generate data that matches some features of the time series data for U.S. stock returns and per capita consumption. We argue that this finding provides support for the hypothesis that the observed excess volatility of asset returns can be explained by changes in investor expectations against a background of relatively small changes in fundamental factors.

JEL Classifications: C62, D83, G12

James Bullard
Research Department
Federal Reserve Bank of St. Louis
411 Locust Street
St. Louis, MO 63102
(314) 444-8576
bullard@stls.frb.org

John Duffy
Department of Economics
University of Pittsburgh
Pittsburgh, PA 15260
(412) 648-1733
jduffy+@pitt.edu

1. INTRODUCTION

1.1. Overview. We present a general equilibrium economy in which the fact that agents are learning can imply persistent volatility in the economy's state variables. The structure of our model is closely related to many models in use in general equilibrium macroeconomics today. We argue that our findings provide a plausible explanation for the observed cross-correlations and levels of volatility in economic data coming from markets where expectations seem to play a large role, such as financial markets. Economists have long argued over whether the data in such markets are consistent with fundamental factors, or whether observed prices are instead consistently deviating from prices based solely upon underlying fundamentals. Our approach provides a way to frame this debate within the context of standard capital theory.

The environment we examine is a general equilibrium lifecycle economy with capital accumulation and exogenous growth. Agents live for many periods, and we use standard specifications for preferences and technology. In this environment, agents face a *multi-step* ahead forecast problem, one that has not often been studied in the learning literature to date. Agents learn by running least squares regressions using data that is endogenously generated by the economy in which they operate: They form forecasts of the future using their least squares rule, and then they take optimal actions given their forecast. The use of a least squares rule is a common choice in the learning literature, but its effects have not been widely studied in models with capital accumulation.

In order to make our system as stark as possible, and to maintain relative simplicity, we eliminate all exogenous uncertainty from the economy we study. The volatility that we isolate is thus entirely due to the effects of learning on the system; in the absence of learning, the rate of return to capital in a stationary equilibrium would be a constant that was entirely pinned down by unchanging economic fundamentals. Thus our volatile return to capital is part of the endogenous fluctuation of the economy under learning—the fluctuations arise because expectations are continually being revised in the face of new data. Volatility can persist because agents' expectations affect actual outcomes, and these outcomes in turn feed back into agents' expectations.

We begin by considering a simple benchmark, a perfect foresight version of the model, and characterize the equilibrium. We then introduce least squares learning, and describe the dynamic system that characterizes the economy under this learning assumption. We show that it is possible to partition the parameter space into two sets—one in which an equilibrium of the model is stable under least squares learning and one in which it is unstable. In parameter regions where instability arises, we show that a neighborhood of the steady state can contain complicated limiting dynamics for the system under learning. These are the excess volatility or *learning equilibria* of our model. We are able to characterize the situations in which these dynamics arise in terms of a vector of parameters that govern specifications of tastes and technology.

We then consider a more realistic version of our model, and explore the possibility of a quantitative case for the learning equilibria that we isolate. We use an evolutionary search

algorithm to find interesting calibrations of the economy under least squares learning. We compare the data generated by these artificial economies to actual data concerning the relationship between consumption and asset returns in the U.S. over the last century. Perhaps the most salient feature of the U.S. data is that the percentage standard deviation of the returns to capital is large relative to the percentage standard deviation of per capita consumption growth. We show that the model with learning is able to generate data which approach the U.S. data on this dimension, whereas the model without learning cannot do so. We also provide a detailed discussion of other facets of the fit between the model and the data.

This research provides some support for the hypothesis that much of the observed volatility in capital asset returns may be due to expectations which are continually being revised, against a backdrop of fundamental factors that are not changing in quantitatively important ways. Our approach provides one method of making this type of argument rigorous and quantifiable, and brings it into contact with a large literature on the nature of asset pricing. While our model is sparse in some obvious ways that enable us to simplify matters considerably (for instance, all assets pay the same rate of return), in many respects the quantitative fit we obtain is impressive and suggests that other persistently volatile phenomena (such as business cycles) might also be characterized in whole or in part by such learning system dynamics. At a methodological level, we have shown how one might go about comparing alternative hypotheses concerning the root causes of asset price volatility to more standard hypotheses, in a way that allows economists to evaluate the relative merits of each explanation.

1.2. Related literature. Arthur, *et al.*, (1997), Brock and Hommes (1997), Bullard (1994), Grandmont (1998), Hommes and Sorger (1998*ab*), and Timmerman (1993, 1996) are among the most recent authors arguing that allowing for adaptive learning behavior might endogenously generate fluctuations. The papers by Arthur, *et al.*, Timmerman, and Brock and Hommes specifically focus on the question that we address in this paper: how the incorporation of adaptive learning behavior might help to explain excess volatility in asset market returns. Arthur *et al.* and Timmerman study learning in the context of the standard, efficient markets “present value” model, where the price of a stock is equal to the expected present discounted value of all future dividends. In these models, the process governing dividends is exogenously given and is therefore unaffected by the fact that agents are learning. Similarly, Brock and Hommes study the implications of learning behavior using a version of the cobweb model with exogenously given demand and supply schedules. We view our analysis as complementary to this earlier work. Our main objective is to discover the implications of learning in the context of a general equilibrium macroeconomic model which is closely related to those used in mainstream macroeconomics, and to explore the implications of our model from a quantitative–theoretic viewpoint.

Our work can be distinguished from “noise trader” explanations of excess volatility (e.g., DeLong, *et al.* (1990)) which posit the coexistence of rational and naive investor types. While the agents in our model are heterogeneous in the sense that agents of

different generations have different planning horizons, all agents' expectations of future rates of return on capital are identical in our framework. By contrast with much of the work on noise trader models, we demonstrate that the boundedly rational agents in our model are unlikely to detect that their learning model is misspecified.

Evans and Honkapohja (1997), Marimon (1997) and Sargent (1993) survey the literature on learning in macroeconomic models. Grandmont (1998) provides a general discussion the stability of rational expectations equilibria under adaptive learning behavior and also addresses the possibility of complicated dynamics under learning in general equilibrium models. Hommes and Sorger (1998*ab*) also study the possibility of model consistent, endogenous fluctuations due to learning. Bullard (1994) has demonstrated the possibility of endogenous fluctuations under the least squares learning dynamic that we consider here for a class of general equilibrium endowment economies. One contribution of this paper is to extend this result to production economies.

A number of researchers have recently used adaptive learning models to help explain empirical macroeconomic phenomena. For instance, Arifovic (1996) proposes an adaptive learning explanation of exchange rate fluctuations, Marcet and Nicolini (1997) suggest that learning may be responsible for recurrent hyperinflationary episodes, Sargent (1998) models the recent history of U.S. inflation and monetary policy as a learning process on the part of the monetary policy authorities, and Arifovic, Bullard and Duffy (1997) build a learning-based explanation of world growth and development patterns. Here, we go a step beyond these qualitative comparisons and directly consider how well a calibrated version of our model fits several features of U.S. data.

There is a large literature that addresses the excess volatility of stock prices relative to changes in fundamentals following Shiller (1981) and LeRoy and Porter (1981). Surveys of this literature are provided in Shiller (1989) and Campbell, Lo, and MacKinlay (1997). West (1988) evaluated much of the literature on stock price volatility tests and concluded that excess volatility is a robust empirical phenomenon that is unlikely to be accounted for by standard present value models. He (along with others, e.g. Shiller (1989)) also dismisses rational bubble explanations as unlikely, and argues that the kind of naive behavior exhibited in "noise trader" models does not play a substantial role in the determination of actual stock prices. West concludes that it might be useful to consider some other models of the determination of excess returns, preferably parametric models, "so that the model potentially could be rejected because of implausible parameter estimates or painfully large test statistics." This paper can be viewed as following up on these suggestions.

A number of papers, for instance Poterba and Summers (1988) among many others, documented that real stock returns have a large forecastable component at longer horizons, even though the predictable component at shorter horizons is relatively small. The equilibria we study have this property, as real returns orbit about a constant mean. Campbell (1991) and Campbell and Ammer (1993) used the forecastability result to decompose stock market volatility into the portions attributable to changing forecasts of stock returns, dividend growth, and real (risk-free) interest rates. Of these, the chang-

ing forecasts of real stock returns were by far the most important factor in explaining stock market volatility. Our model might be viewed as consistent with this evidence on expectations-driven volatility, as our fluctuating learning equilibria involve changes in investor sentiment as the sole driving force.

We organize the paper as follows. The next section describes our general model environment, and presents an argument for the existence of learning equilibria in a simple case. We then proceed in the following section to outline our computational approach to identifying such equilibria in more realistic formulations of the model. Some of the detail behind our approach has been relegated to an appendix. In Section 4 we present our main results and include some remarks on interpretation. The final section offers concluding comments and thoughts on directions for further research.

2. THE ENVIRONMENT

2.1. Preferences and endowments. Time t is discrete and takes on integer values on the real line. At every date t , a new generation of agents is born. These agents live for n periods, where $n \geq 2$ is a positive integer. The size of each new generation of agents grows at a constant gross rate $\psi \geq 1$, where the size of the time $t = 0$ generation is normalized to unity. There is a single, perishable good that is both consumed and used as input into production, which we will call capital, and an unbacked outside asset issued by the government. The representative agent of a generation born at time t seeks to maximize discounted lifetime utility given by

$$U = \sum_{i=0}^{n-1} \delta^i \frac{c_t(t+i)^{1-\rho}}{1-\rho}, \quad (1)$$

where $\delta > 0$, $\rho > 1$ is a preference parameter describing the inverse of the intertemporal elasticity of substitution, and $c_t(t+i)$ denotes the time $t+i$ consumption of the agent born at time t . Agents are endowed with one unit of leisure in every period along with an effective labor productivity coefficient, e_i , where i indexes the period of life. Agents inelastically supply their unit of leisure in every period in exchange for the competitive market wage rate. Their income is then the wage multiplied their labor productivity coefficient. The sequence of lifetime effective labor coefficients, $\{e_i\}_{i=1}^n$, is the same for all agents.

2.2. Production technology. The economy contains a number of perfectly competitive firms that all have access to the same constant returns to scale production technology. Aggregate output, $Y(t)$, is determined according to

$$Y(t) = \lambda^{(t-1)(1-\alpha)} K(t)^\alpha L(t)^{1-\alpha}, \quad (2)$$

where $\lambda \geq 1$ denotes the exogenously given gross rate of labor productivity growth, $\alpha \in (0,1)$ denotes capital's share of output, $K(t)$ denotes the aggregate capital stock at time t , and $L(t)$ denotes the aggregate effective labor supply at date t . This latter

quantity is given by

$$L(t) = \sum_{j=0}^{n-1} \psi^{t-j-1} e_{j+1}. \quad (3)$$

Denoting the ratio of capital to effective labor as $k(t) = K(t)/L(t)$, the marginal products of capital and labor determine the rental rate, $r(t)$, and the wage rate, $w(t)$, as

$$\begin{aligned} r(t) &= \lambda^{(t-1)(1-\alpha)} \alpha k(t)^{\alpha-1}, \\ w(t) &= \lambda^{(t-1)(1-\alpha)} (1-\alpha) k(t)^{\alpha}. \end{aligned} \quad (4)$$

2.3. Government. A government that endures forever plays only one role, which is to supply unbacked liabilities to the economy at a constant growth rate $\theta \geq \lambda\psi$. These liabilities pay no interest. The aggregate, time t stock of this outside asset, denoted $H(t)$, evolves according to:

$$H(t) = \theta H(t-1). \quad (5)$$

The government's real revenue from seigniorage is endogenously determined by $g(t) = [H(t) - H(t-1)]/P(t)$, where $P(t)$ denotes price of the consumption good in terms of the outside asset at time t . Substituting the rule for government liabilities into the expression for $g(t)$, we can rewrite real government revenue as:

$$g(t) = \left(\frac{\theta - 1}{\theta} \right) \frac{H(t)}{P(t)}. \quad (6)$$

It is assumed that government revenue leaves the economy.

Remark on the inclusion of a nominal asset. By including a stock of unbacked government liabilities which grows at a constant rate, we are keeping the model close to the class of economies studied by Bullard (1994) in which learning equilibria were isolated. In this way, we are able to maintain some confidence that the volatile equilibria we seek actually exist in our model with capital accumulation and production—otherwise, we would be left with little guidance as to where and how to find such equilibria. We think this is a sensible way to proceed for the purposes of this paper, even though inclusion of a nominal asset has some clear drawbacks for our quantitative exercise. The unbacked government liabilities can be interpreted as unbacked government debt paying zero nominal return, combined with the monetary base. Because the supply of these liabilities is growing at a constant gross rate θ which is greater than the growth rate of the economy, the price level will rise over time, meaning that the unbacked liabilities will pay a low real rate of return equal to $P(t)/P(t+1)$. Arbitrage will force capital to pay the same real return as these unbacked liabilities, and thus by construction we are going to end up with a counterfactually low mean return to capital. This could be remedied, at the cost of some complication, by inclusion of a simple form of costly financial intermediation for capital in the model, so that the real returns to capital would be higher than the real returns to unbacked government liabilities by a constant sufficient to make the mean return to capital match the data. We could go further, and include a fiat currency reserve

requirement on the intermediaries, effectively splitting the stock of unbacked liabilities into two components, fiat currency paying a low real rate of return, and government debt, paying a real return between those on currency and capital. While we think that our basic results would hold in such an environment, our judgement is that building in such elaborate nominal elements is not the right focus for this paper. Accordingly, we elect to remain with the simple framework and report results which focus on the volatility of real returns to capital relative to the volatility of per capita consumption growth and other related statistics, and to leave explanations for rate of return dominance and other issues to future research.

2.4. Assets and rates of return. Agents can rent capital to firms, borrow or lend in the consumption loan market, and hold government liabilities. Arbitrage ensures that the gross rate of return from holding non-interest bearing government liabilities, $P(t)/P(t+1)$, equals the gross real return from capital or consumption loans, $R(t)$, which in turn equals one plus the rental rate net of depreciation at rate μ . This can be summarized as $P(t)/P(t+1) = R(t) = 1 + r(t+1) - \mu$ for all t .

Since we will be considering the behavior of our model under an adaptive learning assumption where agents lack perfect foresight, we want to draw a sharp distinction between past rates of return that are known at each date t and expected rates of return that must be forecast. We use the notation $R(t-i-1) = P(t-i-1)/P(t-i)$, $i = 0, 1, \dots$ to define past, *realized* gross rates of return. For future *expected* gross rates of return we need to keep track of the date at which these expectations were formed, since when agents are learning forecasts can change from one date to the next. We imagine that agents forecast the future gross *inflation rate* and we use the notation $\beta_{t-i}(t-i+j)$ to denote the commonly held expectation, formed at date $t-i$, of the gross inflation rate at time $t-i+j$, where $i, j = 0, 1, \dots, n-2$. The expected gross inflation rate is the inverse of the expected gross rate of return in this economy.

2.5. The household's problem. The representative agent born at time t faces the following sequence of budget constraints, one for each of the n periods of life:

$$\begin{aligned} c_t(t) &\leq w(t)e_1 - a_t(t), \\ c_t(t+i) &\leq w(t+i)e_{i+1} - a_t(t+i) + a_t(t+i-1)\beta_t(t+i-1)^{-1}, \\ c_t(t+n-1) &\leq w(t+n-1)e_n + a_t(t+n-2)\beta_t(t+n-2)^{-1}. \end{aligned} \quad (7)$$

for $i = 1, 2, \dots, n-2$. Here, $a_t(t+i)$ denotes the asset holdings accumulated by the agent born at time t as of period $t+i$. These budget constraints can be combined into a single lifetime budget constraint which is given by:

$$c_t(t) + \sum_{i=1}^{n-1} c_t(t+i) \prod_{j=0}^{i-1} \beta_t(t+j) \leq w(t)e_1 + \sum_{i=1}^{n-1} w(t+i)e_{i+1} \prod_{j=0}^{i-1} \beta_t(t+j). \quad (8)$$

The household's problem is to maximize (1) subject to (8).

2.6. Perfect foresight equilibria. We first consider the case where households are assumed to have perfect foresight knowledge of rates of return. We will later relax this assumption and introduce learning. The perfect foresight case serves as a useful benchmark for comparison with the results we obtain for the same model under a learning assumption and allows us to characterize the conditions under which a rational expectations equilibrium exists. Under perfect foresight, we do not need to draw a distinction between past, realized gross rates of return and expected future rates of return. For this reason, in this section, we replace the notation of expected gross inflation with the perfect foresight realization, i.e. we set $\beta_{t-i}(t-i+j) = R(t-i+j)^{-1}$ for all i, j . Thus, under perfect foresight, the representative agent's budget constraint may be rewritten as:

$$c_t(t) + \sum_{i=1}^{n-1} c_t(t+i) \prod_{j=0}^{i-1} R(t+j)^{-1} \leq w(t)e_1 + \sum_{i=1}^{n-1} w(t+i)e_{i+1} \prod_{j=0}^{i-1} R(t+j)^{-1}. \quad (9)$$

If we solve the household's problem, we can determine each generation's consumption demand at time t in terms of interest rates and wage rates. Therefore, by writing wage rates as a function of interest rates, we can determine each generation's asset holdings at date t as a function of interest rates alone. Aggregate asset holdings, $A(t)$, is then given by $A(t) = \sum_{i=0}^{n-2} \psi^{t-i-1} a_{t-i}(t)$. This expression is accordingly a complicated function of interest rates dating from $t-n+2$ to $t+n-2$.

The asset market clearing condition is that aggregate asset holdings equals the stock of unbacked liabilities plus the capital stock,

$$A(t) = \frac{H(t)}{P(t)} + K(t+1). \quad (10)$$

Combining (5) and (10), one obtains an equilibrium condition:

$$A(t) - K(t+1) = R(t-1)\theta[A(t-1) - K(t)]. \quad (11)$$

Since both $A(t)$ and $K(t+1)$ can be written as functions of interest rates, this expression is a $2n-3$ order difference equation in $R(t)$. We define a *stationary, competitive, perfect foresight equilibrium* as any stationary sequence of values $\{R(t)\}_{t=-\infty}^{+\infty}$, such that equation (11) holds at every date t . We restrict attention to a class of stationary perfect foresight equilibria for this model in which aggregate asset holdings consist of both capital and holdings of the outside government asset. In steady state, total asset holdings $A(t)$ as well as the capital stock $K(t)$ grow at the gross rate of growth of output, $\lambda\psi$. The existence condition for this steady state equilibrium (which is readily apparent from (11)) is a stationary value of $R = \lambda\psi\theta^{-1}$ such that $A(t) - K(t+1) > 0$ for all t . We study a neighborhood of this steady state equilibrium in the remainder of this paper.

2.7. Learning. In this section we relax the perfect foresight assumption and introduce adaptive learning behavior. The agents in the model under learning form forecasts of the future using a least squares regression, and then take optimal actions given their forecast.

We begin by defining a perceived law of motion for all agents in the model. The time t forecast of the gross inflation rate between time t and time $t + 1$ is denoted by $\beta_t(t)$ and defined by $P(t + 1) = \beta_t(t)P(t)$. We imagine that agents estimate gross inflation by running a first order autoregression on price data available through time $t - 1$. The implied regression coefficient can be recursively updated according to the equation:

$$\beta_{t+1}(t + 1) = \beta_t(t) + \gamma(t) \left[\frac{\theta[A(t - 1) - K(t)]}{A(t) - K(t + 1)} - \beta_t(t) \right], \quad (12)$$

where the gain, $\gamma(t)$, can also be defined recursively as:

$$\gamma(t + 1) = \left\{ \gamma(t)^{-1} \left[\frac{\theta[A(t - 1) - K(t)]}{A(t) - K(t + 1)} \right]^{-2} + 1 \right\}^{-1}. \quad (13)$$

This least squares specification for agent learning behavior is a standard choice in the macroeconomic learning literature.

The agents in this model live for many periods, and so they need to forecast many periods into the future in order to decide how much to consume and save in the present period. The autoregression we have specified implies that $\beta_{t-i}(t - i + j) = \beta_{t-i}(t - i) \forall i, j > 0$, that is, that the agent extrapolates the predicted one-step ahead inflation rate to all future periods for which a forecast is required. This feature of the first-order autoregression means that we can denote the forecast value $\beta_t(t + j)$ more simply by the date at which it is made, namely, we can denote $\beta_t(t + j)$ by $\beta(t) \forall j$. Of course, in the next period, a new piece of information will be available, namely the price level $P(t)$, and so a new forecast will be made that takes account of this additional information; this new forecast will then be extrapolated into the future by all agents who need to forecast one or more periods ahead.

Using this simplified notation, we can write the aggregate capital stock at time t under learning as:

$$K(t) = \lambda^{t-1} \left[\frac{\beta(t-1)^{-1} + \mu - 1}{\alpha} \right]^{\frac{1}{\alpha-1}} \sum_{j=0}^{n-1} \psi^{t-j-1} e_{j+1}. \quad (14)$$

Notice that since $K(t)$ was determined by decisions made at time $t - 1$, $K(t)$ is a function of $\beta(t - 1)$. On the other hand, expected future values of the capital stock *at date* t depend on the current value of $\beta(t)$,

$$K(t + i) = \lambda^{t+i-1} \left[\frac{\beta(t)^{-1} + \mu - 1}{\alpha} \right]^{\frac{1}{\alpha-1}} \sum_{j=0}^{n-1} \psi^{t+i-j-1} e_{j+1}, \quad (15)$$

for all $i > 0$. By the same logic, the wage rate at date t is given by:

$$w(t) = (1 - \alpha) \lambda^{t-1} \left[\frac{\beta(t-1)^{-1} + \mu - 1}{\alpha} \right]^{\frac{\alpha}{\alpha-1}}, \quad (16)$$

and expected future wage rates at date t are given by:

$$w(t + i) = (1 - \alpha) \lambda^{(t+i-1)} \left[\frac{\beta(t)^{-1} + \mu - 1}{\alpha} \right]^{\frac{\alpha}{\alpha-1}}, \quad (17)$$

for all $i > 0$. Accordingly, at time t , we can write all future expected wage rates as a function of next period's expected wage rate, $w(t+1)$:

$$w(t+i) = \lambda^{i-1}w(t+1), \quad (18)$$

for $i \geq 1$.

While all agents form optimal consumption plans at time t on the basis of the inflation forecast $\beta(t)$, when agents lack perfect foresight knowledge, these inflation forecasts will generally be incorrect. Therefore, agents will want to *re-optimize* their consumption decisions for the remainder of their lives at every date, taking into account the new inflation forecast that is available at every date. For this reason, it is important to keep close track of the inflation forecasts that are being used in agent's consumption and savings decisions.

Let us begin by rewriting the consumption decision of the agent born at time t , using the simplified notation of $\beta_t(t+j) = \beta(t)$ for $j = 1, 2, \dots, n-2$. We have:

$$c_t(t) = \frac{w(t)e_1 + \sum_{i=1}^{n-1} w(t+i)e_{i+1}\beta(t)^i}{\sum_{i=1}^n \delta^{\frac{i-1}{\rho}} \beta(t)^{\frac{(i-1)(\rho-1)}{\rho}}}. \quad (19)$$

The asset holdings of this agent may be written as:

$$a_t(t) = w(t)e_1 - \frac{w(t)e_1 + \sum_{i=1}^{n-1} w(t+i)e_{i+1}\beta(t)^i}{\sum_{i=1}^n \delta^{\frac{i-1}{\rho}} \beta(t)^{\frac{(i-1)(\rho-1)}{\rho}}}. \quad (20)$$

Agents who were born prior to period t , in period $t-k$, $k = 1, 2, \dots, n-2$, have additional time t income from existing asset holdings brought over from the previous period. These agents make current consumption and savings decisions on the basis of the current forecast of inflation, $\beta(t)$. The fact that they use this new value for $\beta(t)$ implies that they re-optimize at date t , choosing a new consumption plan for the remaining $n-k$ periods of their lives. For these agents, we have that

$$\begin{aligned} a_{t-k}(t) &= w(t)e_{k+1} + R(t-1)a_{t-k}(t-1) \\ &\quad - \frac{w(t)e_{k+1} + \sum_{i=1}^{n-k-1} w(t+i)e_{k+i}\beta(t)^i + R(t-1)a_{t-k}(t-1)}{\sum_{i=1}^{n-k} \delta^{\frac{i-1}{\rho}} \beta(t)^{\frac{(i-1)(\rho-1)}{\rho}}}. \end{aligned} \quad (21)$$

Define the terms

$$D_k(t) = \sum_{i=1}^{n-k} \delta^{\frac{i-1}{\rho}} \beta(t)^{\frac{(i-1)(\rho-1)}{\rho}} \quad (22)$$

and

$$W_k(t) = \frac{w(t)e_{k+1} + \sum_{i=1}^{n-k-1} w(t+i)e_{k+i+1}\beta(t)^i}{D_k(t)}. \quad (23)$$

The expression for $W_k(t)$ can be written entirely as a function of $\beta(t-1)$, $\beta(t)$ by using our definition for the real wage rates, since only $w(t)$ depends on $\beta(t-1)$:

$$W_k(t) = \frac{(1-\alpha)\lambda^{t-1} \left[\frac{\beta(t-1)^{-1} + \mu - 1}{\alpha} \right]^{\frac{\alpha}{\alpha-1}} e_{k+1}}{D_k(t)} + \frac{\sum_{i=1}^{n-k-1} (1-\alpha)\lambda^{t+i-1} \left[\frac{\beta(t)^{-1} + \mu - 1}{\alpha} \right]^{\frac{\alpha}{\alpha-1}} e_{k+i+1} \beta(t)^i}{D_k(t)}. \quad (24)$$

Then we can write:

$$\begin{aligned} a_t(t) &= w(t)e_1 - W_0(t), \\ a_{t-k}(t) &= w(t)e_{k+1} - W_k(t) + R(t-1) \left[1 - \frac{1}{D_k(t)} \right] a_{t-k}(t-1), \end{aligned} \quad (25)$$

for $k = 1, \dots, n-2$. It follows that

$$a_{t-k}(t-k) = w(t-k)e_1 - W_0(t-k), \quad (26)$$

and

$$a_{t-k}(t-j) = w(t-j)e_{k-j+1} - W_{k-j}(t-j) + R(t-j-1) \left[1 - \frac{1}{D_{k-j}(t-j)} \right] a_{t-k}(t-j-1), \quad (27)$$

for all $j < k$.

Using the above definitions, and noting that $A(t) = \sum_{i=0}^{n-2} \psi^{t-i-1} a_{t-i}(t)$, we deduce an expression for aggregate asset holdings at date t as a function of expectations, β , formed at time t and at dates in the past, and of past, realized rates of return R :

$$\begin{aligned} A(t) &= \sum_{i=0}^{n-2} \psi^{t-i-1} [w(t)e_{i+1} - W_i(t)] \\ &+ R(t-1) \sum_{i=0}^{n-3} \psi^{t-i-2} [w(t-1)e_{i+1} - W_i(t-1)] \left[1 - \frac{1}{D_{i+1}(t)} \right] \\ &+ R(t-1)R(t-2) \sum_{i=0}^{n-4} \psi^{t-i-3} [w(t-2)e_{i+1} - W_i(t-2)] \prod_{j=1}^2 \left[1 - \frac{1}{D_{i+j}(t+j-2)} \right] \\ &+ R(t-1)R(t-2)R(t-3) \sum_{i=0}^{n-5} \psi^{t-i-4} [w(t-3)e_{i+1} - W_i(t-3)] \prod_{j=1}^3 \left[1 - \frac{1}{D_{i+j}(t+j-3)} \right] \\ &+ \dots \\ &+ \prod_{j=1}^{n-2} R(t-j) \psi^{t-n+1} [w(t-n+2)e_1 - W_0(t-n+2)] \prod_{j=1}^{n-2} \left[1 - \frac{1}{D_j(t+j-n+2)} \right]. \end{aligned} \quad (28)$$

Next, we note that the equilibrium condition (11) can be written more generally as:

$$R(t-\ell-1) = \frac{A(t-\ell) - K(t-\ell+1)}{\theta[A(t-\ell-1) - K(t-\ell)]}, \quad (29)$$

for $\ell = 0, 1, \dots, n-3$. Using (29) we can substitute out for $R(t-\ell-1)$ in the above expression for $A(t)$. Collecting terms in $A(t)$ we have a recursive expression for aggregate asset holdings at time t :

$$\begin{aligned}
 A(t) = & \frac{\sum_{i=0}^{n-2} \psi^{t-i-1} [w(t)e_{i+1} - W_i(t)]}{1 - \sum_{i=1}^{n-2} \sum_{j=0}^{n-i-2} \psi^{t-j-i-1} \frac{[w(t-i)e_{j+1} - W_j(t-i)]}{\theta^i [A(t-i) - K(t-i+1)]} \prod_{k=1}^i \left[1 - \frac{1}{D_{j+k}(t+k-i)} \right]} \\
 & - \frac{\sum_{i=1}^{n-2} \sum_{j=0}^{n-i-2} \psi^{t-j-i-1} \left[\frac{K(t+1)[w(t-i)e_{j+1} - W_j(t-i)]}{\theta^i [A(t-i) - K(t-i+1)]} \right] \prod_{k=1}^i \left[1 - \frac{1}{D_{j+k}(t+k-i)} \right]}{1 - \sum_{i=1}^{n-2} \sum_{j=0}^{n-i-2} \psi^{t-j-i-1} \frac{[w(t-i)e_{j+1} - W_j(t-i)]}{\theta^i [A(t-i) - K(t-i+1)]} \prod_{k=1}^i \left[1 - \frac{1}{D_{j+k}(t+k-i)} \right]}. \quad (30)
 \end{aligned}$$

This expression for aggregate asset holding under learning is quite useful, because it depends only on expectations formed at time t and earlier and on past values of aggregate asset holdings. We can therefore combine equation (30) with equations (12) and (13) to define a dynamic system under learning. By substituting appropriately we can create a first order nonlinear system in which $\beta(t+1)$, $g(t+1)$, and $A(t)$ are all functions of past values of these same three variables. We seek to understand the dynamic behavior of this system in a neighborhood of the steady state where $\beta = \theta(\lambda\psi)^{-1}$, $\gamma = 1 - \left(\frac{\theta}{\lambda\psi}\right)^{-2}$, and $A = \bar{A}$ (a complicated expression we do not display here), under the existence condition $A - K > 0$.

2.8. Existence of learning equilibria. The steady state of the system under learning coincides with the steady state of the system under perfect foresight. We now consider the local stability of this steady state under learning. In particular, we show that the parameter space for which the steady state exists may be divided into two regions, one in which the steady state is locally stable under learning, and another in which it is locally unstable under learning. In the parameter regions where instability arises, we show that the system under learning possesses complicated limiting dynamics, which we refer to as the learning equilibria of the model following Bullard (1994) and Grandmont (1998). We can discuss the local stability of the steady state under learning and illustrate the existence of learning equilibria most simply for the $n = 2$ period version of the model; the analysis for higher values of n is similar.

Consider the version of the 2-period model in which the productivity profile, $\{e_1, e_2\} = \{1, 0\}$ and in which five of the seven deep parameters, δ , ρ , λ , ψ and μ are all set equal to 1, leaving the two parameters α and θ free to vary. This much-simplified version of the model merely serves to facilitate our illustration of stability analysis as will become clear below. In this simple case, under least squares learning, the aggregate capital stock and asset holdings are given by:

$$K(t) = [\alpha\beta(t-1)]^{\frac{1}{1-\alpha}}, \quad (31)$$

$$A(t) = \frac{1-\alpha}{2} K(t)^\alpha. \quad (32)$$

The steady state where both inside and outside assets are held is one where $\beta(t) = \theta^{-1}$ for all t , such that $A - K > 0$. Assuming that $\theta \geq 1$, this steady state existence condition may be written as $1 \leq \theta \leq (1 - \alpha)/2\alpha$, with $\alpha \in (0, 1)$. Note that in the limiting case where $\theta = 1$, this condition implies that $\alpha < 1/3$. Let us define the set of free, deep parameters for which the steady state exists in this special case as $\mathcal{E} = \{\theta, \alpha \mid \theta \in (1, (1 - \alpha)/2\alpha), \alpha \in (0, 1)\}$.

The system under least squares learning consists of equations (12–13) in this simple case (because $A(t)$ is not recursive unless $n \geq 3$). Using the above definitions for $K(t)$ and $A(t)$ in (12–13), the two equation learning system can be written entirely in terms of present and past values of β and γ . For stability analysis it will be useful to have a first order representation of this dynamical system. When $n = 2$, we can write the system under least squares learning in the following first order form:

$$\begin{aligned}\beta(t+1) &= f[\beta(t), \beta(t-1), \beta(t-2), \gamma(t)] \\ \beta(t) &= \beta(t) \\ \beta(t-1) &= \beta(t-1) \\ \gamma(t+1) &= g[\beta(t), \beta(t-1), \beta(t-2), \gamma(t)]\end{aligned}\tag{33}$$

Letting $\varphi(t) = [\beta(t+1), \beta(t), \beta(t-1), \gamma(t+1)]$, we can express the dynamical system under learning more compactly as $\varphi(t) = M(\varphi(t-1))$ where M is defined by the right hand side of the system (33). The steady state where both inside and outside assets are held occurs at $\bar{\varphi} = M(\bar{\varphi})$, where $\bar{\varphi} = [\theta, \theta, \theta, 1 - \theta^{-2}]'$. The stability of the system under learning can be assessed by linearizing the system and evaluating the resulting Jacobian matrix at a steady state. Despite the simplicity of the two-period special case that we consider here, the analytic eigenvalues are quite complicated. Therefore, we pursue the following numerical approach to conduct our stability analysis.

Recall that in the special version of the two-period model that we are considering, \mathcal{E} describes the set of free parameters for which the steady state where both inside and outside assets are held exists. We shall therefore consider a grid of values for α and θ that covers this entire set. For each (α, θ) pair in \mathcal{E} , we calculate the maximum of the modulus of the four numerically determined eigenvalues for the linearized system evaluated at the steady state. Local stability of the system under least squares learning requires that all four eigenvalues have modulus less than unity. If the maximum modulus is less than unity, we plot a diamond in Figure 1. Thus area to the southwest in the diagram represents a portion of the region of the parameter space in which the least squares learning system is locally stable, as all four roots of the system have modulus less than unity. The figure indicates that as α increases toward the upper bound of $1/3$, the range of θ values under which the least squares learning system is locally stable steadily shrinks. Similarly, as θ increases beyond the lower bound of 1, the range of α values for which the least squares learning system is locally stable also shrinks.¹ Learning equilibria,

¹In an effort to provide a clear diagram, we have not included the entire feasible parameter set in the figure.




Figure 1: The stability frontier for the simplified, two-period economy. A diamond is plotted if a randomly selected α , θ pair is associated with a set of eigenvalues whose maximum modulus is less than one, creating the gray area in the figure. Thus, when α and θ are sufficiently small, the system will be locally stable under least squares learning. The region where learning equilibria may exist occurs when the largest eigenvalue crosses the unit circle, which is in a neighborhood of the border between the stable and unstable regions, the stability frontier, in the figure.

if they exist, will arise along the border between the stable and unstable regions of the parameter space.² Bullard (1994) has shown that for endowment overlapping generations economies learning equilibria correspond to the existence of a Hopf bifurcation in the learning system dynamics. Figure 1 suggests that similar equilibria will exist here as well, and we use a straightforward numerical simulation strategy to find them in more complicated versions of the model. This strategy works as follows. Suppose the system under learning is initialized at the steady state. If we then perturb the system by adding a small shock to the steady state value of β , we can observe the resulting response of the dynamical system. If the system returns to the steady state then it is locally stable under learning, otherwise the steady state is locally unstable. In the latter case, if the system does not explode—if the dynamics remain bounded—then we have reason to believe that we have found a learning equilibrium. If the forecast errors are stationary and have not diminished to zero, as will be the case throughout this paper, then the learning equilibrium is not a perfect foresight equilibrium.

²In a later section we give some heuristic reasoning on forces we think would drive our systems to this region of the parameter space.

3. COMPUTATIONAL STRATEGY

3.1. Overview. We now turn to the question of whether the learning equilibria discussed in the previous section can be isolated in a more realistic context, where agents live for more than two periods. The multi-period version of the model allows us to interpret each period as a length of time that is conducive to a more realistic calibration of the model as will become clear below. Our strategy is to conduct an evolutionary search over a well-defined parameter space, in an attempt to locate parameter regions in which the implied dynamics under learning are complicated and generate data that match several aspects of the actual data collected on the U.S. economy over the last 100 years.

3.2. Model parameterization. We interpret larger values of n as allowing individuals to update their consumption and savings plans more frequently over their lifetimes. We consider a version of the model where agents live for $n = 11$ periods. Assuming that a typical individual's productive lifetime is approximately 55 years, we can interpret each period in our model as comprising an interval of 5 years. The parameters of our model will be chosen with this interpretation in mind.³ The 11-period model is both computationally feasible and captures the essential insight of modeling learning in a multi-period context, namely that agents frequently reoptimize over their lifetimes, taking into account new information that was unavailable to them when they were younger. Such reoptimization by individual agents is not possible in the standard two period environment.

The lifetime sequence of labor productivity coefficients, $\{e_i\}_{i=1}^{11}$, is assumed to be the same for all agents. The 11-period productivity profile we chose is based on data from Hansen (1993) and is hump-shaped. Hansen reports relative hourly earnings from two different samples for seven different age-sex groupings. We use weighted averages of the age-sex groupings and then fit a polynomial function which provides us with a smooth endowment profile for $\{e_i\}_{i=1}^{11}$. Since we use a model with inelastic labor supply, we need to impose some retirement period on the agents. We do this by setting $e_i = 0$ for the last few periods of life.⁴

Our strategy is to keep n and $\{e_i\}$ fixed, and conduct our search over the remaining parameters of the model, which mainly govern tastes and technology. As suggested by the two period example, changes in any of these parameter values can cause the dynamics of the system under learning to undergo phase changes. In principle, we want to allow these parameters to vary across the entire domain in order to have the best chance of finding a parameter vector that best achieves the goals of our search. However, we also want to search in a reasonable parameter range, in part as a way to speed up the search process. Accordingly, we restrict the range over which we allow variation in each of the parameters of the deep parameter vector. Table 1 provides the parameter ranges that we

³Ideally, we would like to consider a 55-period version of our model so that each period could be interpreted as a single year. However, we have found that our search strategy for the 55-period version of the model is not feasible given our current computational resources.

⁴In some economies we set the productivity profile to zero in the last three periods, while in others we set it to zero during the last two periods. We did not find any qualitative differences in the results that were dependent on this feature of the model.

Table 1. Ranges for model parameters.		
Model parameter	One-year range	Five-year range
δ	0.90 to 1.10	0.59 to 1.61
ρ	1.01 to 10.0	same
α	0.15 to 0.30	same
μ	0.05 to 1.00	0.22 to 1.00
λ	1.0025 to 1.0300	1.0126 to 1.159
ψ	1.001 to 1.020	1.005 to 1.040
θ_{pct}	0.00001 to 0.50	same

Table 1: Ranges for model parameters. Here we list the allowed parameter ranges expressed in annual terms in the second column, and expressed in five year terms in the third column.

adopted for each of the seven model parameters. Our choices for the θ parameter, which governs the growth rate of the outside asset, are based on the following considerations. To ensure that the capital stock is positive, we require that the steady state gross inflation rate $\theta/\lambda\psi < 1/(1 - \mu)$. On the other hand, we also require that the steady state gross inflation rate satisfies $\theta/\lambda\psi > 1$. Since the implied restrictions on θ depend on the choices made for λ , ψ , and μ , we choose a value for θ after choices for these parameters have been made. In particular, for given values of λ , ψ , and μ , we set

$$\theta = \lambda\psi + \theta_{pct} [1/(1 - \mu) - 1] \lambda\psi,$$

where $\theta_{pct} \in (0, 1)$. This formula for θ ensures that the above restrictions always hold. In searching for parameterizations of our model, we chose values for θ_{pct} rather than θ , and then obtained a value for θ using the above formula. Thus in Table 1 we provide the ranges for θ_{pct} rather than for θ . While θ_{pct} can take on values between zero and one, we found by experimenting with our system that lower values tended to be the most relevant, and so for much of our analysis we restricted θ_{pct} to relatively low values in order to speed up our search.

3.3. Data targets. An artificial economy is a tuple $\{\delta, \rho, \alpha, \mu, \lambda, \psi, \theta_{pct}\}$ in our framework, which we sometimes refer to as a *candidate vector*. We wish to choose values for these vectors such that the implied behavior of the dynamic system under learning has aspects that match corresponding aspects of the U.S. data. Before we can assess the properties of the simulated data, however, a candidate vector must meet two necessary conditions, namely (1) the parameterization is such that there exists a steady state where both inside and outside assets are held, and (2) there is persistent volatility in asset returns, that is, the system has achieved some kind of complicated attractor. The first objective was achieved by calculating the values of A and K at the steady state where

Table 2. U.S. data, 1890–1994.	
Statistic	Value
Standard deviation of real stock returns	55.2
Standard deviation of real per capita consumption growth	8.75
First order serial correlation of real per capita consumption growth	−0.24
Contemporaneous correlation between real stock returns and per capita consumption growth	0.24

Table 2: Aspects of the U.S. data on returns to capital and per capita consumption growth, from 1894 to 1994. We use five year, non-overlapping time intervals.

$R = \beta^{-1} = \lambda\psi/\theta$ for each candidate vector of parameter values. A steady state with inside and outside assets exists if $A - K > 0$. The second objective was achieved by checking whether the time path for real returns following a initial displacement from the steady state was asymptotically converging toward the steady state, or exhibiting explosive behavior; in either of these two cases, the objective of persistent volatility was judged to be unsatisfied.

If these two necessary conditions are met, we move to assessing the candidate vector's performance on dynamics. We think the most interesting approach is to try to identify the returns to capital in the model economy with the volatile returns to equity in the U.S. economy. The time series evidence that we use to assess the performance of our model is based on an updated version of the data set used in Grossman and Shiller (1981), which provides annual data on stock prices and dividends as well as per capita consumption for the U.S. from 1890–1997. This is the standard data set used in the macro-finance literature. Using this data, we constructed real stock returns (which include real dividends) and real per capita consumption growth rates at five year, non-overlapping frequencies, covering the period 1890–1994, for a total of 21 observations. The statistics from this data that we will use to assess the performance of our model are given in Table 2.

We chose these particular statistics because they capture several important features of the data that effectively characterize what is commonly called excess volatility. In particular, the standard deviation of real stock returns is more than six times the standard deviation in real per capita consumption growth. Furthermore, the first order serial correlation of real per capita consumption growth is somewhat negative, while the contemporaneous correlation between real stock returns and per capita consumption growth is somewhat positive. The picture painted by these statistics is one in which fundamental factors, typically taken to be represented by real per capita consumption growth, do not vary as much as real stock returns, and there is apparently very little in the way of a

relationship between these two variables.⁵ The four statistics in Table 2 will serve as targets for the 11-period model that we seek to calibrate.

The last of our seven objectives concerns the endogenously determined forecast errors, defined for the contemporaneous case by:

$$e(t) = \beta(t)^{-1} - R(t).$$

The hallmark of learning equilibria is that these forecast errors do not tend to zero as time tends to $\pm\infty$. A reasonable restriction to place on the learning equilibria that we isolate through our parameter search is that these forecast errors are not systematic—random enough so that agents will judge that their linear least squares forecasting model is consistent with the world in which they live. We operationalized this by considering the correlations between $e(t)$ and $e(t - j)$, where $j = 1, 2, \dots, 10$.⁶ If the forecasts are sufficiently random, then forecast errors at 1–10 lags should be uncorrelated with one another. We note that since our learning specification is a first-order autoregression, our objective regarding forecast errors is somewhat more rigorous than one might expect least squares learners to adopt. For each parameterization, we calculated the correlation between $e(t)$ and $e(t - j)$ using the 21-observation sample we drew from the artificial time series generated by the model. We focused on the maximum of these 10 correlation coefficients in absolute value.

Given our two necessary conditions and our five targets, and given the seven parameter value ranges of Table 2, we developed and implemented a genetic algorithm to conduct the search for a parameter vector that could come as close as possible to meeting our objectives and data targets. A genetic algorithm is a population-based, stochastic, directed search algorithm that incorporates basic principles of population genetics.⁷ The algorithm works on a population of “strings.” Each string encodes one candidate solution to some well-specified problem. In our application, strings consist of a seven parameter vector for our model. At the beginning of every “generation,” strings are evaluated according to some fitness criterion. In our application the fitness criterion will be how close the candidate system came to meeting all our objectives and data targets. Following the principle of survival of the fittest, strings with a higher fitness level have a greater probability of advancing to the next generation of candidate solutions. The strings that are selected to be retained in the population of candidate solutions undergo, with some fixed probabilities, naturally occurring genetic operations of crossover and mutation, which serve to advance the search for increasingly higher fitness levels. We chose to use a genetic algorithm to conduct our search for a parameterization of our model because these algorithms are

⁵In our model, there is a sharper definition of fundamental factors, namely preferences, technology, and government policy. We have specified all of these factors to be constant or growing at a constant rate, and not subject to stochastic shocks, so that from this perspective the fundamentals of our model are unchanging over time. Nevertheless, as long as the context is clear, we refer interchangeably to “the fundamentals” of the model economy as either the implied sequence of consumption growth rates or as the unchanging preferences and technology.

⁶Agents only live for 11 periods, so we are considering all of the correlations between forecast errors that occur during agents’ lifetimes.

⁷See, e.g. Michalewicz (1994) for an introduction to genetic algorithms.

known from the artificial intelligence literature to be efficient searchers of large and rugged landscapes, such as the 11-period model we are considering here. Indeed, Holland (1975) has shown that genetic algorithms optimize on the trade-off between exploration of new solutions and exploitation of the best solutions discovered in the past.

A more detailed discussion of this search algorithm is given in the Appendix.

4. MAIN FINDINGS

4.1. Dynamics of artificial economies. The results we report are based on evolutionary searches of the parameter space using our genetic search algorithm. We remark that, despite the use of a rather elaborate methodology, we found this search to be a difficult one, in the sense that separate searches beginning with randomly assigned candidate parameter vectors tended to end up in different portions of the parameter space.⁸ To counteract this effect, we simply searched the parameter space a number of times and here we report several of the interesting economies we found. We think this serves our purpose of illustrating how a learning equilibrium could be behind observed volatility in financial markets, even though it is possible that there may be economies that match the data more closely than those displayed here within our parameter space.

We begin by examining the dynamics of these artificial economies, and then we turn to a discussion of the individual parameter vectors.

We report five illustrative best-of-generation strings—the ones with the best fitness values—from searches we conducted, and for simplicity we refer to these five cases as economies 1, 2, 3, 4 and 5. In Figure 2 we display one view of the attracting sets of these economies. These are the limiting dynamics of systems simulated for a large number of periods following a small shock. We have plotted the deviation of the real return to capital from its mean at time t against the same deviation at time $t - 1$, in order to give a two-dimensional view of the limiting dynamics. It is clear from the figure that Economy 2 has particularly simple dynamics, a three-period cycle. Economies 1 and 5 follow motion on relatively complicated closed curves, while Economies 3 and 4 possess more complicated attracting sets.

In Table 3 we report the time series statistics associated with each of these five artificial economies, as compared to the U.S. data. Economies 1 and 2 are what we call *high volatility economies*, because the returns to capital in these cases are about as volatile as in the U.S. data, in fact exactly so in the case of Economy 2. The other three cases are accordingly what we call *low volatility economies*. If we consider the standard deviation of per capita consumption growth, we see that for Economies 1, 3, and 5, the volatility of this variable is close to the U.S. data, and in most cases substantially less than the volatility of returns to capital. These statistics show that this general equilibrium model under least squares learning can capture an essential feature of the excess volatility phenomenon, namely the much greater variation in real returns to capital compared with underlying fundamentals as captured by per capita consumption growth rates. However, in terms of

⁸In part, this may reflect our inability (because of resource constraints) to let the genetic algorithm search for a sufficiently long period of time. How long such a search might take is unknown, however.

Figure 2

Attracting sets for five artificial economies, viewed in two dimensions.

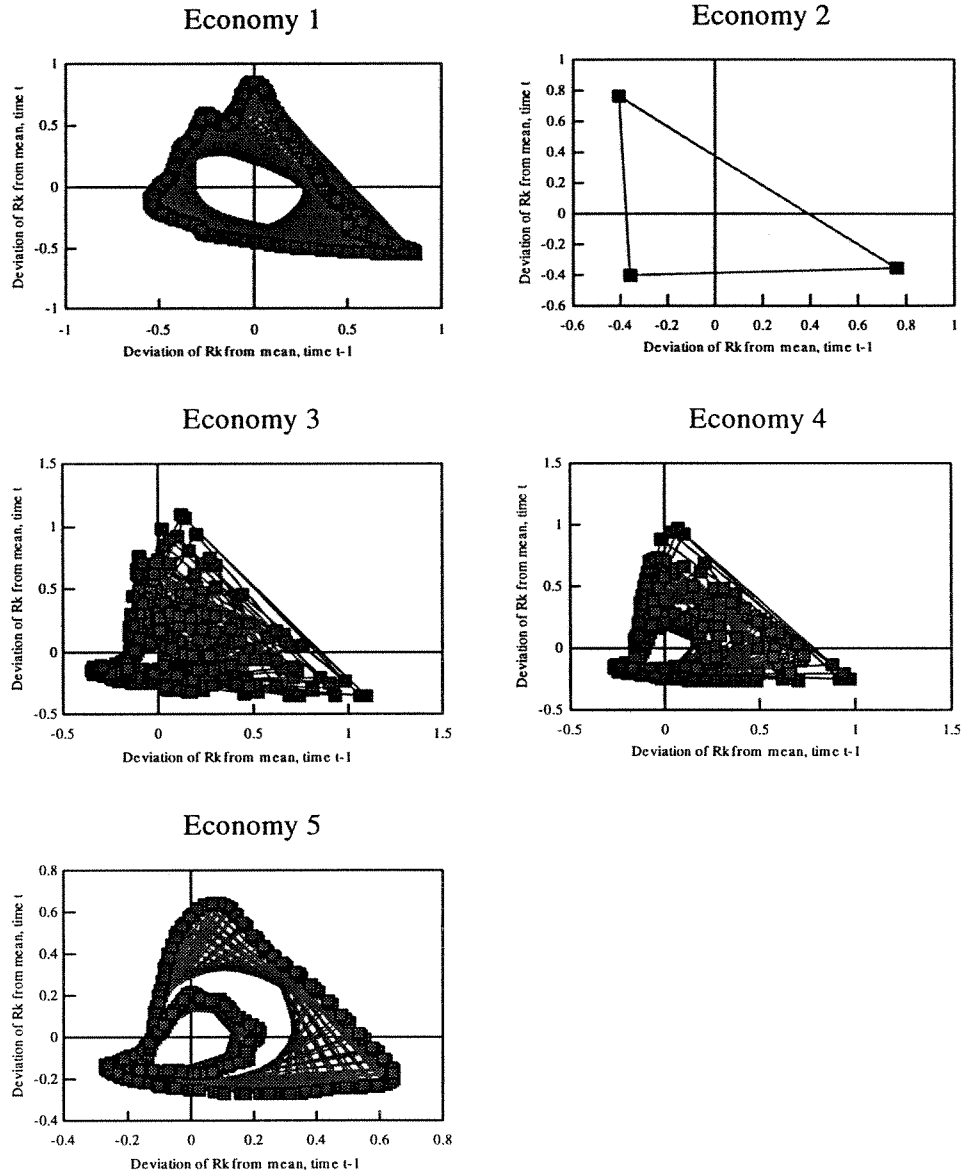


Figure 2: Attracting sets for five artificial economies, viewed in two dimensions. Here we have plotted the the deviation of the real rate of return to capital from its mean at time t against returns at time $t - 1$, for each of the five artificial economies we report.

Table 3. Data statistics associated with artificial economies.						
Statistic	U.S. data	Artificial economy				
		1	2	3	4	5
Standard deviation of returns to capital	55.2	43.8	55.2	24.6	26.5	20.3
Standard deviation of per capita consumption growth	8.75	8.63	21.9	8.17	18.9	8.29
Serial correlation of per capita consumption growth	-0.24	-0.32	-0.53	-0.61	-0.47	-0.29
Contemporaneous correlation between returns to capital and per capita consumption growth	0.24	0.95	0.99	0.43	0.26	.20
Maximum correlation coefficient, forecast errors at 1 to 10 lags	—	0.92	1.00	0.75	0.63	0.68

Table 3: Summary statistics from artificial economies as compared to those from the U.S. data, all in terms of five year time periods. We used a sample size of 21 for the artificial economies, which is the number of data points in the actual data.

relative volatility (the ratio of the standard deviation of returns to that of consumption growth) in these two variables, our model comes up short. None of the artificial economies has a relative volatility of more than about 5, and most are considerably smaller, while the relative volatility in the data is about 6.3. Thus while this model generates excess volatility, it does not generate enough excess volatility to match the data.

The first order serial correlation of per capita consumption growth is negative in the data, and is also negative in each of the five artificial economies. In two cases, Economies 1 and 5, this correlation is relatively close to the data. The contemporaneous correlation between the returns to capital and per capita consumption growth is somewhat positive, .24, in the U.S. data. The high volatility economies we report have far too high a correlation on this dimension to match the data. The low volatility economies do better, with Economy 4 coming particularly close to target.

Our volatile learning equilibria are driven by expectational errors. Agents are using first-order autoregressions to predict the future, and they make mistakes that do not vanish asymptotically. We have no data for the U.S. on forecast errors for five year returns over the last 100 years. However, we do have strong priors that forecast errors should not follow any obvious pattern, lest the agents learn from these errors themselves and change their learning rule, which would in turn change the dynamics of the system. How do our artificial economies fare on this dimension? In the last row of Table 3, we report the maximum correlation for forecast errors at 1 to 10 lags. This is but one way to assess the randomness of the forecast errors in our model. For the two high volatility economies, this

Figure 3
Forecast errors

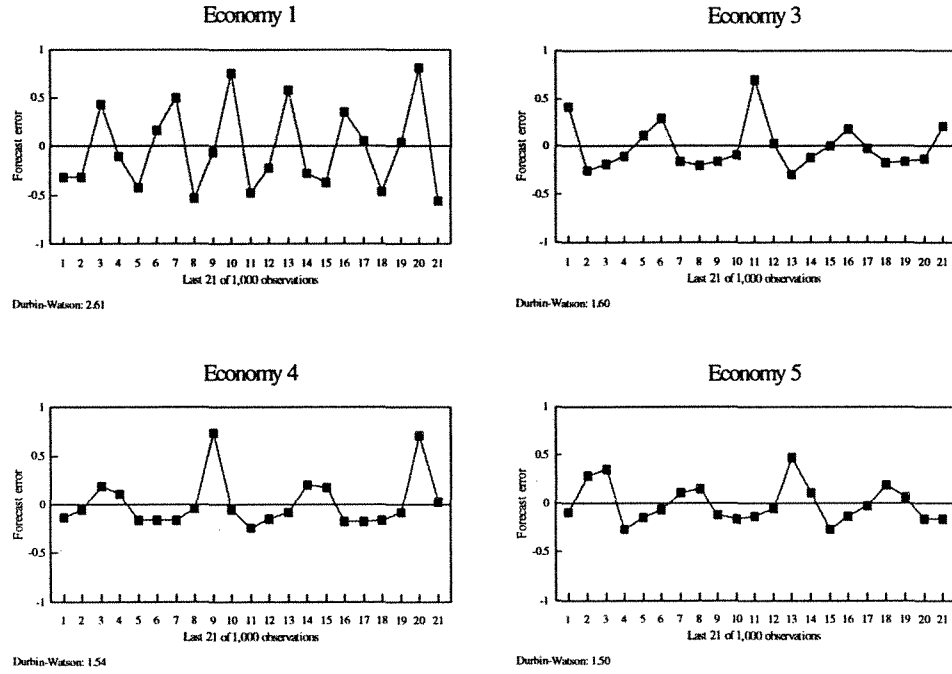


Figure 3: Forecast errors. We plot the last 21 of 1,000 observations for economies 1, 3, 4, and 5. In each case, the Durbin-Watson statistic is given, as if calculated with knowledge of the entire sample.

maximum correlation is very high, and in the case of Economy 2, it is actually 1.0. This is because Economy 2 follows a three-period cycle, so that every third forecast error is exactly the same. The other economies follow more complicated trajectories, with forecast errors that are not perfectly correlated within 10 lags. Our view is that the low volatility economies have forecast errors which are not obviously systematic, enough so that agents could not distinguish them from random errors within 21 observations. We now turn to pursue this point.

Agents who were concerned about possible misspecification of their least squares learning rule might conduct some diagnostic test, such as a Durbin-Watson statistic, for the presence of serially correlated forecast errors.⁹ Using our sample of 21 end-of-run forecast errors, we calculated the Durbin-Watson test statistic for each of our five artificial economies. This statistic d was 2.61, 2.99, 1.60, 1.54, and 1.50, respectively for Economies 1, 2, 3, 4, and 5. A plot of the associated forecast errors for Economies 1, 3, 4 and 5

⁹Bray and Savin (1986) used a Durbin-Watson test as a guide to misspecification in a learning model.

(Economy 2 has perfectly systematic errors and has been omitted) is given in Figure 3. With 21 observations and 1 regressor, the upper and lower bounds at the 5 percent significance level are $d_L = 1.132$ and $d_U = 1.420$.¹⁰ Let's suppose that the agents first test a null hypothesis of no serial correlation against an alternative of positive serial correlation. For economies 1 and 2, the calculated Durbin-Watson statistics are larger than the upper bound and therefore, agents would not reject the null hypothesis of no serial correlation, and would conclude that little evidence of positive serial correlation existed. If they wished to test for negative serial correlation, the corresponding test statistic values would be $4 - d$, or 1.39, 1.01, 2.40, 2.46, and 2.50 for economies 1, 2, 3, 4, and 5, respectively. For Economy 2—the perfectly cyclical economy—the null hypothesis of no serial correlation can be rejected, and for Economy 1, the test is inconclusive. For the three low volatility economies, however, a null hypothesis of no serial correlation again cannot be rejected, and in these cases agents would conclude that there was little evidence of negative serial correlation.

Thus, agents using ordinary least squares regressions might have difficulty detecting that their model was misspecified, even if they had access to data on forecast errors extending back a century or more. Of course, we do not have such data in modern economies. For this reason, it might be more realistic to think of an agent forecasting over his or her lifetime, in our model 11 periods of 5 years each. An agent would presumably recall the forecast errors made over the lifetime, and might change his or her forecast rule late in life based on the errors made early in life. But this would involve an even shorter sample than the 21 observations we have allowed in calculating the Durbin-Watson statistic in Figure 3. We note further that the system we are examining is completely deterministic, an extreme assumption which we have adopted in part for tractability and in part in order to keep the nature of our results clear. If we were to add a reasonable amount of noise to the system, for instance by making output subject to stochastic shocks, the forecast errors generated by our system might appear to be even more complicated to the agents in our model.

4.2. Other characteristics of the artificial economies. We now turn to evaluating other aspects of the artificial economies, based on the best-of-generation parameter vectors at the end of our searches. These parameter vectors, for which growth rates should be interpreted in terms of five-year time periods, are listed in Table 4. Of course, we have already tried to use our seven parameters to meet existence and volatility objectives, as well as five data-based targets, and to ask the model to perform well on a number of additional dimensions is pushing the envelope somewhat—a better approach would be to consider models with more realistic features in order to try to address more aspects of the data. Nevertheless, we think our economies fall short of a completely convincing demonstration of the existence of empirically plausible learning equilibria, mainly because

¹⁰The lower bound for the Durbin-Watson test statistic when there is no intercept term in the regression (as in our case) comes from the tables reported in Farebrother (1980). The upper bound is not affected by the absence of an intercept term.

Table 4. Parameter vectors for artificial economies.							
Economy	δ	ρ	α	μ	λ	ψ	θ_{pct}
1	1.176	2.809	0.1849	0.8940	1.031	1.019	0.01815
2	1.752	3.850	0.1827	0.8630	1.139	1.021	0.02515
3	1.392	1.292	0.1553	0.7507	1.140	1.073	0.07121
4	1.571	1.036	0.1518	0.8033	1.132	1.081	0.08228
5	1.418	1.022	0.1511	0.6616	1.181	1.105	0.08898

Table 4: The parameters associated with five reported artificial economies. These parameter values are expressed appropriately for five-year time periods.

our economies involve capital share parameters which are too low, and depreciation rates which are too high.¹¹ However, we stress that we are only taking a first step in this paper as a means of illustrating the potential of our approach, and that in some ways our parameter vectors are quite reasonable.

Considering Table 4, we begin with the most reasonable aspects and proceed to the less reasonable. In evaluating these parameter vectors, we want to think in terms of modern economies generally, even though we restricted our comparisons to the U.S. data in the previous section. The rate of technological change, λ , and the rate of population growth, ψ , are expressed in gross rates over five-year intervals. Thus the *annual* rate of technological change across the five economies ranges from about one-half of one percent to 3.4 percent. These values are within the range one might expect for industrialized or developing economies.¹² The rate of population growth, on an annual basis, ranges from about four-tenths of one percent to about 2 percent. Again, these values are within the range of observed values for modern economies. The value for θ_{pct} has no direct interpretation, but can be related to the inflation rates observed for these economies. The steady state inflation rates for the five economies range from about 1.1 percent up to about 22 percent, which is consistent with average inflation rates observed in OECD economies during the postwar era. Economy 1 has a steady state inflation rate of 4 percent, the U.S. postwar average.

Regarding preferences, Table 4 reveals that the curvature parameter ρ , which can be interpreted as the coefficient of relative risk aversion, ranged from about unity—logarithmic preferences—up to 3.85. These values are consistent with those often used in the literature. The discount factor δ , was consistently above unity. On an annual basis, the rate of time preference as conventionally measured ranged from -3.2 percent to -8.6 percent.

¹¹When we tried to limit our genetic algorithm search to relatively high capital share, low depreciation cases, we found that it was difficult or impossible to meet our existence check for equilibria where agents hold both inside and outside assets.

¹²We note that, in the model, these rates are constant, but an econometrician evaluating the rate of technological progress would have to estimate them from the volatile data produced by these economies, and that estimate would be subject to some uncertainty. The same would hold true of other parameters.

As is well-known, in overlapping generations models this parameter plays a quite different role from that played in representative agent models, and in particular, there is no requirement that it takes on a positive value. Negative values are consistent with a number of empirical estimates including Hurd (1989), who estimates the difference between the real interest rate facing agents and the rate of time preference as 4.1 percent. The real interest rate is typically taken as the rate on short-term government debt, which averages about 1 percent in the postwar U.S. data, or as that rate after taxes, which averages about 0 percent. This leads to an estimate of the rate of time preference around -3 or -4 percent. Hurd's (1989) alternative estimate of the difference between the rate of time preference and the real interest rate was even larger, 6.1 percent.

The least satisfactory aspect of our economies is that the capital share is too low to be consistent with estimates of the capital share in the U.S. economy, and the depreciation rate is too large. Most capital share calibrations in the real business cycle literature use values of one-third or higher, based principally on inclusion of consumer durables in the measure of the capital stock. Even if one does not include consumer durables in the measured capital stock, capital share is around one-fourth in the postwar U.S. data, higher than in any of our five economies. Estimates of the annual depreciation rate in the U.S. data range from about .04 to .12. In contrast, the five economies listed in Table 4 have annual depreciation rates ranging from a minimum of more than 19 percent up to 36 percent. These are outside the range of U.S. experience.

We leave it as a challenge for future research to develop models with learning equilibria which can perform more satisfactorily on these dimensions.

4.3. A remark on interpretation. So far, we have behaved like econometricians in locating a best fit parameter vector for our model, and interpreting that fit. However, there is an equilibrium question for our analysis, namely, why are these points in the parameter space reasonable? What is it that drives these economies to the stability frontier outlined in Figure 2? Obviously, this type of question is beyond the scope of our analysis, because we have not attempted to model an endogenous process that would keep these economies on that frontier. Instead, we have appealed to the data and argued that the equilibria we describe are the ones that provide the closest match. Nevertheless, we give here a heuristic argument as to why we might observe economies in the volatile region as opposed to, say, in the stable region where the steady state would obtain. We imagine that all parameters are given by nature, except for the government policy parameter θ , which is set by the government to maximize government revenue before setting other explicit taxes. In this scenario, the government moves θ to as large a value as possible. All else equal, larger values of θ will tend to move our systems closer to the stability frontier, and will tend to make observed dynamics more volatile. A government pursuing such a policy would push θ higher but would stop near the stability frontier. At that point, revenue would be maximized on average, but would be volatile. If the government tried to raise still more revenue, the dynamics would collapse (the system would move too far into the unstable region). On the other hand, a lower value of θ would produce lower average revenue. Thus

one might expect certain types of politico-economic equilibria to exist on the stability frontier, and not to exist elsewhere in the parameter space.

5. CONCLUSIONS

Observed levels of volatility in markets where expectations seem to play a large role, such as markets for capital assets, have long been a puzzle for economists. The data we use reveals that the standard deviation in annual returns to equity in the U.S. over the last century or so is around 18 percent. Changes in fundamentals, on the other hand, do not seem to be nearly so pronounced. For instance, the postwar quarterly standard deviation of technological change as measured by the Solow residual is only about 0.7 percent (Cooley and Prescott (1995)) and the standard deviation of annual aggregate per capita consumption growth over the last century is only about 3 percent. Most other processes fundamental to economic value, such as demographics, government policy, or preferences, are usually viewed either as constant or as slowly changing, low-volatility variables. If economic fundamentals are not changing very much from year to year, why are capital asset returns so volatile? Questions like this have led economists to debate whether the data in such markets are consistent with fundamental factors, or whether observed returns are instead consistently deviating from the returns one might expect based on fundamentals alone.

In this paper, we have explored the feasibility of the latter answer to this question, that observed returns are not particularly closely connected with changes in fundamentals, and that they in fact might be driven mostly by changes in investor sentiment. In our framework, the deviations of returns from those suggested by fundamental factors are caused by the fact that agents must learn about the environment in which they operate. This creates a system in which expectational error is a driving force behind economic volatility. In order to illustrate our ideas most vividly (and to keep the analysis relatively simple), we have set up an environment which is very stark: there is no uncertainty whatsoever in the underlying model. In the absence of learning, the steady state return to capital would be completely pinned down by fundamentals, and would exhibit a standard deviation of zero. We have shown that a calibration of our model with learning exists that can deliver time series data that match some of the essential features of the U.S. data on real stock returns and per capita consumption growth. In particular, our model captures a portion of observed excess volatility in the data.

As we have emphasized, our model also has some drawbacks, and our results have to be interpreted as only a first step in this direction. We think further research on quantitative learning equilibria is warranted, and we hope we have provided some ideas about how to pursue such research. Unfortunately, one generally has little guidance from the theoretical literature about which features of quantifiable economies might provide fertile ground for learning equilibria. But on the plus side we were able to extrapolate in the present case from earlier theoretical work in stripped down model frameworks in order to find quantitatively interesting equilibria; perhaps other simple examples can be worked out relatively easily and then be used as a guide for quantitative-theoretic work

like that carried out here.

APPENDIX: THE SEARCH ALGORITHM

The search algorithm proceeds in the following sequence of steps. First, a population of N parameter vectors, or strings, is randomly initialized. We typically set $N = 30, 50$, or 100 , based on suggestions from the literature on genetic algorithms. Each string has seven elements—the seven parameters of our model: $\delta, \rho, \alpha, \mu, \lambda, \psi, \theta$. Denote each of the seven elements of string s_i by ϕ_{ij} , so that ϕ_{i1} is the value for the parameter ρ of string s_i . The initial parameter values for all N strings were drawn with uniform probability from the parameter ranges specified in Table 2.

Second, the values of A and K at the steady state where $R = \lambda\psi/\theta$ were calculated for each parameter string. If $A - K \leq 0$, then the parameter vector is assigned a large number of *penalty points* and no further calculations are made for this string. If, on the other hand, $A - K > 0$, so that the steady state of interest exists, the string is assigned zero penalty points for this step in the algorithm. This step constitutes our existence check. In our search algorithm, high fitness is associated with an absence of penalty points, so a string with a large number of penalty points is not likely to remain long in the population of candidate solutions as will become clear below.

If a string passes the existence check, the next step is to simulate the system for *maxit* iterations using the candidate parameter vector. We found we could get effective simulations by setting the number of iterations as low as 250, but we generally used higher values such as 300, 500, or 1,000. The system is initialized at the steady state and then the initial value $\beta(0)$ is perturbed by a small amount. If the resulting dynamic path for β is determined according to simple criteria to be explosive or convergent, the simulation is stopped (so as to save time) and the candidate string is again assigned a number of penalty points that are inversely related to the number of periods the simulation was continued before being terminated. If the checks for explosive or convergent behavior are satisfied, the model is simulated for *maxit* periods, and earns zero penalty points. This check fulfills our objective of having persistently volatile behavior in asset returns.

If a string passes both the existence and persistent volatility checks, we then take the last 21 observations from the *maxit* simulated observations on rates of return, per capita consumption amounts and forecast errors for this string and we use this sample of 21 observations to construct the statistics that we will then compare with our target values and ranges. In particular, we calculate the standard deviation of returns, the standard deviation of the growth rates of per capita consumption, the serial correlation of per capita consumption growth and the correlation between returns to capital and per capita consumption growth. Finally, we calculate the maximum correlation coefficient of the forecast errors at 1–10 lags. We limit our sample to 21 observations because that is the number of observations from the U.S. data we have on non-overlapping 5-year returns and 5-year growth rates of per capita consumption.

The fitness assigned to a string that passes both the existence and volatility checks is zero. This fitness value may then be further altered according to how well the string

performs with respect to the five statistical targets. Each of the five data statistics gets an equal weight in further altering the fitness of a string. Let s_{it} denote the candidate parameter vector i at generation t , and let θ_{ijt} denote each of the five data statistics ($j = 1, \dots, 5$) that result from simulating the economy with parameter vector s_{it} and analyzing the final 21 observations. Denote the target value for each θ_j by $\hat{\theta}_j$ and the upper and lower bounds by $\bar{\theta}_j$ and $\underline{\theta}_j$. Then the fitness of string s_{it} is given by:

$$F(s_{it}) = \sum_{j=1}^5 p_{ijt},$$

where

$$p_{ijt} = \begin{cases} (\theta_{ijt} - \hat{\theta}_j) / (\underline{\theta}_j - \hat{\theta}_j) & \text{if } \theta_{ijt} < \hat{\theta}_j, \\ (\theta_{ijt} - \hat{\theta}_j) / (\bar{\theta}_j - \hat{\theta}_j) & \text{otherwise.} \end{cases}$$

Given this fitness definition, parameter vectors that come closest to generating data that match the desired targets will have lower fitness values and a vector that delivers an exact fit on all five targets will have a fitness of zero.¹³

Once fitness values have been determined for all N strings in the population, we apply genetic operators which constitute the heart of the genetic algorithm.

First, a selection tournament is held. Two strings are randomly chosen with replacement from the population of strings and their fitness values are compared. The string with the better fitness value wins the tournament and a copy of this string is placed in the next “generation,” $G(t+1)$, of candidate strings. This binary selection tournament is conducted $N-1$ more times so that $G(t+1)$ consists of a population of N strings. The purpose of this selection operation is to direct the search process toward increasingly fit strings. The remaining operations of the genetic algorithm, crossover and mutation, inject the population with new, untried parameter vectors, so as to advance the search for highly fit strings. These operators work on the strings in $G(t+1)$, the strings that have won a selection tournament.

The crossover operation is conducted as follows. The N strings in $G(t+1)$ are considered two at a time, $s_{i,t+1}$, $s_{i+1,t+1}$. With probability p^c , crossover is performed on two vectors; otherwise crossover is not performed. If crossover is performed, we use one of three methods with equal probability. Each of these methods have been shown to have certain strengths in tackling difficult, nonlinear search spaces and therefore we chose to use all three methods to conduct our search. The first method is single-point crossover in which an integer I is chosen uniformly from the set $[1, \dots, 5]$. The two strings are then cut at integer I and the elements of the two strings, ϕ_{ij} and $\phi_{i+1,j}$ to the right of this cut point, i.e. $j > I$, are then swapped. The second method, shuffle crossover, involves 5 draws from a binomial distribution. If the j^{th} draw is a one, then the j^{th} elements of

¹³We also experimented with further penalizing strings that yielded values for θ_j outside the target range $[\underline{\theta}_j, \bar{\theta}_j]$ by adding a quadratic term to the penalty point function, when θ_{ij} values were outside these bounds. However, this modification appeared to make little difference for our results, so we adopted the simpler, linear penalty point mechanism described above.

the two vectors, ϕ_{ij} and $\phi_{i+1,j}$ are swapped, otherwise the j^{th} elements are not swapped. In the third and final method, arithmetic crossover, a random real number $a \in [0, 1]$ is chosen and this number is used to create two new vectors that are linear combinations of the original two vectors: $as_{i,t+1} + (1-a)s_{i+1,t+1}$ and $as_{i+1,t+1} + (1-a)s_{i,t+1}$.

Mutation is performed on the strings in $G(t+1)$ following the application of the crossover operation (i.e. on the recombined strings of $G(t+1)$). The mutation operator makes use of the upper and lower bounds for each of the seven parameter elements of a string, $\bar{\phi}_j$, $\underline{\phi}_j$, that we specified in Table 2, and is applied with probability p^m to every element of every string of $G(t+1)$. If mutation is to be performed on element $\phi_{i,j,t+1}$, then two real numbers, r_1 and r_2 are drawn from $[0, 1]$. The new, mutated value of $\phi_{i,j,t+1}$ is given by:

$$\phi'_{i,j,t+1} = \begin{cases} \phi_{i,j,t+1} + (\bar{\phi}_j - \phi_{i,j,t+1}) \left[1 - r_2^{(1-\frac{1}{T})^b} \right] & \text{if } r_1 > .5, \\ \phi_{i,j,t+1} - (\phi_{i,j,t+1} - \underline{\phi}_j) \left[1 - r_2^{(1-\frac{1}{T})^b} \right] & \text{if } r_1 < .5, \end{cases}$$

where b is a parameter governing the degree to which the mutation operation is non-uniform. This mutation operation is such that the probability of choosing a new parameter element far from the existing value diminishes as $t \rightarrow T$, where T is the maximum number of generations. The purpose of this non-uniform mutation operation is to ensure that, with the passage of time (i.e. following many generations), the genetic algorithm samples more intensively from the neighborhood of existing parameter values, since in the latter stages of a search (close to time T), the parameter vectors should be approaching the optimum.

Following crossover and mutation, the strings of $G(t+1)$ are once again evaluated for their fitness, which involves simulating the economy implied by each $s_{i,t+1} \in G(t+1)$. The genetic algorithm selection tournament was conducted anew followed by another application of the crossover and mutation operators on the winners. This process was repeatedly conducted for some maximum number of generations, which we typically set to 500. The genetic algorithm parameter values we used were $p^c = .95$, $p^m = .20$, and $b = 2$.

REFERENCES

- Arifovic, J. (1996), "The behavior of the exchange rate in the genetic algorithm and experimental economies," *Journal of Political Economy* 104, 510-541.
- Arifovic, J., J. Bullard, and J. Duffy (1997), "The transition from stagnation to growth: an adaptive learning approach," *Journal of Economic Growth* 2: 185-209.
- Arthur, W.B., J. Holland, B. LeBaron, R. Palmer and P. Taylor (1997), "Asset pricing under endogenous expectations in an artificial stock market," in W. B. Arthur, *et al.*, (eds.) *The economy as an evolving complex system II*, Santa Fe Institute Proceedings Vol. 27, Reading, MA: Addison-Wesley, 15-44.

- Bray, M., and N. Savin (1986) "Rational expectations equilibria, learning, and model specification," *Econometrica* 54, 1129-1160.
- Brock, W., and C. Hommes (1997), "A rational route to randomness," *Econometrica* 65, 1059-1095.
- Bullard, J. (1994), "Learning equilibria," *Journal of Economic Theory* 64, 468-485.
- Campbell, J. (1991), "A variance decomposition for stock returns," *Economic Journal* 101, 157-179.
- Campbell, J. (1998), "Asset prices, consumption, and the business cycle," forthcoming in J. Taylor and M. Woodford, eds., *Handbook of Macroeconomics*, North-Holland.
- Campbell, J. and J. Ammer (1993), "What moves the stock and bond markets? A variance decomposition for long-term asset returns," *Journal of Finance* 48, 3-37.
- Campbell, J., A. Lo, and A. MacKinlay (1997), *The Econometrics of Financial Markets*, Princeton: Princeton University Press.
- Cooley, T., and E. Prescott (1995) "Economic growth and business cycles," in T. Cooley, ed., *Frontiers of Business Cycle Research*, Princeton University Press.
- DeLong, J.B., A. Shleifer, L. Summers, and R. Waldman (1990) "Noise trader risk in financial markets," *Journal of Political Economy* 98, 703-738.
- Grossman, S., and R. Shiller (1981) "The determinants of the variability of stock market prices," *American Economic Review* 71, 222-227.
- Evans, G., and S. Honkapohja (1998), "Learning dynamics," forthcoming in J. Taylor and M. Woodford, eds., *Handbook of Macroeconomics*, North-Holland.
- Farebrother, R.W. (1980), "The Durbin-Watson test for serial correlation when there is no intercept in the regression," *Econometrica* 48, 1553-63.
- Grandmont, J.M. (1998), "Expectations formation and stability of large socioeconomic systems," *Econometrica* 66, 741-781.
- Hansen, G.D. (1993), "The cyclical and secular behavior of the labor input: Comparing efficiency units and hours worked," *Journal of Applied Econometrics* 8, 71-80.
- Holland, J.H. (1975), *Adaptation in Natural and Artificial Systems*, Ann Arbor: University of Michigan Press.
- Hommes, C., and G. Sorger (1998a) "Consistent expectations equilibria," *Macroeconomic Dynamics* 2, 287-321.

- Hommes, C., and G. Sorger (1998b) "Consistent expectations equilibria in overlapping generations economies," working paper, University of Amsterdam and University of Vienna.
- Hurd, M. (1989) "Mortality risk and bequests," *Econometrica* 57, 779-813.
- LeRoy, S.F. and R.D. Porter (1981), "The present-value relation: Tests based on implied variance bounds," *Econometrica* 49, 555-74.
- Marcet, A. and J.P. Nicolini (1997), "Recurrent hyperinflations and learning," Economics working paper no. 244, Universitat Pompeu Fabra.
- Marimon, R. (1997), "Learning from learning in economics," in D.M. Kreps and K.F. Wallis, (eds.) *Advances in Economics and Econometrics: Theory and Applications*, Volume 1, 278-315.
- Michalewicz, Z. (1994), *Genetic Algorithms + Data Structures = Evolution Programs*, 2nd ed., New York: Springer-Verlag.
- Poterba, J. and L. Summers (1988) "Mean reversion in stock returns: evidence and implications," *Journal of Financial Economics* 22, 27-60.
- Sargent, T.J. (1993), *Bounded Rationality in Macroeconomics*, Oxford: Oxford University Press.
- Sargent, T.J. (1998), *The Conquest of American Inflation*, Marshall Lectures manuscript.
- Shiller, R.J. (1981), "Do stock prices move too much to be justified by subsequent changes in fundamentals?" *American Economic Review* 71, 421-36.
- Shiller, R.J. (1989), *Market Volatility*, Cambridge: MIT Press.
- Timmerman, A.G. (1993), "How learning in financial markets generates excess volatility and predictability in stock prices," *Quarterly Journal of Economics* 108, 1135-1145.
- Timmerman, A.G. (1996), "Excess volatility and predictability of stock prices in autoregressive dividend models with learning," *Review of Economic Studies* 63, 523-557.
- West, K.D. (1988), "Bubbles, fads and stock price volatility tests: A partial evaluation," *Journal of Finance* 43, 639-660.