

Complex Dynamics in Equilibrium Asset Pricing Models with Boundedly Rational, Heterogeneous Agents

Paul M. Beaumont^a, Yuanying Guan^b, Alec N. Kercheval^{c,*}

^a*Florida State University, Department of Economics, Tallahassee, FL 32306 USA*

^b*Indiana University Northwest, Department of Mathematics and Actuarial Science, Gary, IN 46408 USA*

^c*Florida State University, Department of Mathematics, Tallahassee, FL 32306 USA*

Abstract

We study a simple model based upon the Lucas framework where heterogeneous agents behave rationally in a fully intertemporal setting but do not know other investors' personal preferences, wealth or investment portfolios. As a consequence, agents initially do not know the equilibrium asset pricing function and must make guesses which they update via adaptive learning with constant gain.

We demonstrate that even in this simple environment the economy can, depending on parameters, exhibit either stable convergence to equilibrium, or chaotic dynamical behavior of asset prices and trading volume without converging to the rational expectations equilibrium of the Lucas model. This contradicts the assertion that the Lucas model is stable in the face of modest deviations from the strong assumptions required to compute the equilibrium.

Keywords: learning, financial markets, rational expectations equilibrium, learning, agent-based modeling, equilibrium asset pricing

*Corresponding author

Email addresses: beaumont@fsu.edu (Paul M. Beaumont), guany@iun.edu (Yuanying Guan), kercheval@math.fsu.edu (Alec N. Kercheval)

1. Introduction

During the fall of 2007, financial markets in the United States began to unwind. In the coming year markets would plummet by over fifty percent and then fall an additional twenty-five percent by March 2009; record numbers of banks would fail; Bear Stearns, Merrill Lynch and Lehman Brothers would vanish. “We were seeing things that were 25-standard deviation moves, several days in a row,” said Goldman Sachs’s chief financial officer David Viniar [1].¹ In the ensuing search for guilty parties, fingers pointed toward the complex, equilibrium-based, mathematical models used by Wall Street quants.²

The *equilibrium asset pricing model* in question is based upon the asset pricing theory of 1995 Nobel Prize winner Robert Lucas of Chicago [5]³ and is motivated by the rational expectations theory of John Muth [7] and the efficient markets theory of Fama [8]. Although the quants on Wall Street may not be particularly concerned with the esoteric details of this theory, Farmer and Geanakoplos [9] note that the asset pricing equations that the quants use are largely based upon this framework.

In the Lucas model there are many investors who may buy and sell some set of securities in order to maximize their expected life-time utility. The investors are forward-looking and rational, meaning that they try to implement an investment strategy that will optimize their long-term welfare in the face of an uncertain future, and they know that asset prices should be determined by market fundamentals such as the stream of future earnings and the risks associated with those earnings. Furthermore, each investor believes that all of the other investors in the market are also rational and that those investors, in turn,

¹See [2] for an entertaining discussion of just how rare a 25-sigma event is.

²See, for example, The Economist [3] for critiques of equilibrium models and the response [4] by Robert Lucas.

³The Lucas [5] paper is one of the most cited economics papers of the past four decades [6].

believe the same. In order to determine their optimal portfolio and the market equilibrium, investors must compute future aggregate demands for assets, the future wealth distribution of all other investors, and the transition dynamics of these future wealth distributions [10]. Consequently, the state space of each investor is a very complicated object requiring an enormous amount of information to construct—more information than is available to the econometrician modeling the the system (see [11]).⁴

In fact, Lucas never intended his model to be a description of investor behavior. Rather, he argued that the outcome of the market can be described “as if” the investors were using the model. Lucas writes that the rational expectations assumption

... is not “behavioral”: it does not describe the way agents think about their environment, how they learn, process information, and so forth. It is rather a property likely to be (approximately) possessed by the *outcome* of this unspecified process of learning and adapting. One would feel more comfortable, then, with rational expectations equilibria if these equilibria were accompanied by some form of “stability theory” which illuminated the forces which move an economy toward equilibrium. [5, pg. 1429]

Lucas asks, “will an economy with agents armed with “sensible” rules-of-thumb, revising these rules from time to time so as to claim observed rents, tend as time passes to behave as described. . . [5, pg. 1437]?” Lucas concludes, “A relatively crude use of hindsight, applied in a reasonably stationary physical environment, will lead to behavior well-approximated by rational expectations [5, pg. 1444].”

⁴The problem is greatly simplified if we assume, as did Lucas [5], that all investors are identical. In this *homogeneous investor* case the wealth distribution becomes degenerate and stationary so that the state space and the asset pricing equations for the market equilibrium can easily be computed by the representative agent.

It is this conclusion that we investigate here: that the market equilibrium described by the Lucas model is a local stable attractor and that *reasonable* investors in a *reasonably* functioning market are at least approximately well described by the model. On the contrary, we show that, when the model is populated with heterogeneous investors facing reasonable restrictions in their information set and using sensible rules-of-thumb learning rules, unstable and chaotic asset market dynamics can arise that are inconsistent with the rational expectations equilibrium of the Lucas model.⁵

We make no effort to build a model realistic enough to correctly forecast real market quantities, nor to understand phenomena like bubbles and crashes as in the recent work [14]. On the contrary, we study the simplest case of two agents to show how chaos can emerge robustly from the basic model structure. If chaos can emerge in a two-agent economy, we should expect larger models with similar structure to display at least as much complexity.

2. Related Literature

Our model falls into the broad category of heterogeneous agent learning models reviewed by [15]. More specifically, we are interested in the stability of markets when investors lack perfect foresight but otherwise behave “rationally” and market prices are determined endogenously through market clearing. Our approach is related to [16], who provide review of the relevant literature, but our model structure is quite different.

We use a version of the Lucas model with a single risky asset and heterogeneous investors with differing degrees of risk aversion. All investors are fully rational in terms of their optimizing behavior but are boundedly rational in the sense that no investor knows the private information, such as risk aversion and

⁵Also relevant but beyond the scope of this paper are the issues of computability raised by [12] and [13].

wealth, of other investors. With this restricted state space investors are unable to derive the aggregate asset pricing equation and must base their *optimal* investment strategy upon their best informed guess about the behavior of future asset prices. Given a hypothesized aggregate pricing function, each agent solves the dynamic optimization problem for that agent’s conditional consumption and asset demand equations. None of our investors are “noise traders” (see, for example, [17, 18, 19]) nor do they ignore any observable information about market fundamentals such as the known distribution of dividend payments or market clearing prices.

Market clearing prices are computed using the conditional demand equations. As the market evolves the investors learn how well their aggregate pricing function performs and they adapt their guesses to be more consistent with observed market outcomes. Each investor solves their fully intertemporal dynamic stochastic control problem using the best information that they have in hand at the time and they update that solution continuously as markets evolve. This extends the work of [20] in that our model is fully intertemporal and market prices are fully endogenous.

We define a *correct expectations equilibrium* to be when the market clearing price is consistent with the aggregate pricing function used in common by all investors and at those prices all agents behave optimally. Our definition is most closely related to that of *internally rational expectations equilibrium* used by [21] and to the *consistent expectations equilibrium* used by [22] who require that the autocorrelation function of the aggregate pricing function be consistent with the observed market data.

We will prove that our definition implies the existence of a no trade equilibrium that is consistent with the rational expectations hypothesis. Like [22], we use a constant gain learning mechanism, and find that investors may or may not be able to learn this equilibrium. For a range of risk aversion ($\gamma > 1$), any constant gain parameter turns out to lead to stable convergence to equilibrium.

This contradicts the “E-stability principle” of Evans and Honkapohja [23], because our framework differs from theirs in that our market prices are being determined endogenously via market clearing, rather than via a price process with exogenous shocks.

For other parameter choices, when the gain parameter is not too small, we find the economy converges to a variety of stable, non-equilibria, including cyclical or fully chaotic dynamics. (Deterministic models exhibiting this behavior include [24] and [25].) Our results are consistent with [26] who demonstrates that slow learning can stabilize the market.

We conclude that our model represents a version of the Lucas model where investors use a strategy at least as good as a “relatively crude use of hindsight” but does not always “lead to behavior well-approximated by rational expectations”.

3. A Lucas Model Framework

We begin with a discrete-time market with N traders, or *agents*, and a total of one share of a single risky, infinitely-lived, dividend-paying asset. The ex-dividend price of the asset is denoted p_t at time $t = 1, 2, \dots$, and the dividend is d_t . Agent i can hold at time t a fraction $s_{i,t} \in [0, 1]$ of the asset (no short sales), and therefore at the end of period t has total wealth $s_{i,t}(p_t + d_t)$, which must be allocated to consumption $c_{i,t}$ and reinvestment for next period $s_{i,t+1}p_t$. This is summarized by the budget constraints

$$c_{i,t} + s_{i,t+1}p_t = (p_t + d_t)s_{i,t} \quad \text{and} \quad s_{i,t+1} \geq 0, c_{i,t} \geq 0 \quad (1)$$

which must hold at each time t and for each agent $i = 1, \dots, N$. Furthermore, the stock market must clear in each period, meaning for all t :

$$\sum_{i=1}^N s_{i,t} = 1 \quad \text{which implies, from (1),} \quad \sum_{i=1}^N c_{i,t} = d_t. \quad (2)$$

The dividend is determined by an exogenous *i.i.d.* process z_t via

$$d_{t+1} = d_t z_{t+1}, \quad t = 0, 1, 2, \dots, \quad (3)$$

where d_0 is a known positive initial value, initial portfolios $s_{i,1}$ are known, and z_t has some fixed known positive distribution. The investment share $s_{i,t+1}$ is agent i 's choice variable, with the budget constraint (1) then determining consumption $c_{i,t}$.

3.1. Order of Events in each Period

At the beginning of each period $t \geq 1$, the dividend d_t is announced and agent i knows how many shares $s_{i,t}$ of stock she holds, but not yet the market price p_t . Each agent determines an optimal investment policy function $s_{i,t+1} = \mathbf{s}_{i,t}(p_t)$ that is a monotone decreasing function of market price p_t . The market price is then determined by the unique solution of the market clearing condition

$$\sum_{i=1}^N \mathbf{s}_{i,t}(p_t) = 1 \quad (4)$$

which can be thought of as the outcome of a price auction among the agents. With the market clearing price now announced, agents know their wealth $w_{i,t} = s_{i,t}(p_t + d_t)$ and proceed to trade to the position $s_{i,t+1} = \mathbf{s}_{i,t}(p_t)$ and consume $c_{i,t} = w_{i,t} - s_{i,t+1}p_t$. Here the period t ends and $t + 1$ can begin.

3.2. The Problem of Limited Information

Everything depends on how the agents determine their “optimal” investment/consumption policies. This is a delicate matter that turns on exactly what the agents know. In contrast to the setting of rational expectations, our agents have limited information: they observe the dividend (and know the probability distribution of the dividend shocks z_t) and the current market price, but do not know the holdings or preferences of the other agents, and hence do not know the probability distribution of tomorrow’s market price. For the agents, “optimal” will be defined only in terms of the available information.

Agents wish to maximize a time-separable lifetime expected discounted utility of consumption of the form

$$U_i = E \left[\sum_{t=1}^{\infty} \beta^{t-1} u_i(c_{i,t}) \right] \quad (5)$$

where $0 < \beta < 1$ is a common discount factor and u_i is a single-period increasing concave utility function for agent i . We take u_i to be a CRRA (Constant Relative Risk Aversion) utility,

$$u_i(c) = \frac{c^{1-\gamma_i}}{1-\gamma_i}, \quad (\gamma_i \neq 1) \quad (6)$$

where $\gamma_i > 0$ is a risk-aversion parameter (that can vary by agent). When $\gamma_i = 1$ we use

$$u_i(c) = \log c. \quad (7)$$

Unfortunately, our agents cannot directly maximize (5) because the expectation cannot be computed: the distribution of future prices, hence consumption, depends on future actions of the other agents, and hence is unknown.

Instead, agents will maximize a *conditional* utility, conditional on the choice of a pricing function giving the market price as a function of dividend. Agents can then determine a conditional optimal policy. If all agents choose the same pricing function, which in addition correctly forecasts the market clearing price in each period, then the economy will be in equilibrium. If not, agents may update their pricing function each period according to some learning rule, and the economy becomes a dynamical system whose dynamics we can investigate, and we will see that equilibrium is not always stable. These topics are addressed in the following sections.

3.3. Defining the Agents' Optimization Problem and Market Equilibrium

We assume all agents know the distribution of the exogenous i.i.d. shocks z_t , $t = 1, 2, 3, \dots$. Thus, at time t , the history of dividends d_1, d_2, \dots, d_t is known, and so is the distribution of each future dividend

$$d_{t+s} = d_t z_{t+1} \cdots z_{t+s}. \quad (8)$$

This defines a natural probability measure Q on the space Ω of infinite sequences of dividends. Let \mathcal{F}_t denote the σ -field generated by d_1, \dots, d_t , and define the σ -field $\mathcal{F} = \bigcup \mathcal{F}_t$ and the filtration $F = \{\mathcal{F}_t : t = 1, 2, 3, \dots\}$. We can describe the entire setting by a filtered probability space $(\Omega, \mathcal{F}, F, Q)$ associated to the dividends process. Note that this space does not represent a complete description of the states of our economy, because the preferences and choices of the agents are required to determine the asset price process.

Now we focus on a single agent i , and temporarily drop the subscript i . The agent can observe her holdings s_t and the dividend d_t at the beginning of period t , and the asset price p_t once the market clears. It is convenient to use the agent's wealth w_t as a state variable, given by

$$w_t = (p_t + d_t)s_t = c_t + p_t s_{t+1}. \quad (9)$$

Since $c_{i,t}$ and hence $u_i(c_{i,t})$ are not F -adapted, agents have no choice but to conditionally maximize U under additional assumptions about the hidden variables.

As we will show, an optimal consumption and investment policy can be derived from a pricing function that describes the agent's view of the distribution of market clearing prices in future periods. Since the true equilibrium price depends on unobserved quantities, our agents must hypothesize a pricing function that only depends, at most, on the history of observed prices and dividends. (The derived consumption and investment policies will then depend on these and the agent's current stock holding.) At a no-trading equilibrium, with all agents holdings fixed and for a given initial stock price, we expect the market clearing price to depend only on the history of dividends.

In this paper we choose the simplest form $p : D \rightarrow R^+$, $p(d) = vd$ for the agent's forecast of future market price (D is the set of possible dividends). This is justified because it is the form of the equilibrium price in the simple homogenous agent case, and also is shown in Theorem 3 to be an equilibrium pricing function for the heterogeneous 2-agent case.

In effect, we presume that agents are aware of Theorem 3 and derive their one-time-step optimal policy conditional on a guessed equilibrium pricing function of this form. Later, we will provide agents with a learning rule that allows them to update their pricing function forecast based on the market outcome. In many cases, this will result in a market that converges to a stable equilibrium with agents agreeing on a pricing function $p(d) = v^*d$ that correctly forecasts the market clearing price, and so this learning strategy will be successful. The convergence to equilibrium also justifies, after the fact, the reasonable use of the form $p(d) = vd$ in the agents forecast. However, we will see that success depends on the choice of parameters. In certain cases of low risk aversion, the same assumption leads to behavior that fails to converge to equilibrium in interesting ways, as we describe in section 6.

If we denote by W the set of possible values of wealth (non-negative reals), then we define a consumption policy to be a function $\mathbf{c} : W \rightarrow R^+$, and an investment policy a function $\mathbf{s} : W \rightarrow R^+$. At the beginning of period t , since s_t and d_t are known, wealth is then a function of market price p_t via (9).

At the beginning of period t , for a pricing function $p(\cdot)$, and admissible policy function \mathbf{c} , define the conditional wealth \bar{w}_{t+s} for $s \geq 0$ recursively by

$$\bar{w}_{t+s} = (\bar{w}_{t+s-1} - \mathbf{c}(\bar{w}_{t+s-1})) \frac{p(d_{t+s}) + d_{t+s}}{p(d_{t+s-1})} \text{ for } s \geq 1, \quad (10)$$

and, when $s = 0$,

$$\bar{w}_t = (w_{t-1} - \mathbf{c}(w_{t-1})) \frac{p(d_t) + d_t}{p_{t-1}}. \quad (11)$$

Then we define the conditional utility functional by

$$U(t, \mathbf{c}, p(\cdot)) = E \left[\sum_{s=0}^{\infty} \beta^s u(\mathbf{c}(\bar{w}_{t+s})) | \mathcal{F}_t \right]. \quad (12)$$

Conditional on $p(\cdot)$ and \mathbf{c} , the process \bar{w}_t is adapted to the filtration F , so the conditional utility $U(t, \mathbf{c}, p(\cdot))$ is well-defined.

Denote by \mathcal{C} the set of admissible consumption policies, satisfying $0 \leq \mathbf{c}(w) \leq w$. The optimal utility is then given by the supremum over all admissible con-

sumption policies of the infinite-horizon expected utility, conditional on the pricing function $p(\cdot)$,

$$U^*(t, p(\cdot)) = \sup_{\mathbf{c} \in \mathcal{C}} U(t, \mathbf{c}, p(\cdot)). \quad (13)$$

The solution of this problem will be an optimal consumption policy \mathbf{c}^* , and the corresponding optimal investment policy

$$\mathbf{s}^*(w_t) = \frac{w_t - \mathbf{c}^*(w_t)}{p_t}. \quad (14)$$

An agent's actual time- t wealth and consumption are still not known until the market clearing price p_t is revealed. We now reintroduce the subscript i indicating the i th agent, $w_t = w_{i,t}$.

The market clearing price p_t for time t is determined by interpreting these policy functions (for each agent) as functions of p and applying the market-clearing constraint

$$1 = \sum_i s_{i,t+1} = \sum_i \frac{(p + d_t)s_{i,t} - \mathbf{c}_i^*((p + d_t)s_{i,t})}{p} \quad (15)$$

to solve for $p = p_t$. The agents now know their actual wealth $w_{i,t} = (p_t + d_t)s_{i,t}$, consume $\mathbf{c}_i^*(w_{i,t})$, and invest in $\mathbf{s}_i^*(w_{i,t})$ shares, ending the period.

If the market is at equilibrium, this market clearing price p_t will agree with the value $p(d_t)$ of the pricing function used in common by each agent, and for each i we will have $\bar{w}_{i,t} = w_{i,t}$. We formalize this in a definition.

Definition 1. A correct expectations equilibrium (CEE) consists of an aggregate stock pricing function $p(d)$ and consumption and investment policy functions for each agent, $\mathbf{c}_i(w)$ and $\mathbf{s}_i(w)$, such that

1. For each i , \mathbf{c}_i solves agent i 's optimization problem (13) conditional on $p(\cdot)$,
2. For each i , and t , the consumption and investment policy functions satisfy

$$\mathbf{c}_i(w_t) + p(d_t)\mathbf{s}_i(w_{i,t}) = \mathbf{s}_i(w_{i,t-1})(p(d_t) + d_t), \quad (16)$$

3. $p(d_t)$ is the correct market clearing price resulting from the policies of all the agents, and
4. the stock and consumption markets clear: $\sum_i s_{i,t} = 1$ and $\sum c_{i,t} = d_t$ for each t .

A few comments about this definition are in order. The reason for introducing the CEE in place of the more familiar rational expectations equilibrium (REE) is that for an REE all agents have complete information about the economy, so that the equilibrium pricing and policy functions could depend on variables, such as the aggregate wealth distribution, that are not visible to our boundedly rational agents. Agents cannot evaluate functions that depend on unobserved variables, but our agents can evaluate the policy and pricing functions of a CEE.

If the economy has arrived at a CEE, then this also qualifies as an REE in which the optimal pricing and policy functions do not depend on the unobservable variables (e.g., those variables are constant in time). However, there may be REEs that do not qualify as CEEs.

4. Solving the Single Period Optimal Policy

Again we consider the problem faced by a single agent i in our market, and drop the subscript i for ease of notation.

The agent's objective function at time t is

$$U^*(t, p(\cdot)) = \sup_{\mathbf{c} \in \mathcal{C}} U(t, \mathbf{c}, p(\cdot)). \quad (17)$$

In outline, the infinite-horizon objective function is given by the solution of a recursive value function equation, and the corresponding optimal consumption policy is a solution to a stochastic Euler equation. For CRRA utility with $\gamma > 0$ and i.i.d. dividend growth, as in our model, a result of Duffie [27] (see also [28]) gives us a particularly easy solution at equilibrium: the optimal consumption policy is $\mathbf{c}(w_t) = \delta w_t$ for a known constant δ , and the equilibrium pricing

function takes the form $p(d_t) = vd_t$ for some constant v , the price-dividend ratio. The agent's problem for one time step is solved, and the only open question for the agents is how to adjust the constant v so that $p = vd$ agrees with the observed market clearing price, and hence find a market equilibrium.

Here is the story in more detail, with full proofs relegated to [28].

The recursive version of the optimization problem is known as the value function equation, or Bellman's equation, equation (18) below. In the recursive formulation, we use w, d, c to denote today's values (time t) of the state and choice variables, and w', d', c' to denote tomorrow's values (time $t + 1$).

Theorem 1 (see [28]). *Fix a pricing function $p(d)$ and a time t .*

1. *Suppose $V^*(w)$ is a function satisfying the Bellman equation*

$$V(w) = \max_{c \in [0, w]} \{u(c) + \beta E[V(w') | \mathcal{F}_t]\} \quad (18)$$

where $w' = (w - c) \frac{p(d') + d'}{p(d)}$ and suppose

$$\lim_{\tau \rightarrow \infty} E\beta^\tau V^*(w_\tau) = 0. \quad (19)$$

Then there is a unique continuous policy function $c^(w)$ such that*

- $V^* = U^*$, the optimal utility function, and
- $c^*(w)$ is the optimal policy, attaining the supremum in the definition of U^* , and satisfying

$$V^*(w) = u(c^*(w)) + \beta E[V^*(w') | \mathcal{F}_t] \quad (20)$$

2. *Moreover, the Bellman equation has a unique solution V^* within the class of continuous functions that tend to infinity no faster than u . Also, the function V^* is continuous, concave and differentiable and*

$$\frac{dV^*(w)}{dw} = u'(c^*(w)). \quad (21)$$

3. The optimal policy c^* satisfies the first order (Euler) equation

$$u'(c^*) - \beta E \left[r' \frac{dV}{dw}(w') | \mathcal{F}_t \right] = 0 \quad (22)$$

where $w' = r'(w - c^*)$ defines the return r' .

Equivalently,

$$u'(c^*) = \beta E[r' u'(c^*) | \mathcal{F}_t]. \quad (23)$$

Equation (23) is the familiar form of the Euler equation, which can be used to solve for the optimal policy at each time step.

4.1. The Homogeneous Lucas Model

For comparison, consider the simpler problem when all agents are assumed identical *and know it*. There can be no trading, so each of the N agents holds $1/N$ shares of the stock, and in each period consumes d/N , where d is the current dividend. The problem of deriving the optimal policy is therefore immediately solved, the market is automatically at equilibrium, and the Euler equation can now be used to derive the equilibrium pricing function.

From (23) and the definition of return, we can rewrite the Euler equation for the generic case $\gamma \neq 1$ as

$$\frac{p_t}{d_t} = \beta E \left[\left(1 + \frac{p_{t+1}}{d_{t+1}} \right) \left(\frac{d_{t+1}}{d_t} \right)^{1-\gamma} \right] \quad (24)$$

A pricing function of the form $p_t = v d_t$ for a constant v (the price-dividend ratio) solves this equation. If we write $z_{t+1} = d_{t+1}/d_t$, then recall that z_t is an i.i.d. sequence with known positive distribution. We obtain

$$v = (1 + v) \beta E[z_{t+1}^{1-\gamma}] \quad (25)$$

where the conditional expectation is replaced by an unconditional expectation due to the independence of z_{t+1} , and the easy solution of this equation is

$$v = \frac{\beta E[z^{1-\gamma}]}{1 - \beta E[z^{1-\gamma}]}. \quad (26)$$

The equilibrium pricing function is $p_t = v d_t$ for this constant v .

4.2. The Heterogeneous Lucas Model

In this case it becomes a nontrivial problem to find the consumption policy optimizing the objective function in (13) because the wealth distribution becomes part of the state space (see, [10]). However, a theorem of Duffie [27] comes to the rescue.

Assuming that the equilibrium pricing function takes the form $p(d) = vd$ for some constant v , the one period return becomes

$$r_t = \frac{vd_{t+1} + d_{t+1}}{vd_t} = \left(\frac{v+1}{v} \right) \frac{d_{t+1}}{d_t}, \quad (27)$$

so the returns r_t are themselves i.i.d. In this case, the following theorem provides the optimal consumption policy for the agent:

Theorem 2 (Duffie [27], see also [28]). *Consider the optimization problem (13).*

1. *Suppose*

$$\gamma \in (0, 1) \cup (1, \infty), \quad u(c_t) = \frac{c_t^{1-\gamma}}{1-\gamma}, \quad c_t \in [0, w_t] \quad (28)$$

and suppose the returns r_t are i.i.d. Define $\rho = \beta E[r_t^{1-\gamma}]$ and $\delta = 1 - \rho^{\frac{1}{\gamma}}$.

If $\rho < 1$, then the policy $\mathbf{c}(w) = \delta w$ is optimal at equilibrium.

2. *Suppose again i.i.d. returns r_t but*

$$\gamma = 1, \quad u(c_t) = \log c_t. \quad (29)$$

Define $\delta = 1 - \beta$.

If $E[\log r_t] < \infty$, then the policy $\mathbf{c}(w) = \delta w$ is optimal at equilibrium.

Since our agents have no knowledge of the preferences or holdings of the other agents, but still wish to behave optimally, we suppose our agents follow this strategy:

1. Guess a market clearing pricing function $p_i(d_t) = v_i d_t$ for some constant v_i .

2. Conditional on this guess, solve the infinite horizon optimization problem (13) according to Theorem 2 to obtain consumption and investment policies \mathbf{c}_i and \mathbf{s}_i .
3. Observe the resulting market-clearing price $p_m(t)$, which is announced to the market as the unique solution of the market-clearing equation

$$1 = \sum_{i=1}^N \mathbf{s}_{i,t}(w_{i,t}) = \sum_{i=1}^N \frac{(1 - \delta_{i,t})w_{i,t}}{p_m(t)} = \left(1 + \frac{d_t}{p_m(t)}\right) \sum_{i=1}^N (1 - \delta_{i,t})s_{i,t} \quad (30)$$

where $\delta_{i,t}$ is the consumption policy constant for agent i in time t coming from Theorem 2, and $s_{i,t}$ is the beginning-of-period stock holding of agent i .

4. If $p_i(t) \neq p_m(t)$, the agent knows the market is out of equilibrium and employs a learning rule to revise the guess v_i to a new value for next period.

Agents are all trying to learn the market equilibrium when they can only observe their own private information and the market-clearing price in each period.

Several questions now arise.

1. Does there exist a CEE equilibrium in this market, and if so, what does it look like?
2. Are there reasonable learning rules that lead the market to equilibrium? If the market fails to converge to equilibrium, how can it fail to converge? Can there be periodic motion? Can there be chaos? Can this mechanism lead to excess price volatility and excess trading volume?

We answer these questions in the next sections by focusing on the two-agent ($N = 2$) case.

5. Existence of Equilibrium

We henceforth set $N = 2$. Fix a time t , and for $i = 1, 2$ let s_i, w_i, s'_i, w'_i denote the number of shares of stock and total wealth held by agent i at times t and $t + 1$, respectively. Let us investigate whether there is a CEE of the form $p(d) = v^*d$ and $\mathbf{c}_i(w_i) = \delta_i w_i$, $\mathbf{s}_i(w_i) = s'_i = \frac{w_i - \mathbf{c}_i(w_i)}{p(d)}$ for some constants v^* and δ_i .

Let $z' = d'/d$ denote our i.i.d. dividend growth random variable, and define $\theta_i = E[z^{1-\gamma_i}]$, where γ_i is the risk aversion parameter for agent i .

If the market clearing price is $p(d) = v^*d$, then, as before, the stock return

$$r' = \frac{p(d') + d'}{d} = \frac{v^* + 1}{v^*} z'$$

is i.i.d., so by Theorem 2 the optimal consumption policy for agent i is $\mathbf{c}_i(w) = \delta_i w$, where

$$\delta_i = 1 - \left[\beta E[r'^{(1-\gamma_i)}] \right]^{1/\gamma_i} = 1 - \left(1 + \frac{1}{v^*} \right)^{\frac{1-\gamma_i}{\gamma_i}} (\beta \theta_i)^{1/\gamma_i}. \quad (31)$$

The stock holding for agent i is then updated according to

$$s'_i = \frac{w_i - c_i}{p(d)} = (1 - \delta_i) \left(\frac{p(d) + d}{p(d)} \right) s_i = (1 - \delta_i) \left(1 + \frac{1}{v^*} \right) s_i \equiv C_i s_i, \quad (32)$$

where C_i is defined to be the constant $(1 - \delta_i)(1 + \frac{1}{v^*})$.

Next, we argue that, in equilibrium, unless $s_i = 0$, we must have $C_i = 1$. Since $s'_i = C_i s_i$ for each time period, if $s_i > 0$ then $C_i \leq 1$ since $s_{i,t}$ is bounded above by one.

The budget constraint in periods t and $t + 1$ gives us the equations

$$s_1 + s_2 = 1 \quad (33)$$

$$C_1 s_1 + C_2 s_2 = 1 \quad (34)$$

Case 1. $s_1 > 0$ and $s_2 > 0$. Then equations (33, 34) imply $C_1 = C_2 = 1$, and hence there is no trading and $\delta_1 = \delta_2$. This means either $\gamma_1 = \gamma_2$ (so the

agents are identical except for holdings), or else

$$\left(1 + \frac{1}{v^*}\right)^{\frac{1-\gamma_i}{\gamma_i}} (\beta\theta_i)^{1/\gamma_i} \quad (35)$$

is independent of i , implying a special relationship between γ_1 , γ_2 , and the distribution of z (occurring with probability zero).

Case 2. One agent (say agent 2) holds zero shares of stock at equilibrium. Then $s_1 = 1$, $C_1 = 1$, and there is no trading. This will be the generic case when $\gamma_1 \neq \gamma_2$.

Either way, writing $\delta = \delta_1$, $\theta = \theta_1$, the market clearing condition becomes

$$1 = (1 - \delta)(1 + (1/v^*)) \quad (36)$$

or

$$\delta = 1 - \frac{v^*}{v^* + 1}. \quad (37)$$

Combining this with equation (31) and simplifying, we obtain

$$v^* = \frac{\beta\theta}{1 - \beta\theta} \quad \text{and} \quad \delta = 1 - \beta\theta. \quad (38)$$

These values of v^* and δ now satisfy the budget constraint, market clearing, and optimality conditions from the definition of CEE. Comparing notation from Theorem 2, $\rho = \beta E[r^{1-\gamma}] = \beta\theta(1 + (1/v^*))^{1-\gamma}$, it is easy to show that $\beta\theta = \rho^{1/\gamma}$ and so the condition $\rho < 1$ from Theorem 2 is equivalent to $\beta\theta < 1$.

We summarize these conclusions in the following theorem.

Theorem 3 (Two-agent CEE). *Suppose we have a two-agent economy with a single dividend-paying stock in unit total supply. Suppose $d_0 > 0$ and in period $t + 1$ the stock pays a per-share dividend $d_{t+1} = z_{t+1}d_t$, where $\{z_t\}$ is an i.i.d. sequence of positive random variables. The agents have CRRA utility (6, 7) with risk aversion parameters $\gamma_i > 0$, $i = 1, 2$.*

Let $\theta_i = E[z^{1-\gamma_i}]$, $0 < \beta < 1$ be a common discount factor, and suppose $\beta\theta_i < 1$ for $i = 1, 2$.

Then:

1. There exists positive constants v^*, δ_1, δ_2 such that the pricing function $p(d) = v^*d$ and the consumption policies $\mathbf{c}_i(w) = \delta_i w$ form a Correct Expectation Equilibrium (CEE), in which the consumption policy is optimal for the infinite horizon utility (12) conditional on $p(d)$, and the market clears in each period at $p(d)$.
2. At equilibrium there is no trading. Either (a) one agent (agent $i^* \in \{1, 2\}$) has all the stock, or else (b) $\delta_1 = \delta_2$, $\theta_1 = \theta_2$, and the two agents have the same optimal policy.
3. Define $\delta = \delta_{i^*}$, and $\theta = \theta_{i^*}$ in case (a) above, and $\delta = \delta_1 = \delta_2$, $\theta = \theta_1 = \theta_2$ in case (b). Then at equilibrium,

$$v^* = \frac{\beta\theta}{1 - \beta\theta} \quad (39)$$

and

$$\delta = 1 - \beta\theta. \quad (40)$$

6. Learning, Stability, and Chaos

6.1. Deriving the Dynamical System of the Economy

Our two agents attempt to learn the equilibrium pricing function $p(d) = v^*d$ over time by starting with a guess $p_i(d) = v_i d$, computing optimal policy conditional on this guess, observing the resulting market clearing price p_m , and then updating their guess.

We specify the *learning rule* to be

$$v'_i = \eta_i v_m + (1 - \eta_i) v_i \quad (41)$$

where η_i is a learning parameter controlling the learning speed.

The state variables are each agent's holdings s_1, s_2 , and price-dividend ratio forecast v_1, v_2 . The dynamical system is determined by how these variables get updated in the next period (indicated by primes). Since $s_2 = 1 - s_1$, we have a

three dimensional dynamical system s'_1, v'_1, v'_2 as a function of s_1, v_1, v_2 , which we now derive.

Recall $\theta_i = E[z^{1-\gamma_i}]$, and let

$$\alpha_i = (\beta\theta_i)^{1/\gamma_i}. \quad (42)$$

The conditionally optimal consumption policy is $c_i = \delta_i w_i$, where

$$\delta_i = 1 - \left(1 + \frac{1}{v_i}\right)^{\frac{1-\gamma_i}{\gamma_i}} \alpha_i. \quad (43)$$

The budget constraint is

$$c_i + p_m s'_i = w_i = (p_m + d)s_i \quad (44)$$

or

$$s'_i = (1 - \delta_i)(1 + (1/v_m))s_i. \quad (45)$$

The market clearing conditions are $s'_1 + s'_2 = 1$ and $s_1 + s_2 = 1$, so

$$(1 + (1/v_m))[(1 - \delta_1)s_1 + (1 - \delta_2)s_2] = 1, \quad (46)$$

or

$$1 + (1/v_m) = \frac{1}{1 - \delta_1 s_1 - \delta_2 (1 - s_1)}. \quad (47)$$

Substituting into the expression for s'_i gives

$$s'_1 = \frac{(1 - \delta_1)s_1}{1 - \delta_1 s_1 - \delta_2 (1 - s_1)}. \quad (48)$$

To determine v_m , we can solve (47) to get

$$v_m = \frac{1 - \delta_1 s_1 - \delta_2 (1 - s_1)}{\delta_1 s_1 + \delta_2 (1 - s_1)}. \quad (49)$$

Now equations (41), (43), (48), and (49) are the equations defining the dynamics of the economy, depending on the parameters $\gamma_1, \gamma_2, \eta_1, \eta_2$, and the distribution of z . *We now simplify the situation considerably by choosing $\gamma_2 = 1$, so that the second agent has log-utility.*

Proposition 1. *For a log-utility agent the choice of pricing function, hence the choice of η , has no effect on the agent's policy function, and hence no effect on the market clearing price.*

Proof. By Theorem 2, the optimal policy of a log-utility agent is $c = \delta w$ where $\delta = 1 - \beta$ is independent of the state variables, and therefore not influenced by the pricing function forecast or learning parameters. \square

With the state variable v_2 eliminated we now have, for $s_1 \neq 0$, a two-dimensional dynamical system in the variables s_1, v_1 , depending on the parameters γ_1, η_1 , and α_1 . Note

$$\delta_1 = 1 - \alpha_1 \left(1 + \frac{1}{v_1}\right)^{\frac{1-\gamma_1}{\gamma_1}}, \quad \delta_2 = 1 - \beta. \quad (50)$$

The investment policy for agent 1 is

$$s'_1 = (1 - \delta_1)s_1 \left(1 + \frac{1}{v_m}\right) \quad (51)$$

where, from market clearing,

$$1 + \frac{1}{v_m} = \frac{1}{(1 - \delta_1)s_1 + (1 - \delta_2)(1 - s_1)}. \quad (52)$$

Substituting into equation (51) gives

$$s'_1 = \frac{1}{1 + \frac{1-\delta_2}{1-\delta_1} \frac{1-s_1}{s_1}}, \quad (53)$$

and we obtain the dynamical equation for s_1 :

$$s'_1 = \frac{1}{1 + \left(\frac{\beta}{\alpha_1}\right) \left(1 + \frac{1}{v_1}\right)^{1-1/\gamma_1} \left(\frac{1-s_1}{s_1}\right)}. \quad (54)$$

From (51),

$$1 + \frac{1}{v_m} = \frac{s'_1}{(1 - \delta_1)s_1} = \frac{s'_1}{\alpha_1 \left(1 + \frac{1}{v_1}\right)^{\frac{1-\gamma_1}{\gamma_1}} s_1} \quad (55)$$

so

$$v_m = \frac{1}{\frac{s'_1}{\alpha_1 s_1 (1 + \frac{1}{v_1})^{\frac{1-\gamma_1}{\gamma_1}}} - 1} \quad (56)$$

and hence, from (41),

$$v'_1 = (1 - \eta_1)v_1 + \eta_1 \left[\frac{s'_1}{\alpha_1 s_1 (1 + \frac{1}{v_1})^{\frac{1-\gamma_1}{\gamma_1}}} - 1 \right]^{-1}. \quad (57)$$

Equations (54) and (57) define an explicit two-dimensional dynamical system

$$(s'_1, v'_1) = F(s_1, v_1) \quad (58)$$

describing the evolution of our two-agent economy with state variables s_1, v_1 , agent parameters η_1, γ_1 , and depending on β and $\alpha_1 = (\beta E[z^{1-\gamma_1}])^{1/\gamma_1} = (\beta\theta_1)^{1/\gamma_1}$.

6.2. Dynamical Behavior

We now study the dynamics of this two-agent economy as a function of the first agent's risk aversion parameter $\gamma_1 > 0$ and the first agent's learning speed (or "gain") parameter $\eta_1 \in (0, 1)$. (Recall that agent 2 is assumed to have log utility, or $\gamma_2 = 1$.) We need to select a particular family of dividend growth distributions to define the economy, so we arbitrarily suppose that the dividend growth z is lognormal, $z = e^X$ where X is normal with mean $\mu > 0$ and variance $\sigma^2 > 0$. (Similar analysis could be applied to other choices of distribution for z .)

For Theorem 3 to be applicable, we need the condition $\beta\theta_1 < 1$. Since

$$\theta_1 \equiv E[z^{1-\gamma_1}] = \exp[(1 - \gamma_1)\mu + \frac{1}{2}(1 - \gamma_1)^2\sigma^2], \quad (59)$$

$\beta\theta_1 < 1$ will be true for all $\gamma_1 < 1$ if we choose the discount factor $\beta < \exp(-\mu - \sigma^2/2)$. This is satisfied by reasonable choices of these parameters, for example

for our simulations we choose the values $\beta = 0.95$, $\mu = 0.01752$, $\sigma^2 = 0.00123$.⁶

In this section, we rigorously analyze the equilibria for stability, and also computationally study the global dynamics as a function of the parameters. In summary, the results for $0 < \gamma_1 \leq 1$ are as follows:

- As γ_1 ranges between $\eta_1/2$ and 0, the market dynamics passes through a period doubling cascade in which it exhibits stable periodic behavior of successively higher period, enters a region of stable, chaotic behavior, and, for small enough γ_1 , the system becomes divergent.
- For $\eta_1/2 < \gamma_1 < 1$, the market converges as time tends to infinity to an asymptotically stable equilibrium in which agent 1 holds all the stock.⁷ The price-dividend ratio converges to

$$v^* = \frac{\beta\theta_1}{1 - \beta\theta_1}. \quad (60)$$

- For $\gamma_1 = 1$ the two agents are identical and hold their initial share of stock with no trading. The price-dividend ratio converges to $\beta/(1 - \beta)$.

For $\gamma_1 > 1$, we find that the required condition $\beta\theta_1 < 1$ is lost when

$$\gamma_1 \geq \gamma_{\max} \equiv 1 + \mu/\sigma^2[1 + \sqrt{1 + 2(\sigma^2/\mu^2) \log(1/\beta)}].$$

Therefore our analysis is limited to $\gamma_1 < \gamma_{\max}$. For our parameter choices,

⁶In our simple model of qualitative market behavior, the time scale and market parameters are largely arbitrary, and could be adjusted to suit the investigator without changing the basic conclusions. Since unrealistic market parameters would distract attention from our qualitative conclusions, we chose μ and σ^2 based on typical values for annual consumption growth in the US.

⁷Unlike [29], we do not find that the log utility investor always accumulates all wealth in the long-run as per the ‘‘Kelly rule’’. This is because our investors make endogenous consumption and investment portfolio decisions, as opposed to only making portfolio decisions with fixed consumption shares. Thus, the Kelly rule is not necessarily evolutionarily stable when agents must solve both their consumption and savings decisions endogenously.

this bound is approximately $\gamma_{\max} = 32.3$, safely larger than we would normally consider for a risk aversion parameter.

- For $1 < \gamma_1 < 1 + 2\mu/\sigma^2$ (approximately $1 < \gamma_1 < 29.5$ for our parameter choices), the market converges as time tends to infinity to an asymptotically stable equilibrium in which agent 2 (the log agent) holds all the stock. The price-dividend ratio converges to $\beta/(1 - \beta)$.
- For $1 + 2\mu/\sigma^2 < \gamma_1 < \gamma_{\max}$, the agent 1 equilibrium regains stability.

From a dynamical perspective, we can see that the market converges to a stable equilibrium for all allowed values of the risk aversion parameter $\gamma_1 > \eta_1/2$. Hence we are especially interested in the situation where the risk aversion parameter drops below that value $\eta_1/2$, because this is where we find stable non-equilibrium dynamics corresponding to a chaotic attractor in which both agents have stock holdings bounded away from zero for all time, and there is perpetual trading. Depending on the specific choice of γ_1 , we can also see stable periodic behavior of any period, on the way to the chaotic regime.⁸

We now turn to a rigorous analysis of the stability of the equilibria.

First consider the case $\gamma_1 < 1$. By Theorem 3 there is no trading at equilibrium, which therefore corresponds to a fixed point of the dynamical system,

$$s'_1 = s_1 = s_1^* = 1, \quad v'_1 = v_1 = v_m = v_1^* = \frac{\beta\theta_1}{1 - \beta\theta_1}.$$

The equilibrium point is $(s^*, v^*) = (1, \beta\theta_1/(1 - \beta\theta_1))$. The Jacobian matrix DF evaluated at this point is

$$DF_{(s^*, v^*)} = \begin{bmatrix} 1/\theta_1 & 0 \\ * * * & 1 - \eta_1/\gamma_1 \end{bmatrix}.$$

⁸We are not concerned here with calibrating our model to match specific market data so we do not take a position on what the *correct* coefficient of risk aversion should be. Macroeconomists typically calibrate their models with risk aversion greater than two while experimentalists commonly find risk aversion to be less than one (see, [30, 31, 32]).

An asymptotically stable equilibrium point corresponds to all eigenvalues of this matrix less than one in absolute value, i.e.

$$\theta_1 > 1 \text{ and } |1 - \eta_1/\gamma_1| < 1. \quad (61)$$

The first condition follows from $\gamma_1 < 1$ and the definition of θ_1 . Since $\eta_1 > 0$, the second condition is equivalent to $\gamma_1 > \eta_1/2$, and the stability conditions for the equilibrium point become

$$\frac{\eta_1}{2} < \gamma_1 < 1 \quad (62)$$

(independent of the choices of $\mu \geq 0$ and $\sigma > 0$).

Over the range $\gamma_1 > 1$,

$$\theta_1 = \exp[(1 - \gamma_1)\mu + \frac{1}{2}(1 - \gamma_1)^2\sigma^2] < 1 \quad (63)$$

when

$$\gamma_1 < 1 + \frac{2\mu}{\sigma^2}. \quad (64)$$

Therefore, the stability of the equilibrium point $(s^*, v^*) = (1, \beta\theta_1/(1 - \beta\theta_1))$ is lost. However, we see next that the other no-trading equilibrium point $(s^*, v^*) = (0, \beta/(1 - \beta))$ gains stability. This corresponds to the market equilibrium in which agent 2 holds all the stock.

The two dimensional dynamical system F specified in equations (54) and (57) appears to have a singularity at $s = 0$; however the limit as $s \rightarrow 0$ exists, so the singularity is removable and F extends to a continuous function near the fixed point $P = (0, \beta/(1 - \beta))$. Moreover, the partial derivatives can be shown to have a finite limit as $(s, v) \rightarrow P$, and hence by L'Hopital's Rule F extends continuously to a C^1 function at the fixed point P , whose stability is determined by the Jacobian matrix DF evaluated at P , which (after some computation) is

$$DF_{(s^*, v^*)} = \begin{bmatrix} \theta_1^{1/\gamma_1} & 0 \\ * & * & * \\ * & * & * \end{bmatrix}.$$

An asymptotically stable equilibrium point corresponds to all eigenvalues of this matrix less than one in absolute value, i.e.

$$\theta_1^{1/\gamma_1} < 1 \text{ and } |1 - \eta_1| < 1. \quad (65)$$

The second condition is automatically satisfied, and the first condition is equation (64). Thus the stability conditions for this equilibrium point become

$$1 < \gamma_1 < 1 + \frac{2\mu}{\sigma^2}. \quad (66)$$

In the small range $\gamma_{\max} > \gamma_1 > 1 + \frac{2\mu}{\sigma^2}$, we again obtain

$$\theta_1 = \exp[(1 - \gamma_1)\mu + \frac{1}{2}(1 - \gamma_1)^2\sigma^2] > 1, \quad (67)$$

so the equilibrium point $(s^*, v^*) = (1, \beta\theta_1/(1 - \beta\theta_1))$ regains stability.

Notice we have shown that when $\gamma_1 < \eta_1/2$, neither the $s = 1$ nor $s = 0$ equilibrium is stable. This indicates the potential likelihood of non-equilibrium (i.e. non-fixed-point) attractors in the system, which we in fact observe.

To investigate numerically the global dynamics of the economy for various values of the parameters $0 < \gamma_1, \eta_1 < 1$, we fix the values of β, μ, σ^2 mentioned above and simulate the dynamical system by computing F -orbits for various initial conditions, giving the time evolution of stock holding and pricing function forecast for agent 1.

When $\gamma_1 > \eta_1/2$ we observe, for initial conditions around the equilibrium, that the economy converges to (s_1^*, v_1^*) and therefore this equilibrium is stable and attracting under our adaptive learning scheme.

When $\gamma_1 < \eta_1/2$, we know that the equilibrium is no longer asymptotically stable. The observed behavior can be summarized in a two parameter bifurcation plot (Figure 1), a plot of the (η, γ) -plane, where $\eta = \eta_1$ and $\gamma = \gamma_1$, indicating the nature of the observed attractor for each η and γ value. We see that the economy may stably cycle with various periods, and even may exhibit stable chaotic behavior or diverge with $v \rightarrow \infty$.

We also plot a similar diagram for the behavior of the stock holding s of agent 1 in Figure 2. The white region corresponds to parameters for which the system converges to equilibrium, and red represents parameters for which both agents survive (no convergence to equilibrium).

To illustrate in more detail what kind behavior is observed for γ, η parameters in the chaotic region, we plot time series for $(\gamma, \eta) = (0.18, 0.6)$ and $(0.22, 0.8)$, two values in the chaotic region. In Figures 3 and 4, we plot segments of a time series for v and s with $(\gamma, \eta) = (0.18, 0.6)$. The initial 500 time steps beginning with initial conditions $v_0 = 19$ and $s_0 = 0.4$ are plotted in Figure 3, to illustrate the initial transient behavior, and then in Figure 4 we plot 500 time steps starting from $t = 5000$ to illustrate the long-term behavior. There is perpetual trading of small increments, and the price-dividend ratio is seen to fluctuate between about 22 and 70. Figures 5 and 6 tell a similar story for $(\gamma, \eta) = (0.22, 0.8)$.

In Figures 7 and 8, we examine the initial transient behavior a little more closely, this time plotting price and return for the first 100 time steps. In Figure 7 we have a period of seemingly stable upward trend of the price, which then gradually becomes more volatile, breaking into an oscillatory behavior, which, we know from previous graphs, will begin to oscillate with higher and higher period until reaching a long term chaotic attractor. In Figure 8, the chaotic behavior is visible earlier, within only a few time steps.

We point out that we are not observing bifurcation in these plots: the parameters are fixed. Rather, the natural behavior of the system in Figure 7 is to begin with an appearance of stability with a slow trend, followed by gradual encroachment of chaos reflected in the endogenous market clearing price. The long term asymptotic behavior is chaotic.

Another view of the asymptotic attractors is shown by the phase portraits in Figures 9 and 10. Here we fix $\gamma = 0.15$ and let η vary from 0.1 to 0.6, plotting the attractors we observe after all transients have dissipated. As expected, we begin

with an attracting fixed point, bifurcate to a periodic orbit, and by $\eta = 0.48$ observe a chaotic attractor.

For yet another view, Figure 11 shows an orbit diagram focusing on the price-dividend ratio v , with $\eta = 0.5$ and γ varying along the horizontal axis. On the vertical axis are plotted the numerically observed limit points of F -orbits. Above $\gamma = 0.25$, we expect only a single point, the stable equilibrium. As γ decreases below 0.25, the picture reproduces the classic period-doubling cascade toward chaos observed in generic one-dimensional dynamical systems. (See [28] for a rigorous analysis confirming that this is actually what is happening.)

7. Conclusions

We have demonstrated that even in the smallest heterogeneous version of the Lucas model, with a single asset and two boundedly rational investors, it is easy to produce dynamical behavior of the endogenous asset prices and trading volume that is inconsistent with rational expectations behavior. Our investors are boundedly rational only in the sense that their state space is restricted to exclude private information of other investors, such as wealth and degrees of risk aversion, but they continue to use publicly available market information and to be forward-looking when making their investment decisions.

In our two-investor market with learning, one investor has log-utility and the other has CRRA risk aversion parameter $\gamma > 0$ and learning (gain) parameter η , $0 < \eta < 1$. Among our results, we find that when $\gamma < \eta/2$ the market avoids equilibrium and can be chaotic; when $\eta/2 < \gamma < \gamma_0$ the market converges over time to an asymptotically stable equilibrium in which the least risk-averse investor (according to whether or not $\gamma > 1$) holds all the stock. (Here $\gamma_0 = 1 + 2\mu/\sigma^2 \approx 29.5$ for our parameter choices.) When $\gamma > 1$ market stability is highly likely. While this may offer some comfort to financial economists, we note that market instability is a fundamental property of the model for some parameter ranges even in this very simple version of the model. The stability

of more complex models over broad parameter ranges remains an open question but the recent work of [16] casts some doubt that stability will be ubiquitous.

One way to interpret our results is that they contribute to the discussion of constant-gain learning, as discussed at length by Evans and Honkapohja [23], who examine various constant gain learning models as sources of persistent dynamics. Their models differ somewhat from ours in that they incorporate exogenous price shocks, while in our model the market price is determined endogenously through market-clearing from an exogenous dividend shock. In the Evans-Honkapohja framework, convergence to the rational expectations equilibrium with stochastic shocks can only happen with decreasing-gain learning ([23], p. 48). Our stochastic framework provides an alternate situation in which constant gain learning can frequently lead to convergence to equilibrium.

Least squares learning and Bayesian learning correspond to decreasing gains learning rules. Indeed, other numerical experiments we conducted suggest that convergence to equilibrium is stable for all positive values of risk aversion, in our framework, if the gains process corresponds to least squares learning (decreasing like $1/t$).

However, a constant gain learning rule, as pointed out in [23], is a prudent rule if structural market changes are expected in the future. It is also a reasonable rule for the following reason. If we denote by v_t the current price-dividend ratio estimate by agent 1, and v_t^m the ratio of the realized market-clearing price to the dividend d_t , then our learning rule can be restated as

$$v_t = v_{t-1} + \eta(v_{t-1}^m - v_{t-1}),$$

where η is our gain parameter. This can be solved recursively as

$$v_t = (1 - \eta)^t v_0 + \eta \sum_{i=0}^{t-1} (1 - \eta)^i v_{t-1-i}^m,$$

meaning that constant gain learning corresponds asymptotically to estimating v by means of an exponentially weighted average of historical lagged market values. This kind of forecasting is common in market practice.

Another note worth making concerns what we have learned about existence of equilibria. Our CEE-style equilibria are also REE-style equilibria, because endowing agents with complete market information after the CEE has been reached would not change their behavior, since they are already behaving optimally and agree on the true price forecast. Most of our equilibria are “corner” solutions in which one of the two agents asymptotically holds all the wealth and the other holds zero, corresponding to the homogenous agent case. However, Theorem 3 identifies an additional kind of equilibrium. When $\theta_1 = \theta_2$, it is possible that the agents have differing risk aversion parameters, but have the same optimal policy and co-exist at a no-trading equilibrium in which both agents hold a positive fraction of the supply of stock, agree on the pricing function (39), and invest and consume optimally in every period.

For the log-normal dividend growth process we considered above, from Equation (59) we can conclude that this 2-agent equilibrium takes place whenever

$$(1 - \gamma_1)\mu + \frac{1}{2}(1 - \gamma_1)^2\sigma^2 = (1 - \gamma_2)\mu + \frac{1}{2}(1 - \gamma_2)^2\sigma^2.$$

For any $\gamma_1 < 1$ there will be a “dual” value $\gamma_2 > 1$ satisfying this equation. Therefore any CRRA agent has a distinct “dual” agent for which the pair form a heterogeneous economy with nontrivial REE and explicit pricing and policy functions given by the formulas in Theorem 3.

As a final note, we remark that our results illustrate that chaos can be deceptive. Looking back at Figures 7 and 3, agents observing the apparent bifurcations might assume that the system parameters are changing, when in fact they are not. Instead, the economy is gradually falling into the basin of attraction of a chaotic attractor, but unpredictable behavior takes more than 100 time steps to begin to be visible. An observer of prices in the setting of Figure 7 cannot tell the difference, during the first 40 periods, between a gentle price trend and the early stages of a descent into a chaotic basin. This should caution observers of vastly more complex real economies to be wary of drawing long-term conclusions from observations of finite time series.

8. Acknowledgements

The authors are grateful to the referees for helpful comments and suggestions improving this paper.

References

- [1] P. T. Larson, Goldman pays the price of being big, *Financial Times* (August 13, 2007).
- [2] K. Dowd, J. Cotter, C. Humphrey, M. Woods, How unlucky is 25-sigma?, *Journal of Portfolio Management* 34 (4) (2008) 76–80.
- [3] Opinion, *The Economist* (July 16, 2009).
- [4] R. Lucas, In defense of the dismal science, *The Economist* (August 6, 2009).
- [5] J. Lucas, Robert E, Asset prices in an exchange economy, *Econometrica* 46 (6) (1978) 1429–45.
- [6] E. H. Kim, A. Morse, L. Zingales, What has mattered to economics since 1970, *The Journal of Economic Perspectives* 20 (4) (2006) pp. 189–202.
- [7] J. F. Muth, Rational expectations and the theory of price movements, *Econometrica* 29 (3) (1961) pp. 315–335.
- [8] E. F. Fama, Efficient capital markets: A review of theory and empirical work, *The Journal of Finance* 25 (2) (1970) pp. 383–417.
- [9] J. D. Farmer, J. Geanakoplos, The virtues and vices of equilibrium and the future of financial economics, *Complexity* 14 (3) (2009) 11–38.
- [10] J. Miao, Competitive equilibria of economies with a continuum of consumers and aggregate shocks, *Journal of Economic Theory* 128 (1) (2006) 274 – 298.

- [11] Z. Feng, J. Miao, A. Peralta-Alva, M. Santos, Numerical simulation of nonoptimal dynamic equilibrium models, Tech. Rep. 0912, University of Miami, Department of Economics (February 2009).
- [12] D. S. Graça, Noncomputability, unpredictability, and financial markets, *Complexity* 17 (6) (2012) 24–30.
- [13] Y.-F. Kao, V. Ragupathy, K. V. Velupillai, S. Zambelli, Noncomputability, unpredictability, undecidability, and unsolvability in economic and finance theories, *Complexity* 18 (1) (2012) 51–55.
- [14] W. A. Branch, G. W. Evans, Learning about risk and return: A simple model of bubbles and crashes, *American Economic Journal: Macroeconomics* 3 (3) (2011) 159–191.
- [15] C. H. Hommes, Heterogeneous agent models in economics and finance, in: L. Tesfatsion, K. Judd (Eds.), *Agent-Based Computational Economics*, Vol. 2 of *Handbook of Computational Economics*, Elsevier, 2006, Ch. 23, pp. 1109 – 1186.
- [16] G. Bottazzi, P. Dindo, Evolution and market behavior with endogenous investment rules, LEM Working Paper Series 2010/20, Scuola Superiore Sant’Anna, Pisa, Italy (January 2011).
- [17] C. Chiarella, R. Dieci, L. Gardini, Speculative behaviour and complex asset price dynamics: a global analysis, *Journal of Economic Behavior & Organization* 49 (2) (2002) 173 – 197.
- [18] L. Neuberg, K. Bertels, Heterogeneous trading agents, *Complexity* 8 (5) (2003) 28–35.
- [19] C. Chiarella, R. Dieci, L. Gardini, Asset price and wealth dynamics in a financial market with heterogeneous agents, *Journal of Economic Dynamics and Control* 30 (2006) 1755 – 1786.

- [20] W. A. Brock, C. H. Hommes, Heterogeneous beliefs and routes to chaos in a simple asset pricing model, *Journal of Economic Dynamics and Control* 22 (1998) 1235 – 1274.
- [21] K. Adam, A. Marcet, Internal rationality, imperfect market knowledge and asset prices, *Journal of Economic Theory* 146 (3) (2011) 1224 – 1252.
- [22] C. Hommes, G. Sorger, Consistent expectations equilibria, *Macroeconomic Dynamics* 2 (1998) 287–321.
- [23] G. W. Evans, S. Honkapohja, *Learning and Expectations in Macroeconomics*, Princeton University Press, Princeton, New Jersey, 2001.
- [24] K. Nishimura, G. Sorger, Optimal cycles and chaos: A survey, *Studies in Nonlinear Dynamics & Econometrics* 1 (1) (1996) 11–28.
- [25] J. Benhabib, On cycles and chaos in economics, *Studies in Nonlinear Dynamics & Econometrics* 1 (1) (1996) 1–2.
- [26] B. LeBaron, Short-memory traders and their impact on group learning in financial markets, *Proceedings of the National Academy of Sciences of the United States of America* 99 (Suppl 3) (2002) 7201–7206.
- [27] D. Duffie, *Security Markets: Stochastic Models*, Academic Press, 1988.
- [28] Y. Guan, Asset market dynamics of heterogeneous agent models, Ph.D. thesis, Florida State University (July 2011).
- [29] I. V. Evstigneev, T. Hens, K. R. Schenk-Hoppé, Globally evolutionarily stable portfolio rules, *Journal of Economic Theory* 140 (1) (2008) 197 – 228.
- [30] E. Asparouhova, P. Bossaerts, N. Roy, W. Zame, 'Lucas' in the laboratory, Working Paper 19068, National Bureau of Economic Research (May 2013).
- [31] P. Bossaerts, W. Zame, Risk aversion in laboratory asset markets, in: J. C. Cox, G. W. Harrison (Eds.), *Risk Aversion in Experiments*, Vol. 12 of *Research in Experimental Economics*, JAI Press, 2008, pp. 341–358.

- [32] C. A. Holt, S. K. Laury, Risk aversion and incentive effects, *American Economic Review* 92 (5) (2002) 1644–1655.

Bifurcation plot for Price-Dividend ratio v

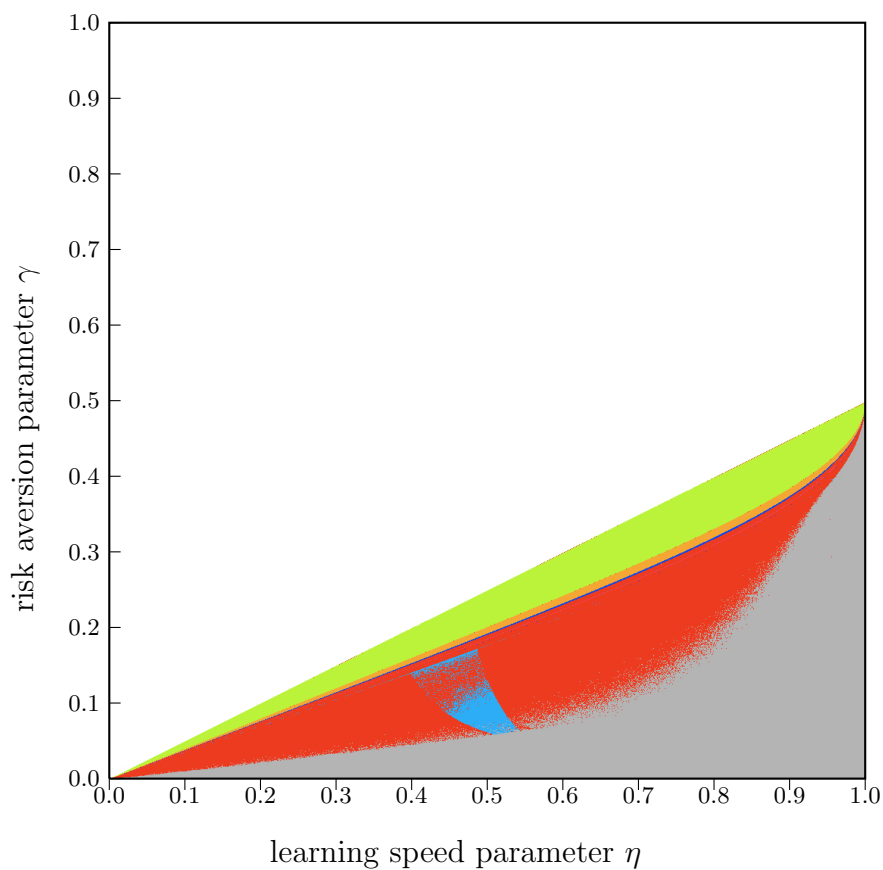


Figure 1: This bifurcation plot of the price-dividend kernel v summarizes the asymptotic v dynamics as a function of the risk aversion parameter γ and the learning speed parameter η . Each point is colored to indicate the kind of attractor we observe. *Color dictionary:* white: the price-dividend ratio converges to a fixed point (equilibrium); light green: converges to attracting period 2 orbit; orange: period 4; dark blue: period 8; purple: period 6; light blue: period 5; red: bounded, either chaos or period > 8 ; grey: divergent. The initial conditions used are $s_0 = 0.4$, $v_0 = 30$, with parameters $\beta = 0.95$, $\mu = 0.01752$, $\sigma^2 = 0.00123$.

Bifurcation plot for Holdings s




Figure 2: Bifurcation plot of the stock holding variable s . *Color Dictionary:* white: converges to $s^* = 1$ (agent 2 dies); red: s remains bounded below 1 (both agents survive); grey: price diverges to infinity. Parameters as in Figure 1.




Figure 3: The graphs of price-dividend ratio and stock holding when $\gamma_1 = 0.18$ and $\eta = 0.6$.
 A single path simulation with $v_0 = 19$, $s_0 = 0.4$, time step 0 – 500, β, μ, σ^2 as before.




Figure 4: The graphs of price-dividend ratio and stock holding when $\gamma_1 = 0.18$ and $\eta = 0.6$.
 A single path simulation with $v_0 = 19$, $s_0 = 0.4$, time step 5000 – 5500, β, μ, σ^2 as before.




Figure 5: The graphs of price-dividend ratio and stock holding when $\gamma_1 = 0.22$ and $\eta = 0.8$.
 A single path simulation with $v_0 = 25$, $s_0 = 0.7$, time step 0 – 500, β, μ, σ^2 as before.




Figure 6: The graphs of price-dividend ratio and stock holding when $\gamma_1 = 0.22$ and $\eta = 0.8$.
 A single path simulation with $v_0 = 25$, $s_0 = 0.7$, time step 5000 – 5500, β, μ, σ^2 as before.




Figure 7: The graphs of price and stock return when $\gamma_1 = 0.18$ and $\eta = 0.6$. A single path simulation with $v_0 = 19$, $s_0 = 0.4$, time step 0 – 100, β, μ, σ^2 as before.




Figure 8: The graphs of price and stock return when $\gamma_1 = 0.22$ and $\eta = 0.8$. A single path simulation with $v_0 = 25$, $s_0 = 0.7$, time step 0 – 100, β, μ, σ^2 as before.




Figure 9: Plots of the attractor for $\gamma = 0.15$ and various values of the learning parameter η . Each figure shows the attractor by plotting the orbit from $t = 5000$ to $t = 6000$ in order to remove initial transients. In all cases the initial values were $s_0 = 0.4$ and $v_0 = 30$. Here we see the attractors are fixed points or periodic orbits.

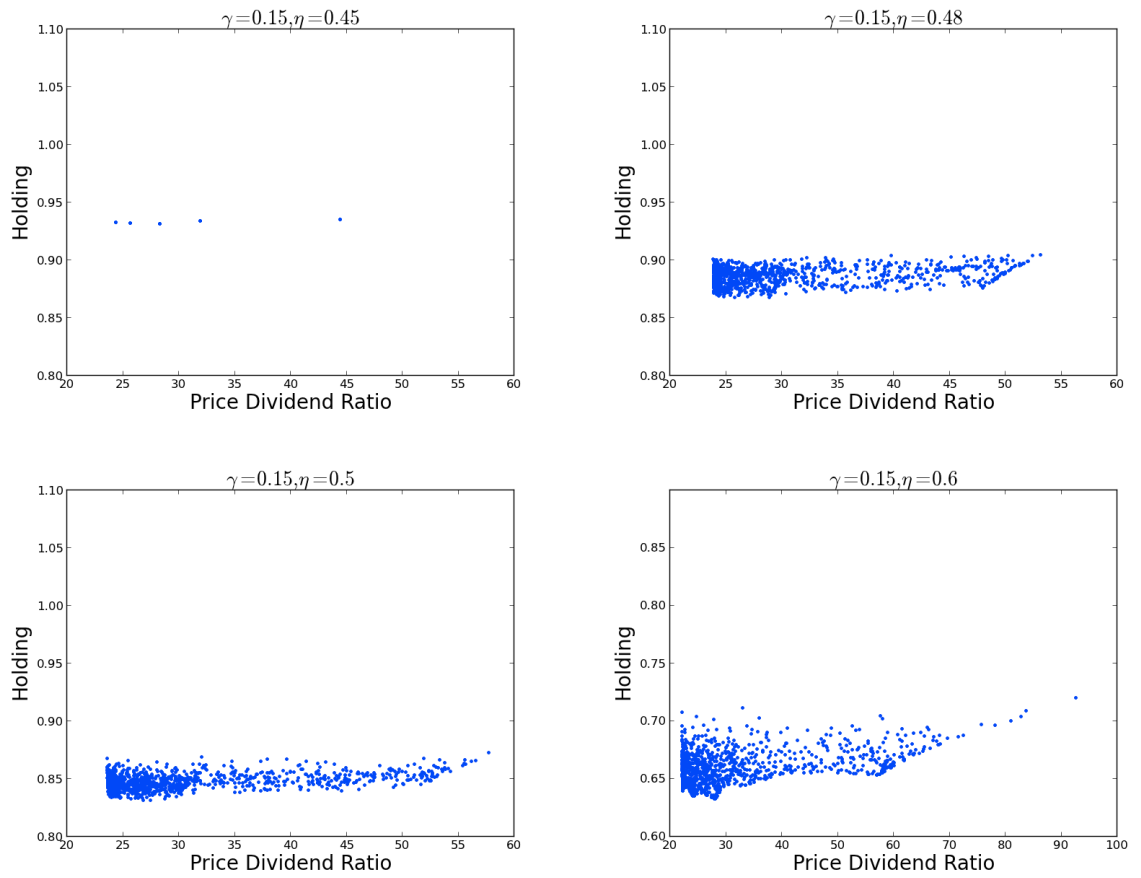


Figure 10: Plots of the attractor for $\gamma = 0.15$ and various values of the learning parameter η larger than those of the previous figure. Each figure shows the attractor by plotting the orbit from $t = 5000$ to $t = 6000$ in order to remove initial transients. In all cases the initial values were $s_0 = 0.4$ and $v_0 = 30$. Here we see the attractors become complex, indicating chaos.




Figure 11: The orbit diagram for the price-dividend ratio function in the system and using $\eta = 0.5$ and initial seed $v_0 = 30$ and $s_0 = 0.4$. A point is plotted above a particular value of risk aversion γ if it is an attracting value of v for the parameters $(\gamma, 0.5)$. For $\gamma > 0.25$ we begin with an attracting fixed point, and as γ decreases we bifurcate to more complex attractors.