

Asset Market Dynamics in Equilibrium Models with Heterogeneous Agents: Analytical Results

P. M. Beaumont^{1,*}, A. J. Culham², A. N. Kercheval³

¹Department of Economics, Florida State University, Tallahassee, Florida, 32306, United States

²Constellation Energy, Baltimore, Maryland, United States

³Department of Mathematics, Florida State University, Tallahassee, Florida, 32306, United States

*Corresponding Author: beaumont@fsu.edu

Copyright ©2013 Horizon Research Publishing All rights reserved.

Abstract We examine market dynamics in a discrete-time, Lucas-style asset-pricing model with heterogeneous, utility-optimizing agents. Finitely many agents trade a single asset paying a stochastic dividend. All agents know the probability distribution of the dividend but not the private information such as wealth and asset holdings of other agents. The market clearing price is determined endogenously in each period such that supply always equals demand.

Our aim is to determine whether and how the pricing function evolves toward equilibrium. In the special case where all agents have logarithmic utility, but possibly different holdings and discount factors, we completely describe the market dynamics, including the evolution of the pricing and demand functions, and asset holdings of the agents. The market converges to a stable equilibrium state where only the most patient agents remain, and the equilibrium pricing function is the same as the one arising in the simple homogeneous setting.

Keywords Market Dynamics, Market Disequilibrium, Correct Expectations Equilibrium

1 Introduction

In a general equilibrium asset pricing model, the existence of a rational expectations equilibrium (REE) implies that each trader solves a utility-optimization problem that incorporates all current information about the market. In particular, traders implicitly must: (i) have full knowledge of the problem including knowing the preferences and holdings of all the other agents, (ii) be able to deduce their optimal behavior no matter how complex, and (iii) share common expectations including that all the other agents are themselves rational and know all these things. In a sense, these agents must be “super-rational”: they are “experienced masters” in a stable world where the past is a good indicator of the future, and full rationality of all agents is common knowl-

edge [1, ch. 1].

Access to full information has the advantage of making the REE, in many cases, unique and computationally feasible [2]. However, real traders seldom know this much, because real markets are always changing unpredictably. Since super-rational agents cannot evolve, we also cannot use the REE concept to study market states away from equilibrium, or the dynamics of convergence toward equilibrium. Important prior work on REE-style equilibrium models, including [3, 4, 5], has naturally focused solely on the equilibrium states themselves and do not illuminate the nearby disequilibrium states and their evolution, which are typical initial states for boundedly-rational agents.

The agent-based computational economic (ACE) modeling literature, where agents follow formal, often myopic, rules, which can evolve by natural selection takes a more general, econophysical view of market dynamics [6, 7, 8]. This literature examines the resulting emergent properties of the market, including the dynamics of prices and holdings as the market evolves and agents with poor rules die and those with good rules prosper. These systems are typically studied by simulation. An excellent and recent review of this literature is provided in [9].

In this paper we study a discrete-time, general equilibrium Lucas tree economy with one asset and $N > 1$ infinitely lived agents who act to optimize an infinite-horizon utility. The asset (stock) is in fixed supply and pays a random dividend in each period. The asset price is determined endogenously in each period such that total (price-dependent) demand is equal to total supply. The dividend is assumed to follow a Markov process whose distribution is known to the agents, but we need make no further assumptions about it – its distribution may be discrete or continuous.

Each agent’s income (available for consumption or investment) in each period is just the stock dividend multiplied by the number of shares held that period. In this sense we have a pure Lucas-style economy with no exogenous income stream. The stock price, and therefore wealth of agents, is an “emergent” property coming from the joint behavior of all the agents.

Agents determine how much to consume and how much to invest in the stock each period by optimizing a time-separable constant relative risk aversion (CRRA) utility function. As in the standard Lucas model, at any time, agent i , $1 \leq i \leq N$, holds a certain number of shares of stock $s_i \geq 0$, which can vary over time if the agent trades. (Short positions are not allowed.) Agents also have discount factors β_i and risk-aversion parameter γ_i . All of which are allowed to vary across the population of agents.

Our agents are “boundedly rational” in the specific sense that they do not know the holdings or parameters of the other agents. Therefore they cannot incorporate that information into the solution of their personal optimization problems. As we will see, this allows us to model the economy both at and away from equilibrium.

To compute their optimal asset demand and consumption functions each agent must solve an expectation involving future market-clearing prices. Tomorrow’s market clearing price P' depends on tomorrow’s unknown dividend D' (the only source of randomness in our model). While agents know the probability distribution of D' , they do not know how the market clearing price $P(D')$ depends on D' since this depends in part on the behavior of other agents.

Therefore, solving for the demand equations requires agents, individually, to hypothesize “best guess” market-clearing pricing functions $p_i(D')$. This allows them to compute optimal demands, conditional on $p_i(\cdot)$, and to trade if indicated. The market will actually clear today at a price $P_m(D)$, such that supply equals demand. (Here we make explicit the dependence of $P_m(\cdot)$ on the dividend D , but in general, in disequilibrium, this function also depends on all the agents’ holdings, parameters, and pricing functions p_i [2].)

Using the terminology from [1], we will say that the market is at a “correct expectations equilibrium” (CEE), if all agents are using the same pricing function $P_m(\cdot)$ to solve for optimal demand, and this pricing function is correct in the sense that the actual resulting market clearing price is $P_m(D)$ in every period. This framework has the advantage that agent behavior is well-defined whether or not the market is at equilibrium.

We may now ask whether CEE exist in our framework and, if so, whether they are unique and how we can characterize market dynamics. Specifically, how do the asset holdings and market clearing pricing function P_m evolve over time, *both at and away* from equilibrium?

In general, we do not expect to find closed formulas describing the dynamics of such markets and so would investigate these questions via numerical simulation, as is common in the ACE literature. However, we can give rigorous and complete answers to these questions in the special (but reasonable) case where all agents have log utility ($\gamma_i = 1$).

In summary, we will find that there is a unique CEE pricing function given by

$$P^*(D) = \frac{\beta}{1 - \beta} D, \quad (1)$$

where β is the discount factor of the most patient agent. CEE stock holdings are zero for agents with discount factor less than β , and can be distributed arbitrarily among

agents with discount factor equal to β . This result is consistent with the long-known principle (e.g. [3]) for REE equilibria that the most patient agent eventually holds all the wealth. Our work confirms this principle in the CEE boundedly rational context typical of ACE and furthermore we provide a formal proof of convergence to the familiar pricing function well-known in the homogeneous case.

Finally, we formally prove that asset holdings and market pricing functions evolve according to explicit formulas given below (equations (11), (12)). The market dynamics do not depend on the individual pricing functions p_i , on the learning mechanism used, if any, or on the probability distribution of dividends. Each CEE is a fixed point of the dynamics. From any disequilibrium initial condition, markets will converge deterministically and exponentially fast to a CEE given by equations (1) and (15). Therefore we are able to use methods of deterministic dynamics to analyze the evolution in this paper, rather than the statistical methods used in many ACE models with stochastic inputs.

Since [3] there has been a renewed interest in the role of heterogeneous discount rates in determining market dynamics. In his 2011 Presidential address to the American Finance Association, Cochrane [10] emphasizes the importance of discount rates in discussing many of the “puzzles” in asset pricing. The role of discount rate heterogeneity in explaining wealth inequality is discussed by [11], [12, 13] focus specifically on heterogeneity of time preferences and asset prices, and [14] considers the effects of heterogeneous patience on the terms structure of interest rates. An interesting example of how heterogeneous rates of time preference can help describe personal behaviors such as the decision to smoke is given in [15]. Our work contributes to this literature by adding a formal mathematical analysis of the dynamical behavior of asset markets in a general equilibrium model with heterogeneous agents.

2 Model framework

Consider the standard Lucas asset pricing model [16, 17] with N possibly heterogeneous agents and a single risky asset with period t market clearing spot price denoted P_t . The number of shares of the asset is normalized to be N and the asset pays a random dividend D_t per share, described by some Markov process and so determined solely by the observed state of the world at the beginning of each period t . All agents are assumed to know the distribution of dividend payments across states, which can be arbitrary, discrete or continuous, stationary or not.

There is no production in this economy so in time period t agent i will choose optimal consumption $c_{i,t} \geq 0$ and investment in the asset $s_{i,t+1} \geq 0$ based upon the agent’s preferences and period budget constraint

$$c_{i,t} + P_t s_{i,t+1} \leq w_{i,t} = (P_t + D_t) s_{i,t}, \quad \text{for all } t, \quad (2)$$

where $w_{i,t}$ is the agent’s period t wealth.

The order of events is as follows. Agent i begins period t knowing her asset holdings $s_{i,t}$. Next, today’s stock dividend D_t is announced to all agents. At this point the

agent does not yet know her current wealth because the market price P_t has yet to be determined through market clearing. Each agent must first compute her optimal demand as a function of price, $s_i(P)$, representing the optimal number of shares demanded at any given price P . These demand functions then collectively determine the unique price P_t that clears the market according to

$$\sum_{i=1}^N s_i(P_t) = N. \tag{3}$$

The actual mechanism of market clearing is not important; we can imagine that there is some market maker who collects the demand functions from each agent and publicly declares the market price satisfying (3). (This price will be unique if, for example, demands are all monotone in P .) This declared market price is the only source of information about the other agents.

Each agent is assumed to be an expected utility maximizer with constant relative risk aversion preferences and risk aversion $\gamma_i > 0$. The one-period utility of consumption is

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

or

$$u_i(c) = \log c, \quad \gamma_i = 1,$$

and consumption and investment polices must be found that optimize the time-separable expected utility

$$\max_{\{c_{i,\tau}, s_{i,\tau+1}\}_{\tau=t}^{\infty}} E_t \sum_{\tau=t}^{\infty} \beta_i^{\tau-t} u_i(c_{i,\tau}) \tag{4}$$

subject to the budget constraint (2). The agents' discount factors, $\beta_i \in (0, 1)$, may differ and the expectation in (4) is over the distribution of dividends and is conditional upon the information available to the agent at the beginning of period t .

Denote by $\mathbf{s}_t \in \mathbb{R}^N$ the vector of all agents' time- t asset holdings. We explicitly include the time subscript to emphasize that these holdings may change over time. We describe two different equilibrium concepts which we label REE and CEE.

Definition 1 A rational expectations equilibrium (REE) for this economy consists of an aggregate pricing function $P^*(D_t, \mathbf{s}_t)$ for the risky asset and a set of agent consumption demand functions $c_{i,t} = c_i(D_t, \mathbf{s}_t)$ and asset demand functions $s_{i,t+1} = s_i(D_t, \mathbf{s}_t)$ such that, for all future times t , the asset market clears at $P_t = P^*(D_t, \mathbf{s}_t)$, the budget constraint is satisfied for each agent, and the demand functions solve the agents' optimization problems (incorporating full knowledge of the market).

Our boundedly rational agents, however, do not know \mathbf{s}_t , and must optimize based upon more limited knowledge. Therefore we use instead the notion of a CEE:

Definition 2 A correct expectations equilibrium (CEE) for this economy consists of an aggregate pricing function $P^*(D_t)$ for the risky asset and a set of agent consumption demand functions $c_{i,t} = c_i(D_t, s_{i,t})$

and asset demand functions $s_{i,t+1} = s_i(D_t, s_{i,t})$ such that, for all future times t , the asset market clears at $P_t = P^*(D_t)$, the budget constraint is satisfied for each agent, and the demand functions solve the agents' optimization problems (incorporating available knowledge).

To clarify the difference between "available knowledge" and "full knowledge", we need to look at the optimization problem more carefully.

Solving the Bellman equation from the usual recursive dynamic programming approach leads to the following standard Euler equation as viewed from time t , which must be satisfied by the optimal consumption demand function $c_{i,t+1}$:

$$P_t = E_t \left[\beta_i \left(\frac{c_{i,t+1}}{c_{i,t}} \right)^{-\gamma_i} (P_{t+1} + D_{t+1}) \right]. \tag{5}$$

Using the budget constraint (2) to eliminate c , and using the notation $s'_{i,t} = s_{i,t+1} = s_i(D_t, s_{i,t})$, and $s''_{i,t} = s_i(D_{t+1}, s_{i,t+1}) = s_i(D_{t+1}, s_i(D_t, s_{i,t}))$, we may rewrite (5) as

$$P_t = \beta E_t \left[\left(\frac{s_{i,t}(P_t + D_t) - s'_{i,t}P_t}{s'_{i,t}(P_{t+1} + D_{t+1}) - s''_{i,t}P_{t+1}} \right)^{\gamma_i} (P_{t+1} + D_{t+1}) \right]. \tag{6}$$

Each agent, at equilibrium, must solve for the asset demand function $s_i(D, s)$ that satisfies this optimality condition.

Now our boundedly rational agents will have a problem computing the conditional expectation in the Euler equation because they need to know the equilibrium pricing function P^* in order to know the probability distribution of tomorrow's prices which depends on the holdings and preferences of the other agents. Worse, it depends on knowing that the other agents also know the equilibrium pricing function, which will not be true in disequilibrium.

For illustration, consider a simple case where agents are super-rational and the REE is computable. Suppose all agents are identical with log utility ($\beta_i = \beta$, $\gamma_i = 1$, $s_{i,t} = 1$, for all i and t) and all agents are aware of this. Knowing they are identical, agents can deduce that there will never be any trading and so the budget constraint will thus imply that $c_{i,t} = D_t$ and $s_{i,t+1} = s_{i,t} = 1$ for each agent. The Euler equation (5) simplifies to

$$P_t = E_t \left[\beta \left(\frac{D_t}{D_{t+1}} \right) (P_{t+1} + D_{t+1}) \right] \tag{7}$$

which is now easy to solve. The REE aggregate pricing function satisfying this equation is

$$P^*(D_t) = \frac{\beta}{1 - \beta} D_t.$$

and holdings $\mathbf{s} = (s_1, \dots, s_N)$ remain fixed over time.

Note that the demand function s' has vanished from the Euler equation, since demand s' is fixed at s . Furthermore, agents know *a priori* that the market is in

equilibrium so the REE pricing function, and hence the market clearing price, is known by all agents as soon as the dividend is announced so there is no need to compute demand functions for a range of prices. This is a significant simplification of the problem, but its operational validity is questionable: only if the agents know that all other traders are identical to themselves can they justify setting $c_{i,t} = D_t$ in their Euler equation [18].

In our boundedly rational world, agents do not know holdings or preferences of the other agents, nor what pricing functions the other agents are using. Therefore they cannot know whether the market is at equilibrium. They can only know whether their own pricing function is correctly predicting market clearing prices.

Our way forward is to allow agents to use private estimates $p_i(D_{t+1})$ of the period $t + 1$ aggregate pricing function in their Euler equations (5). We assume that these functions depend upon D alone. Although agents are aware that there are unobserved variables \mathbf{s}_t and functions $\{p_j\}_{j \neq i}$ influencing market prices away from equilibrium, we can imagine they implicitly operate under the hypothesis that the market will converge to a CEE, so that the dependence on these variables will vanish with time. Thus, it is pragmatic for the agents to use $p_i(D_{t+1})$ in the Euler equation.

At a CEE, the demand functions s_i depend only on current holdings $s_{i,t}$ and the current dividend D_t , because the asset price P_t is implicitly determined by D_t . Away from equilibrium this is no longer the case. In order for a market clearing price to be determined according to equation (3), each agent's demand function must include price P_t as an explicit independent variable. In other words, *agents must determine how much to demand at any potential price*. These demand equations will then determine the current period's market clearing price, and, in turn, each agent's actual realized investment.

This framework now provides agents with enough information to derive optimal consumption and asset demand functions $c_{i,t} = c_i(D_t, s_{i,t}, P | p_i)$ and $s_{i,t+1} = s_i(s_{i,t}, D_t, P | p_i)$, where we indicate explicitly the dependence on the choice of pricing function p_i , which is allowed to change over time, e.g. via some learning mechanism.

Now we can make precise the spaces and dynamics we are studying.

- Let \mathbb{D} be the set of possible dividends and \mathbb{P} the set of possible asset prices; in this case both can be taken to be the set \mathbb{R}^+ of nonnegative real numbers.
- \mathcal{P} will denote the set of all possible pricing functions $p : \mathbb{D} \rightarrow \mathbb{P}$. This can be thought of as a space of random processes depending on the underlying dividend process D_t .
- $\mathcal{S} \subset \mathbb{R}^N$ is the set of all possible holdings vectors $\mathbf{s} = (s_1, \dots, s_N)$
- $\mathcal{A} = \mathcal{P}^N$ is the set of possible vectors of agents' individual pricing functions (p_1, \dots, p_N) .

The market clearing function $M : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{P}$, given the agents' holdings and pricing functions, returns the

resulting market clearing pricing function implied via (3) by the resulting agents' optimal demand functions.

The *market dynamical state space* is

$$X = \{(\mathbf{s}, p, \mathbf{a}) \in \mathcal{S} \times \mathcal{P} \times \mathcal{A} : M(\mathbf{s}, \mathbf{a}) = p\}.$$

An element of X specifies asset holdings \mathbf{s} , individual pricing functions \mathbf{a} , and the resulting market pricing function p . (Note we suppress the space \mathbb{D} , since the underlying dividend process is exogenous; effectively we are studying the dynamics of random processes measurable with respect to the filtration generated by $\{D_t\}$.)

The *market dynamical system* is the mapping $f : X \rightarrow X$ corresponding to updating the state variables by one time step. The holdings vector \mathbf{s} will change due to trading and the vector \mathbf{a} of agent pricing functions will change due to the particular learning mechanism chosen, if any. The passage of time is tracked by the trajectory of an f -orbit $\{f^n(x) : n = 0, 1, 2, 3, \dots\}$ of an initial state $x \in X$, where f^n denotes the n -fold composition of f with itself.

In this framework, a CEE corresponds to a state $(\mathbf{s}, p^*, \mathbf{a})$ where $\mathbf{a} = (p^*, p^*, \dots, p^*)$ and $f(\mathbf{s}, p^*, \mathbf{a}) = (\mathbf{s}', p^*, \mathbf{a})$ for some \mathbf{s}' . If there is no trading, then $(\mathbf{s}, p, \mathbf{a})$ will be a fixed point of f .¹

We will now focus on a special case where it is possible to derive tractable analytical results.

3 The log utility case

Suppose all $N > 1$ agents have CRRA utility with a common risk aversion, $\gamma_i = \gamma > 0$ for all i , but possibly differing discount factors β_i and initial holdings $s_{i,0}$.

We assume the aggregate supply of stock is N shares, and that agents have limited information in the sense that they may not make any assumptions about the preferences or holdings of the other agents when solving for their optimal consumption in each time period. Agents are presumed to have private estimates $p_i(\cdot)$ of the pricing function, but we impose no assumptions yet on what these are or how they evolve.

Using the budget constraint to substitute for c , the i^{th} agent's Euler equation (5) is given by the following equation, where for convenience we use the notation $D = D_t$ and $D' = D_{t+1}$:

$$\frac{P}{(s_{i,t}(P + D) - P s_i(D, s_{i,t}, P))^{\gamma_i}} = \beta_i E_t \left[\frac{(p_i(D') + D')}{(s_i(D, s_{i,t}, P)(p_i(D') + D') - p_i(D') s_i(D', s_i(D, s_{i,t}, P), p_i(D')))^{\gamma_i}} \right]. \quad (8)$$

Here, P represents any current market price, $s_{i,t}$ is the time- t number of shares of stock of agent i , p_i is the i^{th} agent's current pricing function, and $s_i(D, s, P)$ is the i^{th} agent's optimal demand function, which is the unknown here. The agent's optimization problem is to find a function s_i satisfying (8) for all possible values of

¹If there is perpetual trading at equilibrium, none of these points will be fixed, but we could view a CEE as the smallest f -invariant set containing such a point.

P ; the actual market price will be determined after all agents have done this.

Notice that the unknown demand function s_i appears in (8) in a highly nonlinear way. Nonetheless, agents must solve this equation for s_i in order to be able to participate in the implicit price-calling auction used to arrive at the market clearing price. (Though this poses a computational problem in simulations, we assume our agents can solve this mathematical problem at no cost.)

In the log-utility case ($\gamma_i = 1$), it is easy to verify that (8) has the simple explicit solution

$$s_i(D, s, P) = \beta_i(1 + (D/P)) s. \quad (9)$$

Unlike the general case $\gamma_i \neq 1$, this solution has the special property that it does not depend upon the agent's pricing function p_i . This is a very important feature of log utility and is what makes this special case analytically tractable.

Because the pricing functions p_i do not influence agent behavior, the specific learning mechanism used by the agent to update her pricing function guess is rendered irrelevant, and the dynamical system f projects down to a mapping

$$g : Y \rightarrow Y$$

on the *reduced dynamical state space*

$$Y = \{(s, p) \in \mathcal{S} \times \mathcal{P} : (s, p, \mathbf{a}) \in X \text{ for some } \mathbf{a} \in \mathcal{A}\},$$

consisting of just asset holdings and market clearing price functions.

The market clearing price, which we now denote $P_m = P_m(D_t)$, is determined from the market clearing condition

$$\sum_{i=1}^N s_i(s_{i,t}, D_t, P_m) = N. \quad (10)$$

Substituting the demand function (9) for s_i and solving for the market clearing price, gives

$$P_m = \frac{\sum_1^N \beta_j s_j}{N - \sum_1^N \beta_j s_j} D. \quad (11)$$

Substituting (11) into the demand function (9) gives the agent's next period holdings at market clearing prices as

$$s_{i,t+1} = s_i(s_{i,t}, D, P_m) = \beta_i \left(\frac{N}{\sum_1^N \beta_j s_{j,t}} \right) s_{i,t}. \quad (12)$$

The evolution of holdings has now become a deterministic dynamical system $h : \mathcal{S} \rightarrow \mathcal{S}$ given by $h = (h_1, \dots, h_N)$ where

$$h_i(s_{1,t}, \dots, s_{N,t}) = \beta_i \left(\frac{N}{\sum_1^N \beta_j s_{j,t}} \right) s_{i,t}. \quad (13)$$

At any time, the market pricing function P_m is determined by the asset holdings according to (11), so the market dynamical system is completely described by h . We now need only determine the behavior of h -orbits on \mathcal{S} in order to completely understand the time evolution of our market.

The following theorem establishes that this dynamical system converges and reports the limiting asset holdings and pricing function.

Theorem 1 Consider a pure exchange economy of N infinitely-lived agents and N shares of a single risky asset paying stochastic dividend D at the beginning of each period. Each agent maximizes her discounted, expected life-time utility subject to the period budget constraint $c_i + P s'_i \leq (P + D) s_i$. All agents have log utility and have discount factors β_i and initial asset holdings s_i° , where $\sum_{i=1}^N s_i^\circ = N$. Agents know the probability distribution of dividends but not the asset holdings, discount factors, or utility functions of other agents.

For convenience, order the agents by decreasing discount factor and let k be the number of agents who share the maximum discount factor β , so that

$$1 > \beta = \beta_1 = \dots = \beta_k > \beta_{k+1} \geq \dots \geq \beta_N > 0.$$

Then the dynamic behavior of holdings and market clearing prices is given by the deterministic equations (11) and (12). This system converges exponentially fast to

$$P^*(D) = \frac{\beta}{1 - \beta} D \quad (14)$$

and

$$s_i^* = \frac{N s_i^\circ}{s_1^\circ + \dots + s_k^\circ}, \quad i \leq k, \quad (15a)$$

$$= 0, \quad i > k. \quad (15b)$$

The theorem states that the asset holdings of all agents with less than the maximum subjective discount factor converge to zero at an exponential rate. The asset holdings of the remaining most patient agents, with the highest discount factor, converge to a limit proportional to the initial holdings of this subset of agents. The patience of these agents is eventually rewarded by accumulating all of the wealth in the economy while the impatient agents are driven out of the market as their wealth is asymptotically driven to zero. Furthermore, the economy eventually collapses to a set of agents with differing holdings but a common discount factor. The market clearing price globally converges to the pricing function obtained in the classical and more restricted case of homogeneous, super-rational agents.

In the special case where agents (unknowingly) have identical discount factors but possibly different initial holdings, there is never any trading and the market clears immediately in the first time step at the familiar rational expectations equilibrium price

$$P_m = \frac{\beta}{1 - \beta} D. \quad (16)$$

3.1 Proof of Theorem 1

Proof: It is easy to verify algebraically that the demand function (9) solves the Euler equation (8), and therefore (11) and (12) describe the market clearing price and new stock holdings in each time step.

Also, it is easy to see that the market clearing price P_m is given by the P^* in (14) if the stock holdings are such that the only non-zero holdings are for agents with $\beta_i = \beta$. Therefore it remains to prove that holdings globally converge to the values described by (15).

It is convenient to rewrite the dynamical system (13) in terms of the relative holdings $x_i = s_i/N$:

$$x'_i = \frac{\beta_i x_i}{\sum_1^N \beta_j x_j}. \quad (17)$$

Here $x_j \in [0, 1]$ for all j and $\sum_j x_j = 1$, so the state (x_1, \dots, x_N) lies on the $(N-1)$ -dimensional unit simplex

$$\Delta^N = \{(x_1, \dots, x_N) \geq 0 : \sum_i x_i = 1\} \quad (18)$$

in the positive orthant of \mathbb{R}^N . Since $\sum_j x'_j = 1$, we can describe the dynamics as an iteration of the mapping $T : \Delta^N \rightarrow \Delta^N$ where the i th coordinate of $T(x)$ is defined to be x'_i given by (17).

If $\Gamma \subset \Delta^N$ denotes the k -dimensional simplicial face

$$\Gamma = \{(x_1, \dots, x_k, 0, \dots, 0) : \sum_{j=1}^k x_j = 1\}, \quad (19)$$

then it is easy to see that every point of Γ is fixed by T . From (17), if $x_j \neq 0$ and $\beta_i = \beta_j$, then $T(x_i)/T(x_j) = x_i/x_j$. Hence T always preserves the relative sizes of the coordinates x_1, \dots, x_k . Therefore the limiting holdings must be given by (15) if we can show that every forward T -orbit $\{T^n(x)\}$ converges to Γ .

Define $\pi_\Gamma : \Delta^N \rightarrow \Gamma$ to be the projection fixing the first k coordinates and setting the remaining $N-k$ coordinates to zero. Let $\Delta^{N+} = \{x \in \Delta^N : \pi_\Gamma(x) \neq 0\}$.

Lemma 2 *Define $F : \Delta^N \rightarrow \mathbb{R}$ by $F(x) = \sum_i \beta_i x_i$, where $\beta = \beta_1$ and the β_i are ordered as in Theorem 1. Then for any $x \in \Delta^{N+}$, $F(T^n(x))$ increases monotonically with limit β as $n \rightarrow \infty$.*

Proof: Let Υ denote the $(N-k)$ -dimensional sub-simplicial face

$$\Upsilon = \{(0, \dots, 0, x_{k+1}, \dots, x_N) : \sum_{j=k+1}^N x_j = 1\}, \quad (20)$$

and $\pi_\Upsilon : \Delta^N \rightarrow \Upsilon$ the projection fixing the last $N-k$ coordinates and setting the first k to zero.

$F(x)$ is simply a weighted average of the β 's, weighted by the x 's.

Using the definition of T , we have, for any x ,

$$F(T(x)) = \frac{\sum_i \beta_i^2 x_i}{F(x)} \quad (21)$$

and so $F(x)F(T(x)) = \sum_i \beta_i^2 x_i$. Also, $F^2(x) = (\sum_i \beta_i x_i)^2$.

Now $F(T(x)) \geq F(x)$ follows from Jensen's inequality

$$\phi\left(\sum_i \beta_i x_i\right) \leq \sum_i \phi(\beta_i) x_i \quad (22)$$

for the convex function $\phi(x) = x^2$. The inequality is strict if both $\pi_\Gamma(x)$ and $\pi_\Upsilon(x)$ are nonzero.

Fix $x \in \Delta^{N+}$. If $\pi_\Upsilon(x) = 0$ then $x \in \Gamma$, $T(x) = x$, and $F(x) = \beta$, so there is nothing further to prove. Hence suppose $\pi_\Upsilon(x)$ is nonzero. This means $\pi_\Gamma(T(x))$ and $\pi_\Upsilon(T(x))$ are also nonzero, so $F(T^n(x))$ is a strictly

monotone sequence bounded by β . It must therefore converge to its supremum, call it β^* .

Suppose $\beta^* < \beta$. By compactness of Δ^N , the sequence $\{T^n(x)\}$ has a convergent subsequence $y_k = T^{n_k}(x) \rightarrow x^* \in \Delta^N$, and by continuity of F , $F(x^*) = \beta^*$. By the definition of T , $(T^n(x))_i$ is monotone increasing for $i = 1, \dots, k$. Therefore $x^* \in \Delta^{N+}$. Since $F(x^*) < \beta$, $\pi_\Upsilon(x^*) \neq 0$, and so

$$F(T(x^*)) > F(x^*) = \beta^*. \quad (23)$$

However, we also have $F(T(y_k)) \leq \beta^*$, and since $y_k \rightarrow x^*$ this contradicts the continuity of F and T . Therefore we must have $\beta^* = \beta$. This completes the proof of the lemma.

□

From the lemma above, the limit of $F(T^n(x))$ is β for all $x \in \Delta^{N+}$. Since F is continuous and $F^{-1}(\beta) = \Gamma$, every forward T -orbit starting in Δ^{N+} must converge to Γ . From equation (12), we see that if holdings are close to zero for agents $j = k+1, \dots, N$, then we have, approximately,

$$s'_j = \frac{\beta_j}{\beta_1} s_j, \quad (24)$$

which gives us, asymptotically, an exponential rate of convergence to zero all $j > k$. This completes the proof of Theorem 1.

□

3.2 An Illustration with Two Agents

For the special case of two agents with different discount factors $\beta_1 > \beta_2$ it is possible to describe the market dynamics with a simple diagram. The asset demand functions are

$$s_i(s_{1,t}, s_{2,t}) = \beta_i \frac{2s_{i,t}}{\beta_1 s_{1,t} + \beta_2 s_{2,t}}, \quad i = 1, 2. \quad (25)$$

Using the market clearing constraint $s_{1,t} + s_{2,t} = 2$ gives

$$s_1(s_{1,t}) = \frac{\beta_1 s_{1,t}}{\beta_2 + (\beta_1 - \beta_2)(s_{1,t}/2)} \quad (26a)$$

and

$$s_2(s_{2,t}) = \frac{\beta_2 s_{2,t}}{\beta_1 + (\beta_2 - \beta_1)(s_{2,t}/2)}. \quad (26b)$$

These two functions are plotted in Figure 1 for the discount factors $\beta_1 = 0.95$ and $\beta_2 = 0.7$. A discount factor of β implies that an agent could only be induced to save rather than consume if the return on savings were at least $\frac{1-\beta}{\beta}$. Thus, a $\beta = 0.95$ implies a very patient agent who would accept a return of 5.26% on savings to forgo consumption whereas an agent with $\beta = 0.7$ would be very impatient and could not be induced to save unless returns were about 43%. Of course, we are not interested in calibrating our simple model to actual investors and we select these discount factors merely as illustrations for a numerical example. Theorem 1 proves that our example holds for any $\beta_1 > \beta_2$ no matter how close the values are although the effect is easier to illustrate for fairly small β_2 .




Figure 1. For a market with two agents having different discount factors $\beta_1 = 0.95$ and $\beta_2 = 0.7$, we plot the demand function of agent 1 (solid), and the diagonal (dashed). The demand function of agent 2 is symmetrically below the diagonal because of the market-clearing requirement $s_1 + s_2 = 2$. Iteration of agent 1's demand function, describing the passage of time, is illustrated with the arrows. Holdings for agent 1 converge to 2. Symmetrically, holdings for agent 2 converge simultaneously to zero.

Iteration of the upper function for agent 1, the most patient agent with the higher discount factor, is illustrated with the arrows showing that asset holdings will converge to $s = 2$. Similarly, the asset holdings of the less patient agent 2 will decrease monotonically to zero along the lower function. This behavior is common to any choice of discount factors as long as $\beta_1 > \beta_2$. When $\beta_1 = \beta_2$ both graphs are along the diagonal and the asset holdings of both agents remain fixed and there is no trading.

4 Conclusions

Theorem 1 and the surrounding discussion provide the details of the answers summarized in the Introduction. Since, in our log-utility case, the decisions of the agents turn out to be uninfluenced by the chosen pricing functions $\mathbf{a} \in \mathcal{A}$, the full dynamical system $f : X \rightarrow X$ conveniently reduces to a mapping $g : Y \rightarrow Y$ on the reduced state space. This allows us to completely specify the market dynamics both at and away from equilibrium.

Our model is fully consistent with a rational expectations general equilibrium framework and does not rely upon agents who ignore market fundamentals or behave irrationally such as in [20, 19]. At each period in our model markets clear according to (10) which sets the aggregate demand for assets equal to the aggregate supply. The market price is then determined endogenously according to (11) which depends upon the individual asset demand equations of each agent.

We are able to completely specify the asset market dynamics of this system as the economy moves toward

equilibrium. We show that the most patient agents will asymptotically accumulate all wealth with a convergence rate that depends upon the ratio of the discount factor of the less patient agent to the most patient agent. In an environment without idiosyncratic wealth shocks, this convergence will be monotonic. Although our conclusions are consistent with the existing literature we have provided a rigorous proof of these results.

It's natural to ask about what happens when agents are allowed to have other risk aversion parameters $\gamma_i > 0$. In general, optimal demands will then depend on the individual pricing functions \mathbf{a} as well as the wealth and asset holdings distributions [2]. Therefore, the mapping $f : X \rightarrow X$ will not be fully specified until the pricing function updating (learning) mechanism is specified.

Simulation studies are reported in [21] for this general case where the learning mechanism is a simple least squares updating of the pricing function based on the observed history of market prices. In all cases studied, market prices are observed to converge to a no trading equilibrium which, in our language, is a fixed point $(\mathbf{s}, P^*, \mathbf{a})$ of $f : X \rightarrow X$, where $\mathbf{a} = (P^*, \dots, P^*)$. These simulation studies are numerically challenging because each agent's optimal demand must be solved numerically in every time step.

The log utility case makes for an interesting comparison between our two equilibrium concepts, CEE and REE. Assuming that agents know all the pricing functions, holdings, and preferences of the other agents will not change their behavior in each time step. Equations (11) and (12) describe the market price $P_m(D, \mathbf{s})$ and demands as a function of D and \mathbf{s} , so the market is at an REE all all times, even as holdings evolve. When agents do not observe holdings \mathbf{s} , the pricing function

$P_m(D)$ evolves over time and reaches the equilibrium $P^*(D)$ only in the limit.

Acknowledgements

This work is partially supported by the FSU Council on Research and Creativity, IDS grant #014895. The authors are grateful for the helpful suggestions of the anonymous referee.

REFERENCES

- [1] Magill, M. and Quinzii, M. Theory of Incomplete Markets, volume 1. MIT Press, 1996.
- [2] Miao, J. Competitive equilibria of economies with a continuum of consumers and aggregate shocks. *Journal of Economic Theory*, Vol. 128, No. 1, 274 – 298, 2006.
- [3] Becker, R. A. On the Long-Run Steady State in a Simple Dynamical Model of Equilibrium with Heterogeneous Households. *Quarterly Journal of Economics*, Vol. 95, No. 2, 375–382, 1980.
- [4] Hernández D., A. and Santos, M. Competitive Equilibria for Infinite-Horizon Economies with Incomplete Markets. *Journal of Economic Theory*, Vol. 71, 102–130, 1996.
- [5] Magill, M. and Quinzii, M. Incomplete Markets Over an Infinite Horizon: Long-Lived Securities and Speculative Bubbles. *Journal of Mathematical Economics*, Vol. 26, 133–170, 1996.
- [6] Palmer, R. G., Arthur, W. B., Holland, J. H., LeBaron, B., and Tayler, P. Artificial Economic Life: A Simple Model of a Stockmarket. *Physica D: Nonlinear Phenomena*, Vol. 75, No. 1–3, 264–274, 1994.
- [7] Arthur, W. B., Holland, J. H., LeBaron, B., Palmer, R., Tyler, P. Asset Pricing Under Endogenous Expectations in an Artificial Stock Market. In: Arthur, W. B., Durlauf, S. N., Lane, D. A. (Eds.), *The Economy An an Evolving Complex System II*. Addison-Wesley, pp. 15–44, 1997.
- [8] Chiarella, C., Dieci, R., Gardini, L. Asset Price and Wealth Dynamics in a Financial Market with Heterogeneous Agents. *Journal of Economic Dynamics and Control*, Vol. 30, 1755–1786, 1998.
- [9] Chiarella, C., Dieci, R. and He, X-Z. Heterogeneity, Market Mechanisms, and Asset Price Dynamics. In: Thorsten Hens IV and Klaus Reiner Schenk-Hoppe, editors, *Handbook of Financial Markets: Dynamics and Evolution*, Ch. 5, 277–344, North-Holland Amsterdam, 2009.
- [10] Cochrane, J. H. Presidential address: Discount rates. *The Journal of Finance*, Vol. 66, No. 4, 1047–1108, 2011.
- [11] Hendricks, L. How important is discount rate heterogeneity for wealth inequality?. *Journal of Economic Dynamics and Control*, Vol. 31, No. 9, 3042–3068, 2009.
- [12] Gollier, C., Zeckhauser, R. Aggregation of heterogeneous time preferences. *Journal of Political Economy*, Vol. 113, No. 4, 878–896, 2005.
- [13] Lengwiler, Y., Malamud, S., Trubowitz, E. Asset pricing in heterogeneous economies. ETH Zurich Department of Mathematics Working Paper, 2005.
- [14] Lengwiler, Y. Heterogeneous patience and the term structure of real interest rates. *The American Economic Review*, Vol. 95, No. 3, 890–896, 2005.
- [15] Scharff, R. L., Viscusi, W. K. Heterogeneous rates of time preference and the decision to smoke. *Economic Inquiry*, Vol. 49, No. 4, 959–972, 2011.
- [16] Lucas, Jr, R. E., Asset prices in an exchange economy. *Econometrica*, Vol. 46, No. 6, 1429–1445, 1978.
- [17] Lyungqvist, L., Sargent, T. J., *Recursive Macroeconomic Theory*, 2nd Edition. The MIT Press, Cambridge, MA, 2004.
- [18] Frydman, R. Individual rationality and decentralization. In: Roman Frydman and E. S. Phelps, editors, *Individual Forecasting and Aggregate Outcomes: Rational Expectations Examined*, Ch. 5, pp. 97–122, Cambridge University Press, UK, 1984.
- [19] Guidolin, M., Timmermann, A. Properties of equilibrium asset prices under alternative learning schemes. *Journal of Economic Dynamics and Control*, Vol. 31, 161–217, 2007.
- [20] Jouini, E., Napp, C. Heterogeneous beliefs and asset pricing in discrete time: An analysis of pessimism and doubt. *Journal of Economic Dynamics and Control*, Vol. 30, 1233–1260, 2006.
- [21] Culham, A. J. Asset Pricing in a Lucas Framework with Boundedly Rational, Heterogeneous Agents. Ph.D. thesis, Florida State University, Department of Mathematics, Tallahassee, Florida, USA, 2007.