

Asset price dynamics with heterogeneous beliefs [★]

P. M. Beaumont ^{a,*}, A. J. Culham ^b, A. N. Kercheval ^c

^a*Department of Economics, Florida State University*

^b*FPL Energy, Juno Beach, Florida*

^c*Department of Mathematics, Florida State University*

Abstract

We examine market dynamics in a Lucas-style, asset-pricing model with heterogeneous traders who know the distribution of dividends but not the private information of other traders. Agents optimize a CRRA utility function while learning about aggregate states in order to better estimate the equilibrium pricing function. Our goal is to determine whether and how prices evolve toward equilibrium. In the case where all agents have logarithmic utility, but possibly different holdings and discount factors, we completely describe the market dynamics and show that the familiar equilibrium pricing formula also applies in this more general setting.

Keywords: Heterogeneous agents; Asset pricing; Rational Expectations; Learning

JEL classification: C61; D53; D83; G12

[★] This work is partially supported by the FSU Council on Research and Creativity, IDS grant #014895.

* Corresponding Author: Department of Economics, Florida State University, Tallahassee, FL 32306 USA; Tel.: +001 850 644 7085; fax: +001 850 644 4535.

Email addresses: beaumont@coss.fsu.edu (P. M. Beaumont),
aculham@math.fsu.edu (A. J. Culham), kercheval@math.fsu.edu (A. N.

1 Introduction

Price volatility and high trading volume are pervasive features of asset markets that are difficult to capture with standard equilibrium asset pricing models. For well over a decade financial economists have investigated the role that learning plays in explaining these asset market properties. Notable early examples in the literature include Arthur et al. (1997); Brock and Hommes (1998); Marcet and Sargent (1989); Sargent (1993) and Timmermann (1993, 1996). Of the many approaches used, most make either strong assumptions to obtain analytical solutions or employ ad hoc behavioral assumptions about agents.

In this paper we use an equilibrium asset pricing model consistent with the standard asset pricing theory of Lucas (1978), but we allow for heterogeneous agents who must learn essential aspects of aggregate markets while making optimal decisions based upon the best available information. This approach implies a complex interaction between the learning mechanism and the market dynamics of the model.¹

The complexity of this problem suggests that few analytical results will be available and that numerical methods will be necessary. Agent-based models (Arthur et al., 1997; Chiarella et al., 1998; LeBaron et al., 1999) often use myopic agents who use *ad hoc* learning rules and do not attempt to solve an optimal first-order condition to a dynamic optimization problem. Although these models are quite promising for analyzing complex market arrangements,

Kercheval).

¹ In a similar context, Sargent (1993) notes, “Verifying the convergence of the system is technically difficult because the firms are learning about a moving target, a law of motion that is influenced by their own learning behavior (pg. 123). ”

the market dynamics that emerge from such models are difficult to interpret because agent behavior is often not consistent with established asset pricing theory. Evans and Honkapohja (1995, 2001), Adam et al. (2006), Guidolin and Timmermann (2007) and Carceles-Poveda and Giannitsarou (2007b) are able to derive some analytical results in learning models by assuming that agents know their correct demand functions even though they do not know the correct aggregate pricing function for the asset. As a consequence, agents do not trade at market clearing prices during the learning process so that market dynamics are greatly simplified. Scheinkman and Xiong (2003) and Li (2007) present two agent models that do allow for trading during the learning process but, in order to achieve tractable analytical results, the models use specialized assumptions about agents' beliefs of the diffusion process driving dividends.

Our point of view is that agents are forward-looking and good optimizers but find themselves in an environment where they have insufficient information to deduce equilibrium asset pricing functions. In particular, even though agents may know the distribution of dividends and some aggregate information such as total market volume and the supply of assets, they cannot observe the asset holdings or preferences of other agents. It is likely, therefore, that traders will form differing views of the distribution of future asset prices and are likely to trade with one another based upon those differing beliefs.

In our model a single stock paying a stochastic dividend is traded and we assume that agents know only their own parameters and the aggregate dividend distribution. We study each agent's optimization problem in order to discover how they behave within this context. Specifically, we are interested in whether markets will evolve toward equilibria as agents observe and learn

market characteristics.

This contrasts with the usual approach in which the market equilibrium is computed without consideration for how the agents find it—in effect, assuming that agents have full information about the market and all other agents. In our context, since the market clearing price will depend not only upon the observable dividend but also upon unobservable variables such as the holdings of the other agents, it is not clear in advance that agents will be able to solve their individual optimization problems in each time period.

We resolve this difficulty by allowing our agents to make a “pragmatic compromise” in the computation of their asset demand functions. Specifically, our agents approximate the market clearing function with a function of the aggregate dividend alone. Our agents operate under the hypothesis that the distribution of future market clearing prices will depend, at equilibrium, only on the dividend process—just as is the case in the homogeneous agent Lucas model. This hypothesis will be eventually satisfied if the equilibrium pricing function is independent of the wealth distribution or if the market converges over time to a no-trading equilibrium (Judd et al., 2003).

With this approach, agents can in principle solve their Euler equations for the optimal consumption and stock demand in each time period. The trader’s learning problem becomes one of finding the correct equilibrium pricing function with which to compute lifetime expected utility in her dynamic optimization problem. Traders begin with an initial guess of the pricing function and update their guess over time using some learning algorithm based upon observed market prices. In this situation, for a given learning mechanism, we want to study whether agents will eventually learn the correct equilibrium pric-

ing function that forecasts actual market clearing prices in some no-trading equilibrium and, if so, what are the price and holdings dynamics along the way?

Although we assume that our agents are able to solve any well-posed mathematical problem, it appears that with multiple distinct agents an exact analytical solution of the behavior of this market is out of reach. Indeed, even numerical solutions are difficult because the individual optimal demand functions must be computed in each time step to find the market clearing prices. However, as described in this paper, it turns out that the case of log utility can be solved completely. These results both stand on their own and are useful as a point of comparison to help validate numerical simulation results for non-log utility cases.

What makes the log-utility case analytically tractable is that the pricing function is not needed in the agent's optimization problem because it drops out of the agent's Euler equation. Therefore the learning process is *a posteriori* irrelevant to the market evolution. Indeed, our log-utility results apply independently of how the agents think they should forecast prices, so our results apply to any single asset, log-utility, heterogeneous agent Lucas economy. Since the resulting pricing function is the familiar one coming from the homogeneous agent model, this paper also provides an extension of the applicability of that function to a more general heterogeneous, limited-information setting.

2 An equilibrium model with heterogeneous agents and learning

Consider the standard Lucas asset pricing model (Lucas, 1978; Lyungqvist and Sargent, 2004) 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 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.

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

$$c_{i,t} + P_t s_{i,t+1} \leq w_{i,t} = (P_t + D_t) s_{i,t} + e_i, \quad \forall t, \quad (1)$$

where $w_{i,t}$ is the agent's period t wealth and e_i is a constant endowment received by the agent each period.

Agent i begins period t knowing her asset holdings $s_{i,t}$ and endowment e_i . Next, today's stock dividend D_t is announced to all agents. At this point the agent does not know her 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_t . These functions then determine the unique price P that clears the market according to

$$\sum_{i=1}^N s_i(P_t) = N. \quad (2)$$

The actual mechanism of market clearing is not important. We imagine that there is some market maker who receives the demand functions from each agent and publicly declares the market price satisfying (2).

Each agent is assumed to be an expected utility maximizer with constant relative risk aversion preferences and risk aversion $\gamma_i > 0$. The demand function must be solved to optimize expected utility,

$$\max_{\{c_{i,\tau}, s_{i,\tau+1}\}_{\tau=t}^{\infty}} E_t \sum_{\tau=t}^{\infty} \beta_i^{\tau-t} \frac{c_{i,\tau}^{1-\gamma_i} - 1}{\gamma_i} \quad (3)$$

subject to the budget constraint (1) and initial conditions for wealth, asset holdings and dividends. The agents' discount factors, $\beta_i \in (0, 1)$, may differ and the expectation in (3) is over the distribution of dividends and is conditional upon the information available to the agent at the beginning of period t .

For such markets with heterogeneous agents, the existence and uniqueness of equilibria, or even how to compute them, may not be clear (Frydman, 1983). Furthermore, under the dynamics of trading by utility-maximizing agents who need to learn aggregate properties of the market, it is not clear whether the market will converge to an equilibrium. Our primary interest lies in understanding this market dynamics, which includes describing the equilibria, as well as how the market moves when away from equilibrium. First, we need to be precise about what we mean by an equilibrium.

Denote by $\mathcal{S}_t \in \mathfrak{R}^N$ the vector of all agents' time- t asset holdings, which we call the "distribution of holdings". We have explicitly included the time subscript to emphasize that the distribution of holdings may change over time.

Let μ_i denote the triple of each agent's parameters β_i , γ_i , and e_i , all assumed constant in time, and let \mathcal{M} denote the vector of parameters (μ_1, \dots, μ_N) .

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

The REE pricing and demand functions are useful to the omniscient economist, but are not available to the agents themselves who do not observe the quantities \mathcal{S}_t and \mathcal{M} . Therefore, in general, agents will not be able to compute the REE functions but must instead arrive at equilibrium, if at all, by other means.

We can view the individual agent's optimization problem in the usual recursive formulation which leads to the standard Euler equation

$$P_t = E_t \left[\beta_i \left(\frac{c_{i,t+1}}{c_{i,t}} \right)^{-\gamma_i} (P_{t+1} + D_{t+1}) \right]. \quad (4)$$

Using the budget constraint (1) to eliminate c , dropping the subscript i , and using the notation $s = s_t$, $s' = s_{t+1}$, and similarly for the other variables, we may rewrite (4) as

$$P = \beta E \left[\left(\frac{s(P + D) + e - s'P}{s'(P' + D') + e - s''P'} \right)^\gamma (P' + D') \right]. \quad (5)$$

Each agent must solve for consumption and asset demand functions that satisfy this optimality condition.

In general, to evaluate the integral represented by the conditional expectation operator E_t in the Euler equation and to compute their demand functions, agents must know the distribution of future spot prices P_{t+1} of the asset. This, in turn, generally requires that the agents have a complete knowledge of the other agents' asset holdings, \mathcal{S}_t , and preferences, \mathcal{M} .²

Before continuing with the general model, it is useful to discuss an important special case for which the REE is easy to compute. Suppose all agents are identical with log utility and no endowments ($\beta_i = \beta$, $\gamma_i = 1$, $s_{i,t} = 1$, and $e_i = 0$, for all i and t) and all agents are aware of this. Knowing they are identical, agents can deduce that there will be no asset trading and that 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 (4) simplifies to

$$P_t = E_t \left[\beta_i \left(\frac{D_t}{D_{t+1}} \right) (P_{t+1} + D_{t+1}) \right] \quad (6)$$

and it is easy to check that the REE aggregate pricing function satisfying this equation is $P^*(D_t) = \frac{\beta}{1-\beta} D_t$.

Note that the aggregate distribution \mathcal{S}_t is degenerate in the homogeneous agent case, so does not appear in the pricing function $P^*(D_t)$, which now depends on D_t alone.³ Also, the demand function s' has vanished from the

² Nontrivially, each agent also would have to know that all other agents know that all agents know \mathcal{S}_t and \mathcal{M} and behave optimally (Frydman and Phelps (1983); Townsend (1983)). The Lucas (1975) argument that this degree of knowledge is applicable in cases where the aggregate distributions had already settled down to their stationary values and were thus observable by all agents seems to have been embraced by most researchers. However, Frydman (1983) argues that this reasoning is circular because it implies that “the markets are in the rational expectations equilibrium if and only if every agent forms its expectations according to its rational expectations equilibrium forecast function (pg. 111).”

³ When the parameters \mathcal{M} are fixed, we may always omit them as arguments in

Euler equation and is no longer needed to determine the market clearing price.

Since the REE pricing function is known *a priori*, 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 and this elegant solution is often employed in the financial economics literature. But its operational validity is questionable in the sense that 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 (Frydman, 1983).

When agents lack the certainty that they are identical or, as we assume, agents do not know the aggregate distribution of holdings \mathcal{S}_t , they cannot compute the equilibrium pricing function because it depends upon unobservable variables. The situation is even worse away from equilibrium because agents' pricing functions must then also depend on the pricing functions of all the other agents, since those functions help determine market demand and therefore the market clearing price. In other words, as ? notes, "A right forecast must take into account the possibly wrong forecasts of others (pg. 1254)." These are fundamental stumbling blocks for defining the behavior of heterogeneous agents in limited-information Lucas-style asset models.

Our solution to this problem is to allow agents to use private estimates $\tilde{p}_{i,t+1}(D_{t+1})$ of the aggregate pricing function in their Euler equations (4). We assume that these estimates depend upon D_t alone. Although agents are aware that there are unobserved variables \mathcal{S}_t and functions $\{\tilde{p}_{j,t}\}_{j \neq i}$ influencing

the pricing and demand functions. The dependence on fixed parameters \mathcal{M} is of interest to the economist, not the agents.

market prices away from equilibrium, they operate under the hypothesis that the market will converge to an equilibrium so that the dependence on these variables vanishes with time. In other words, our agents hypothesize that all agents' personal pricing functions will converge to the REE pricing function and that the distribution of holdings \mathcal{S}_t will converge to a constant distribution at equilibrium. Thus, it is pragmatic for the agents to use the approximation $\tilde{p}_{i,t+1}(D_t)$ for the aggregate pricing function.

This provides agents with enough information, in principle, to solve for their optimal consumption and asset demand functions $c_{i,t} = \tilde{c}_i(s_{i,t}, D_t, P | \tilde{p}_i)$ and $s_{i,t+1} = \tilde{s}_i(s_{i,t}, D_t, P | \tilde{p}_i)$, where we have adopted the notation \tilde{c}_i and \tilde{s}_i to denote the i^{th} agent's demand functions. Note that these demand functions are conditional on the agent's estimated pricing function $\tilde{p}_{i,t+1}$.

We assume that our agents are aware that their aggregate pricing function estimates $\tilde{p}_{i,t+1}$ may not be accurate so they will attempt to improve upon them by observing actual market prices over time. The *actual* market price P_t will in general depend not only on D_t but also on the private quantities \mathcal{S}_t , \mathcal{M} , and all the agents' personal pricing functions $\tilde{p}_{i,t+1}$. Therefore, market dynamics will depend upon the chosen learning mechanism that drives the time evolution of $\tilde{p}_{i,t+1}$, as well as the other parameters of the agents. In general, it is not clear that such a market will reach an equilibrium.⁴

Our model as described above, with the use of recursive least squares learning, can be solved for multiple agents using a discrete numerical solution of each agent's optimization problem in each time step. Preliminary work (Cul-

⁴ ? argue that learning is an unsatisfactory foundation for rational expectations but ? provide an aggregation procedure to construct a consensus agent that allows them to analyze equilibrium characteristics in models with heterogeneous beliefs.

ham, 2007) suggests that convergence to the REE is quite robust. In the next section we analyze the log-utility case which is independent of the learning mechanism of the agent and thus admits an analytical solution. Beyond its own interest, this will be useful as a base case for comparison with the purely numerical simulation results when agents are allowed to possess more general utility functions.

3 Analytical results for the log-utility, zero endowment case

Suppose all $N > 1$ agents have CRRA utility with a common risk aversion and zero endowment ($\gamma_i = \gamma > 0$ and $e_{i,t} = 0$ for all i and t) 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 $\tilde{p}_i(D)$ 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 (4) is

$$\frac{P}{(s_i(P+D) - P\tilde{s}_i(s_i, D, P))^\gamma} = \beta_i E_t \left[\frac{\tilde{p}_i + D'}{(\tilde{s}_i(s_i, D, P)(\tilde{p}_i + D') - \tilde{p}_i\tilde{s}_i(\tilde{s}_i(s_i, D, P), D', \tilde{p}_i))^\gamma} \right]. \quad (7)$$

Notice that the unknown demand function \tilde{s}_i appears in (7) in a highly non-

linear way. Nonetheless, agents must solve this equation for \tilde{s}_i in order to be able to participate in the implicit price-calling auction used to arrive at the market clearing price.

Remarkably, in the log-utility case, $\gamma = 1$, (7) turns out to have the simple explicit solution

$$\tilde{s}_i(s_i, D, P) = \beta_i(1 + (D/P)) s_i. \quad (8)$$

Critically, and unlike the general case for $\gamma \neq 1$, this solution has the special property that it does not depend upon the agent's estimated pricing function \tilde{p}_i nor directly upon the distribution of holdings \mathcal{S} . This is a very important feature of log utility and is what makes this special case analytically tractable.

Consequently, given an announced asset price, agents can act optimally without knowing the distribution of the aggregate state variables because their optimal asset demand functions are independent of the mechanism that they have chosen for learning the equilibrium pricing function. The agents themselves do not know this and so believe that the distribution \mathcal{S} affects the market clearing prices and that their own demand functions are only approximations based upon the simplifying assumption that aggregate states may be ignored during the learning process. We, the omniscient modelers, however, know that every agent has log utility, and so we can analyze how prices and holdings will evolve as these agents behave optimally in each time step.

The market clearing price, which we now denote P_m , is determined from the market clearing condition

$$\sum_{i=1}^N \tilde{s}_i(s_i, D, P_m) = N. \quad (9)$$

Substituting the demand function (8) for \tilde{s}_i and solving for the market clearing price, gives

$$P_m = \frac{\sum_1^N \beta_j s_j}{N - \sum_1^N \beta_j s_j} D. \quad (10)$$

Substituting (10) into the demand function (8) gives the agent's next period holdings at market clearing prices as

$$s'_i = \tilde{s}_i(s_i, D, P_m) = \beta_i \left(\frac{N}{\sum_1^N \beta_j s_j} \right) s_i. \quad (11)$$

Although the agent's demand function (8) in the log utility case does not depend directly upon the distribution of assets \mathcal{S} , it does depend upon this distribution indirectly through the observed market clearing price (10). Also, s'_i does not depend upon the dividend D except indirectly through the market price. The market clearing price and stock holdings will evolve according to the dynamical system described by (10) and (11).

Conditional on the number of agents N and the discount factors β_i , the demand function (11) may be written as a function of the asset distribution \mathcal{S} alone. Thus, for all N agents we may write an N -dimensional dynamical system $\mathcal{S}' = \tilde{s}(\mathcal{S})$ describing how the evolution of holdings is determined by iteration of the N -dimensional function \tilde{s} .

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 equations (10) and (11). This system converges exponentially fast to

$$P^*(D) = \frac{\beta}{1 - \beta} D \tag{12}$$

and

$$s_i^* = \frac{N s_i^\circ}{s_1^\circ + \dots + s_k^\circ}, \quad i \leq k, \tag{13a}$$

$$= 0, \quad i > k. \tag{13b}$$

Proof: See Appendix.

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 same rational expectations equilibrium pricing function obtained under the classical and more restricted assumption that agents are identical and have perfect information.

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

$$P_m = \frac{\beta}{1 - \beta} D. \quad (14)$$

3.1 *An illustration with two agents*

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

$$\tilde{s}_i(s) = \beta_i \frac{2s_i}{\beta_1 s_1 + \beta_2 s_2}, \quad i = 1, 2. \quad (15)$$

Using the market clearing constraint $s_1 + s_2 = 2$ gives

$$\tilde{s}_1(s_1) = \frac{\beta_1 s_1}{\beta_2 + (\beta_1 - \beta_2)(s_1/2)} \quad (16a)$$

and

$$\tilde{s}_2(s_2) = \frac{\beta_2 s_2}{\beta_1 + (\beta_2 - \beta_1)(s_2/2)}. \quad (16b)$$

These two functions are plotted in Figure 1 for the discount factors $\beta_1 = 0.95$ and $\beta_2 = 0.7$. 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.

Since consumption is a fixed share of wealth in the log utility case, the most patient agent's consumption increases with investment holdings. Specifically, for the two agent case, consumption of the more patient agent is given by

$$c_1 = \frac{(1 - \beta_1)s_1 D}{1 - \beta_2 - (\beta_1 - \beta_2)(s_1/2)} \quad (17)$$

which is increasing and convex in s_1 . As $s_1 \rightarrow 2$, $c_1 \rightarrow 2D$ so that the patient agent eventually consumes all of the dividend payments. Consumption of the less patient agent converges to zero. If $\beta_1 = \beta_2$ then $c_i = s_i D$ so the agents do not trade and simply consume their share of the dividend payments.

Substituting the asset demand functions (16) into the market clearing price (10) gives

$$P_m = \left(\frac{2\beta_2 + (\beta_1 - \beta_2)s_1}{2(1 - \beta_2) - (\beta_1 - \beta_2)s_1} \right) D \quad (18)$$

which reduces to $P_m = \beta/(1 - \beta)D$ when $\beta_1 = \beta_2 = \beta$. The price/dividend ratio P_m/D is increasing and convex in s_1 . This makes sense since, from (10), we see that the market price is an increasing function of the weighted average of the β_i 's with the investment shares s_i/N as weights. As the investment share of the most patient agent increases, the weight on the highest β increases making the market price higher. More intuitively, as the most patient agent becomes wealthier his consumption and investment demands increase driving up the price of the asset.

4 Conclusion

In this paper we have described a heterogeneous agent, equilibrium asset pricing model in which agents' information about the aggregate economy is realistically limited. In a dynamic equilibrium model agents need aggregate asset pricing functions in order to solve their Euler equations for their goods and asset demand functions. To compute the rational expectations equilibrium pricing functions, agents need to know the utility functions and asset holdings of all other agents in the economy. Additionally, agents need to know that all other agents know this information and use it optimally. This informational requirement is trivial in a homogeneous agent model, since the aggregate distributions are degenerate, but becomes overwhelming in the heterogeneous agent case with non-degenerate aggregate distributions. This leads to the question of how heterogeneous agents can come to know the equilibrium pricing function without having access to an unrealistic amount of information about other agents.

In our model we assume that agents know their own preferences and in-

vestments as well as the stochastic process driving asset dividends but that they do not know the preferences and investments of other agents. However, our agents operate under the hypothesis that, at equilibrium, aggregate distributions will become stationary so that asset prices will eventually become independent of these distributions and thus depend only upon the observed dividend. Consequently, our agents make personal approximations of the equilibrium aggregate pricing function and base their estimated demand functions upon those approximations. This is a critical feature that differentiates our model from many other learning models.

As markets evolve, the consumption and investment decisions of our agents, which are based upon observed market clearing prices, will be revealed to be suboptimal as the agents learn that the market clearing prices have deviated from their expectations. This prompts our agents to update both their approximations to the aggregate pricing functions as well as their demand functions. During the learning process the market will exhibit complex dynamics that will depend upon the specific type of learning used by the agents (Carceles-Poveda and Giannitsarou, 2007a,b).

In general, numerical methods will be required to analyze the dynamics and convergence properties of such an economy. However, we have derived analytical results for the special case where all agents have log utility. The critical feature is that the agents' pricing function approximations cancel out of their first-order conditions for the log utility case. This means that the agents' consumption and asset demands are independent of the learning rule that they choose and we are able to completely characterize the resulting market dynamics and convergence properties for this case.

The characteristics of the market dynamics and the convergence properties of models with differing degrees of risk aversion remains an open question. Preliminary numerical results (Culham, 2007) suggest that convergence to the rational expectations equilibrium is quite robust.

Appendix. Proof of Theorem 1

It is easy to verify algebraically that the demand function (8) solves the Euler equation (7), and therefore (10) and (11) 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 (13).

It is convenient to rewrite the dynamical system (11) in terms of the relative holdings $x_i = s_i/N$:

$$x'_i = \frac{\beta_i x_i}{\sum_1^N \beta_j x_j}. \quad (\text{A.1})$$

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

$$\Delta^N = \{(x_1, \dots, x_N) \geq 0 : \sum_i x_i = 1\} \quad (\text{A.2})$$

in the positive orthant of \mathfrak{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 (A.1).

If $U \subset \Delta^N$ denotes the k -dimensional sub-simplex

$$U = \{(x_1, \dots, x_k, 0, \dots, 0) : \sum_{j=1}^k x_j = 1\}, \quad (\text{A.3})$$

then it is easy to see that every point of U is fixed by T . Likewise, let V denote the $(N - k)$ -dimensional sub-simplex

$$V = \{(0, \dots, 0, x_{k+1}, \dots, x_N) : \sum_{j=k+1}^N x_j = 1\}. \quad (\text{A.4})$$

From (A.1), 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 (13) if we can show that every forward T -orbit $\{T^n(x)\}$ converges to U .

Define $\pi_U : \Delta^N \rightarrow U$ to be the projection fixing the first k coordinates and setting the remaining $N - k$ coordinates to zero, and similarly $\pi_V : \Delta^N \rightarrow V$ the projection fixing the last $N - k$ coordinates and setting the first k to zero. Let $\Delta^{N+} = \{x \in \Delta^N : \pi_U(x) \neq 0\}$.

Lemma 1 *Define $F : \Delta^N \rightarrow \Re$ 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: $F(x)$ is simply a weighted average of the β 's, weighted by the x 's. Using the definition of T , we have, for any x ,

$$F(T(x)) = \frac{\sum_i \beta_i^2 x_i}{F(x)} \quad (\text{A.5})$$

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

$$\phi\left(\sum_i \beta_i x_i\right) \leq \sum_i \phi(\beta_i) x_i \quad (\text{A.6})$$

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

Fix $x \in \Delta^{N+}$. If $\pi_V(x) = 0$ then $x \in U$, $T(x) = x$, and $F(x) = \beta$, so there is nothing further to prove. Hence suppose $\pi_V(x)$ is nonzero. This means $\pi_U(T(x))$ and $\pi_V(T(x))$ 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 $F(x^*) < \beta$, $\pi_V(x^*) \neq 0$, and so

$$F(T(x^*)) > F(x^*) = \beta^*. \quad (\text{A.7})$$

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 $\beta^* = \beta$. \square

From the lemma above, the limit of $F(T^n(x))$ is β for all $x \in \Delta^{N+}$. Since F is continuous and $F^{-1}(\beta) = U$, every forward T -orbit starting in Δ^{N+} must converge to U . From equation (11), we see that if holdings are close to zero for agents $j = k + 1, \dots, N$, then we have, approximately,

$$s'_j = \frac{\beta_j}{\beta_1} s_j, \quad (\text{A.8})$$

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

References

- Adam, K., Marcet, A., Nicolini, J. P., April 2006. Stock market volatility and learning.
- Arthur, W. B., Holland, J. H., LeBaron, B., Palmer, R., Tyler, P., 1997. 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.
- Brock, W. A., Hommes, C. H., 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control* 22, 1235–1274.
- Carceles-Poveda, E., Giannitsarou, C., 2007a. Adaptive learning in practice. *Journal of Economic Dynamics and Control* 31, 2659–2697.
- Carceles-Poveda, E., Giannitsarou, C., February 2007b. Asset pricing with adaptive learning.
- Chiarella, C., Dieci, R., Gardini, L., 1998. Asset price and wealth dynamics in a financial market with heterogeneous agents. *Journal of Economic Dynamics and Control* 30 (1755–1786).
- Culham, A. J., 2007. Asset pricing in a lucas framework with boundedly rational, heterogeneous agents. Ph.D. thesis, Florida State University, Department of Mathematics, Tallahassee, Florida, USA.
- Evans, G. W., Honkapohja, S., January 1995. Local convergence of recursive learning to steady states and cycles in stochastic nonlinear models. *Econometrica* 63 (1), 195–206.

- Evans, G. W., Honkapohja, S., 2001. *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, NJ.
- Frydman, R., 1983. Individual rationality and decentralization. In: Frydman and Phelps (1983), Ch. 5, pp. 97–122.
- Frydman, R., Phelps, E. S., 1983. *Individual forecasting and aggregate outcomes: Rational expectations examined*. Cambridge University Press, Cambridge, UK.
- Guidolin, M., Timmermann, A., 2007. Properties of equilibrium asset prices under alternative learning schemes. *Journal of Economic Dynamics and Control* 31, 161–217.
- Judd, K. L., Kubler, F., Schmedders, K., October 2003. Asset trading volume with dynamically complete markets and heterogeneous agents. *The Journal of Finance* LVIII (5), 2203–2217.
- LeBaron, B., Arthur, W. B., Palmer, R., 1999. Time series properties of an artificial stock market. *Journal of Economic Dynamics and Control* 23, 1487–1516.
- Li, T., 2007. Heterogeneous beliefs, asset prices, and volatility in a pure exchange economy. *Journal of Economic Dynamics and Control* 31, 1697–1727.
- Lucas, R. E., 1975. An equilibrium model of the business cycle. *Journal of Political Economy* 83, 1113–1144.
- Lucas, Jr, R. E., 1978. Asset prices in an exchange economy. *Econometrica* 46 (6), 1429–1445.
- Lyungqvist, L., Sargent, T. J., 2004. *Recursive Macroeconomic Theory*, 2nd Edition. The MIT Press, Cambridge, MA.
- Marcet, A., Sargent, T. J., 1989. Convergent of least squares learning mechanisms in self-referential linear stochastic models. *Journal of Economic Theory* 48 (2), 337–368.

- Sargent, T. J., 1993. *Bounded Rationality in Macroeconomics*. Oxford University Press.
- Scheinkman, J. A., Xiong, W., December 2003. Overconfidence and speculative bubbles. *Journal of Political Economy* 111 (6), 1183–1219.
- Timmermann, A., 1993. How learning in financial markets generates excess volatility and predictability in stock prices. *Quarterly Journal of Economics* 108, 1135–1145.
- Timmermann, A., 1996. Excess volatility and predictability of stock prices in autoregressive dividend models with learning. *Review of Economic Studies* 63 (523–557).
- Townsend, R., 1983. Forecasting the forecast of others. *Journal of Political Economy* 91, 546–588.




Fig. 1. The demand functions of two agents, $\beta_1 = 0.95$, $\beta_2 = 0.7$. Iteration of the upper function is shown with the arrows; holdings converge to $s_1 = 2$. Likewise, holdings for the other agent converge to zero.