Asset pricing with a continuum of belief types*

Cees Diks & Roy van der Weide

CeNDEF, Faculty of Economics and Econometrics, University of Amsterdam, and Tinbergen Institute, Roetersstraat 11, 1018 WB, Amsterdam, The Netherlands

Abstract

We propose a general framework for describing the evolution of heterogeneous beliefs in dynamic settings. Taking a continuous beliefs space as a starting point, beliefs distributions are introduced, as well as a dynamic updating rule based on the continuous choice model. The methodology is illustrated using dynamical asset pricing models, in which agents optimize their strategies based on past performance measures. This approach gives rise to price dynamics in which the beliefs distribution evolves together with realized prices. For several models, the role of the performance measure, the class of predictor functions, and the memory parameter is examined both analytically and numerically. By considering aggregate beliefs, conditions can be derived under which the dynamics tends to a deterministic or a stochastic law when the number of agents tends to infinity. In the case of stochastic dynamics, the distributional properties of the endogenous random price fluctuations can be expressed in terms of the beliefs distribution. Finally, conditions are given under which the dynamics are finite dimensional, and the beliefs distributions are observable from realized prices time series.

Keywords: Heterogeneous agents; Expectation formation; Continuous beliefs; Continuous choice; Endogenous noise;

JEL classification: C50, G12, D84

*Research supported by the Netherlands Organization for Scientific Research (NWO) under a MaG-Pionier grant.

1 Introduction

As little as it takes one to realize that expectations about future prices are essential for investment decisions made by market participants, as difficult it is to study the interaction between prices and expectations empirically. Expectations, let alone the beliefs on which they are based, are hardly ever directly observable in practice. For stock markets, for example, one typically observes realized prices, reflecting market aggregated expectations rather than individual expectations. We are interested in issues such as the interaction between information available to market participants and the distribution of their beliefs, the time evolution of the heterogeneity among agents, and the question to which extent it is possible to extract information about the distribution of beliefs from empirical data. The aim of this paper is to provide a framework for studying these types of questions. We propose a methodology and illustrate it with agent-based asset pricing models, since they are typical examples of the dynamic feedback systems we have in mind.

Until recently, the literature on agent-based asset pricing models has been mainly concerned with price dynamics in settings where agents are allowed to choose among a finite number of strategies. Brock and Hommes (1997) showed that, with a discrete choice model for the distribution of strategies, these models give rise to deterministic price dynamics for a continuum of agents. Brock *et al.* (2001) examined conditions under which similar results are obtained for models with agents that are allowed to employ strategies from a continuum of strategies. This was done by considering a limit, in which the number of strategies which the agents can employ tends to infinity. This approach was motivated by the question whether an extension of the diversity of belief types leads to more complex dynamics and possibly more realistic price dynamics than that observed in few type systems.

As noted above, the motivation behind the methodology we develop here is different. We have in mind a framework which provides insights into the way agents update their beliefs, which can be used to study the interaction between realized prices and beliefs, and might provide information about the beliefs distribution from observed realized prices. Therefore, in our approach, the evolution of the probability distribution over the belief parameter space is accounted for explicitly by a continuous choice model. The beliefs distribution, which is updated using past realized prices, plays the role of a population distribution, rather than the actual distribution of strategies among the market participants. It represents the probability distribution from which the strategies eventually employed by the individual agents, are drawn.

Since the continuous choice model utilizes random components on individual utility functions the price dynamics is stochastic if the number of agents is finite. Depending on the details of the model, such as the performance measure used to evaluate strategies, the price dynamics either remain stochastic or tend to a deterministic law when the number of agents tends to infinity. This endogenous randomness is a direct consequence of the use of a continuum of strategies. The type of randomness arising in these models is extremely relevant for the study of deterministic price dynamics with endogenous dynamic noise, since it allows the explicit derivation of distributional properties of the endogenous randomness and the exact way where the noise enters the price dynamics. To the best of our knowledge, in the literature, these type of dynamics have only been studied with random shocks with ad-hoc distributional properties, and which were implemented as additive noise on the price equation without further theoretical justification.

This paper is organized as follows. In section 2 the concept of a beliefs distribution in a continuous beliefs space is introduced, as well as its form given by the continuous choice model. Section 3 describes the the evolution of the beliefs distribution in a dynamic setting where beliefs are updated when new information becomes publically available. This section also describes the dynamics this implies in a standard asset pricing context. In section 5 the mechanism by which endogenous noise arises from the dynamic continuous choice dynamics is described. Section 6 summarizes and discusses the results.

2 Continuous choice

Agent based models represent market participants as (typically a large number of) agents, who can select among a number of alternative strategies. If the strategies among which the agents can choose consists of a finite set strategies, $s_1, \ldots, s_k$ say, then agents employing strategy s_i, $i = 1, \ldots, k$ are said to be of type i. McFadden (1973) derived an expression for the fractions n_i of agents employing strategy i starting from the concept of random utility functions. It is assumed that the utility function of agent j can be written as

$$V_j(s) = U(s) + \epsilon_j(s)$$

where $U(s)$ is a non-stochastic "common" utility function representing the tastes of the population, and $\epsilon_j(s)$ is stochastic and reflects the idiosyncrasies of individuals in tastes. The individuals choose the alternative which optimizes their utility. Under the assumption that the disturbances of the utility function follow a Gnedenko extreme value distribution, it can be shown that this leads to the *multinomial logit distribution* of choices:

$$n_i = \frac{\mathrm{e}^{\beta U(s_i)}}{\sum_{l=1}^{k} \mathrm{e}^{\beta U(s_l)}},$$

where $U(s_j)$ is the utility associated with alternative j. The parameter β is referred to as the *intensity of choice*, and is related to the scale of the noise term $\epsilon_j(s)$. The larger the value of β, the smaller the noise, and the larger the probability that an agent chooses the option which actually optimizes $U(s)$. This is why $1/\beta$ is sometimes interpreted as the propensity of agents to err, assuming they actually all wish to optimize $U(s)$.

In the presence of a continuum of belief types it is convenient to introduce a finite dimensional parameter space Θ, containing all possible strategies that can be employed by the agents. Each element θ in Θ uniquely determines a possible strategy. We will refer to Θ as the beliefs space. Note that the choice of the beliefs space is not unique, since any one-to-one transformation of the beliefs space Θ into another space, Θ', say, will again yield a suitable representation of the beliefs space. In analogy with the discrete choice model, we wish to represent the diversity of belief types by a probability distribution function over the beliefs space. The distribution of strategies can be obtained from the generalization of the discrete choice model referred to as mixed discrete/continuous choice models. Let us denote the pdf associated with the beliefs distribution by $\phi(\theta)$. As in the discrete choice setting, it is convenient to adopt a random utility approach (Hanemann, 1984; Dagsvik, 1994). The random part of the utility function of an agent will eventually determine which strategy that particular agent considers optimal. Therefore, the strategies employed by individual agents in a random utility framework are random variables (note that this does not imply that individual agents perceive their own utility functions to be random, only that they are random to the econometrician). If we let θ_i denote the strategy adopted by agent i, then the pdf $\phi(\theta)$ is the probability density function associated with the random variables θ_i. Dagsvik (1994) derived the following pdf for the θ_i:

$$\phi(\theta) = \frac{\mathrm{e}^{\beta U(\theta)}\varphi(\theta)}{Z}, \tag{1}$$

where Z is a normalization constant, given by

$$Z = \int_{\Theta} \mathrm{e}^{\beta U(\vartheta)}\varphi(\vartheta)\mathrm{d}\vartheta.$$

As in the discrete choice setting, β represents the intensity of choice. The function $\varphi(\theta)$ is nonnegative, and can be used to put different weights on different parts of the beliefs space. We refer to $\varphi(\theta)$ as the *opportunity function*. The introduction of $\varphi(\theta)$ is not an additional expansion of the model, but necessary for obtaining a representation that is consistent in that it is independent of the parameterization of the beliefs space Θ. Mathematically, $\varphi(\theta)$ fixes the integration measure in the beliefs space Θ. In economic applications the opportunity function can be thought of as reflecting the *a priori* faith of individuals in parameters within certain regions of the parameter space. In an asset pricing framework, if agents have a tendency to avoid strategies in certain parts of the parameter space, for example because extremely large or small parameter values are implausible to agents, this will be reflected by small values of $\varphi(\theta)$ in those regions of the parameter space.

The beliefs distribution is independent of the parameterization of Θ, provided that the opportunity function is transformed properly when moving from one parameterization to another, as stated in the following proposition.

Proposition 1 *Under the convention that $\varphi(\theta)$ transforms as a density function under one-to-one measurable coordinate transformations, $g : \Theta \mapsto \Theta'$, that is, $\varphi'(\theta') = \mid J_{g^{-1}}(\theta') \mid \varphi(g^{-1}(\theta'))$, where $J_{g^{-1}}(\theta')$ is the Jacobian of $g^{-1}(\cdot)$ evaluated at θ', the beliefs representation is independent of the parameterization of Θ.*

Proof: Requiring that the pdf $\phi(\theta)$ transforms as a density leads to

$$\phi'(\theta') = \mid J_{g^{-1}}(\theta') \mid \phi(g^{-1}(\theta')).$$

From the definition of $\phi(\theta)$, after substituting $U(g^{-1}(\theta')) = U'(\theta')$, one obtains

$$\mathrm{e}^{\beta U(g^{-1}(\theta'))}\varphi'(\theta') = \mid J_{g^{-1}}(\theta') \mid \mathrm{e}^{\beta U(g^{-1}(\theta'))}\varphi(g^{-1}(\theta')).$$

By defining $\varphi'(\theta') = \mid J_{g^{-1}}(\theta') \mid \varphi(g^{-1}(\theta'))$ it can be easily seen that this holds for all possible utility functions $U(\theta)$. □

Generally speaking, we neither need to require that $\varphi(\theta)$ be a pdf, nor that it is integrable in order for the distribution of beliefs $\phi(\theta)$ to be well-defined. For example, if $\Theta = \mathbf{R}^m$, $\phi(\theta) = 1$, and $U(\theta)$ is a quadratic form in θ with a single maximum, then $\phi(\theta)$ is a multivariate normal probability density function.

As stated above the distribution of beliefs does not always exist. The following proposition describes the necessary and sufficient conditions for the existence of the probability density of beliefs over Θ.

Proposition 2 *The beliefs distribution $\phi(\theta)$ given in Eq. (1) is well-defined as a pdf if and only if $Z = \int_\Theta \mathrm{e}^{\beta U(\vartheta)}\varphi(\vartheta)\,\mathrm{d}\vartheta$ is positive and finite.*

Proof: Since $\mathrm{e}^{\beta U(\theta)}$ and $\varphi(\theta)$ are non-negative, the function $\phi(\theta)$ is nonnegative if and only if $\int_\Theta \mathrm{e}^{\beta U(\vartheta)}\varphi(\vartheta)\,\mathrm{d}\vartheta$ is positive and finite, in which case $\int_\Theta \phi(\vartheta)\,\mathrm{d}\vartheta$ equals one. □

An equivalent way of stating this proposition is that $\phi(\theta)$ is well-defined if and only if $\mathrm{e}^{\beta U(\theta)}$ is φ-integrable. It is possible to give some sufficient conditions for $\mathrm{e}^{\beta U(\theta)}\varphi(\theta)$ to be integrable. For example, if either $\varphi(\theta)$ or $\mathrm{e}^{\beta U(\theta)}$ is bounded, and the other integrable, $\mathrm{e}^{\beta U(\theta)}\varphi(\theta)$ is integrable, so that Z is finite and $\phi(\theta)$ is well-defined. Another sufficient condition is $U(\theta)$ to be bounded from above in Θ and $\varphi(\theta)$ to be integrable.

3 Continuous choice dynamics

In this section we introduce updating of the beliefs distribution in a continuous choice framework. The methodology is illustrated within a standard asset pricing context, but applicable to any dynamic setting in which beliefs distributions are dynamically updated. As noted in the introduction, few belief types systems were among the first dynamic asset pricing models. In reality we would expect a high degree of heterogeneity to be present in the population of market participants, and a small number of belief types will probably be insufficient to model the rich behavior of price fluctuations observed in real life. Typical questions we have in mind are the following. What is the role of the performance measure used to evaluate strategies? How does the heterogeneity among agents' beliefs evolve with realized prices? Under which conditions will a small subset of the initial set of belief types eventually attract most of the population of traders? The latter questions are particularly interesting since they provide insight in the behavior of economic agents, in particular their expectation formation.

We denote the information available to agents at time t by $\mathcal{F}_t = \{p_{t-1}, p_{t-2}, \dots\}$. In an asset pricing context, for example, p_s can stand for the price at time t, so that the information set consists of a historic record of prices p_s up to and including the price at time $t-1$. The price p_t at time t will be determined by the market equilibrium equation which depends on the agents' expectations, at time t, about the price the next time they can sell the asset again, $t+1$. The possible strategies from which the agent can choose to predict future prices are represented by a function of past observables, parameterized by θ, which we shall denote by $f_\theta(p_{t-1}, p_{t-2}, \dots)$. For example, the class of d-th order linear predictors consists of all predictors of the form

$$f_\theta(p_{t-1}, ..., p_{t-d}) = \theta_0 + \theta_1 p_{t-1} + ... + \theta_d p_{t-d},$$

in which case the beliefs are represented in $\mathbf{R}^{d+1}$ and the beliefs distribution is a probability distribution over this space. Generally speaking, the function $f_\theta(p_{t-1}, p_{t-2}, \dots)$ is referred to as the *predictor function.* The prediction of agent i using strategy $\theta_{i,t}$ at time t, for a given price history $\mathcal{F}_t$, will be denoted by $p^e_{i,t+1} = f_{\theta_{i,t}}(p_{t-1}, p_{t-2}, \dots)$. That is, $p^e_{i,t+1}$ represents the agent's *prediction*, at time t, of p_{t+1}, conditional on the information set at time t.

Let n denote the number of agents. The aggregated expectations are defined as the mean expectation

$$\bar{p}^e_{n,t+1} = \frac{1}{n}\sum_{i=1}^{n} p^e_{i,t+1} = \frac{1}{n}\sum_{i=1}^{n} f_{\theta_{i,t}}(p_{t-1}, p_{t-2}, \dots).$$

This definition of average predictions requires the predictions to be representable in a space where averages are well defined, such as $\mathbf{R}^d$, or a convex subset thereof. Note that, more generally, one could examine agents predicting some function of the observables instead of the observable p_{t+1} itself, but this is possibility is not examined further here.

We consider a context in which the price realized today is a function of aggregate expectations of agents about future prices, that is, the price at time t is given by

$$p_t = h\left(\bar{p}^e_{n,t+1}\right), \tag{2}$$

where h is a time independent and continuous function.

Once p_t becomes part of the public information set, the beliefs distribution is updated. This is where the continuous choice framework is incorporated in the model. The common (nonrandom) part of the utility function, conditioned upon information available at time t, is denoted by $U_t(\theta) = U(\theta; \mathcal{F}_t)$. Typically, $U(\theta; \mathcal{F}_t)$ will be based on a fitness measure of strategies θ, such as the last net ex post profits or squared prediction errors. More generally, one might introduce dependence on the further past performances by introducing memory in the model. The evolution of the fitness measure for predictor f_θ can for example be modeled as:

$$U_t(\theta) = \alpha U_{t-1}(\theta) + \pi_t(\theta),$$

where $\alpha \in [0, 1)$ is a memory parameter. The fitness measure $U_t(\theta)$ then becomes a geometrically weighted sum of ex post performances of predictor f_θ. Given a new observation p_t, $U_t(\theta)$ can be updated, and the pdf of the belief distribution is determined by the continuous beliefs model:

$$\phi_t(\theta) = \frac{\mathrm{e}^{\beta U_t(\theta)} \varphi(\theta)}{Z_t}.$$

After this, the updating process of prices and the beliefs distribution can be started all over again.

Since we have described how a new price and a new beliefs distribution can be obtained from the previous beliefs distribution and the price history, the feedback loop is closed. At first glance it might seem that given the beliefs space Θ, the utility function $U(\theta|\mathcal{F})$, and the function g, the model is fully specified by the equations for prices in terms of average expectations, and the evolution of the distribution of beliefs (or equivalently, strategies, or predictors). The model is not complete, however, without specifying how, given the beliefs distribution, the strategies actually used by the agents are drawn from this distribution. There are many beliefs selection mechanisms compatible with a given beliefs distribution, differing in dependence among strategies used among agents and over time. Dependence among traders can arise for example by the exchange of private information. It is also possible that although the idiosyncratic noise of each agent changes only slowly over time, giving rise to temporal dependence of the idiosyncrasies. In many cases, however, the following simplifying assumption is not very restrictive:

Assumption 1 *(Independent agents) The strategies $\theta_{i,t}$ employed by agent i at time t, for each fixed time t are independent random variables, distributed according to the population distribution of beliefs at time t, $\phi_t(\theta)$.*

Assuming independence over agents seems reasonable, since it is always possible to consider expectations of groups of correlated agents as expectations of a single agent representative of this group. The effect of dependence then is merely a reduction in the effective number of agents. A stronger assumption is made if one additionally assumes independence of the idiosyncratic noise of each agent over time:

Assumption 2 *(Independent agents with temporal independence of idiosyncrasies) The strategies $\theta_{i,t}$ employed by agent i at time t, conditionally on $\mathcal{F}_t$ are jointly independent random variables, distributed according to the population distribution of beliefs at time t, $\phi_t(\theta)$.*

This assumption, which is not essential for the main results derived in this section, and will be used only in the example models, is reasonable if the time interval corresponding to one time step in the model is large compared to the time scale on which idiosyncratic preferences of single agents change over time. The conditioning on the information available at time t is necessary to allow for temporal dependence in strategies arising through the price dynamics.

The following lemma is is concerned with the almost sure behavior of the model in the limit where the number of agents tends to infinity. Again we use the short-hand notation

$p^e_{t+1} = f_{\theta_{i,t}}(p_{t-1}, p_{t-2}, \dots)$ for the expectation, at time t, of agent i with parameter θ_i about the information x_{t+1} that will become known at time $t+1$. $E_{\phi_t(\theta)}$.

Lemma 1 *Under Assumption 1, and given $\mathcal{F}_t$, if the number of traders, n, tends to infinity, the aggregate expectation $\bar{p}^e_{n,t+1} = \frac{1}{n}\sum_{i=1}^{n} f_{\theta_{i,t}}(p_{t-1}, p_{t-2}, \dots)$ converges almost surely to $\bar{p}^e_{t+1} = E[f_{\theta_{i,t}}(p_{t-1}, p_{t-2}, \dots)|\mathcal{F}_t]$, if and only if $E[f_{\theta_{i,t}}(p_{t-1}, p_{t-2}, \dots)|\mathcal{F}_t] < \infty$.*

Proof: By Assumption 1, the strategies $\theta_{i,t}$ are IID random variables in Θ, which implies that the predictions $E[f_{\theta_{i,t}}(p_{t-1}, p_{t-2}, \dots)|\mathcal{F}_t]$ of agents are IID random variables. The result is immediate from Kolmogorov's strong law of large numbers for the random variables $f_{\theta_{i,t}}(p_{t-1}, p_{t-2}, \dots)$, given $\mathcal{F}_t$ (see e.g. Resnick, 1998, p. 220). □

Note that $E[f_{\theta_{i,t}}(p_{t-1}, p_{t-2}, \dots)|\mathcal{F}_t] = E[p^e_{t+1}|\mathcal{F}_t] = \int_\Theta f_\vartheta(p_{t-1}, p_{t-2}, \dots)\phi_t(\vartheta)\mathrm{d}\vartheta$. Lemma 1 states that, given the information public at time t, a sufficient and necessary condition for the aggregate expectation about x_{t+1} to to converge a.s. to the mean expectation over the beliefs distribution, is the existence of this mean. This result is used to prove the following theorem about the continuous choice dynamics.

Theorem 1 *(Functionally deterministic dynamics) Under Assumption 1 the price p_t is a deterministic functional of the information $\mathcal{F}_t$ available at time t if $E_{\phi_t(\theta)}[p^e_{t+1}|\mathcal{F}_t] < \infty$.*

Proof: According to Lemma 1, with probability one, $p^e_{n,t+1}$ converges in probability to $E_{\phi_t(\theta)}[p^e_{t+1}|\mathcal{F}_t]$, if and only if $E_{\phi_t(\theta)}[p^e_{t+1}|\mathcal{F}_t] < \infty$. Since h was assumed to be continuous, the right hand side of Eq. (2) which gives the dynamic equation for the observable x_t, then converges in probability to $h\,(g_t)$. □

We coin the term functional determinism here to make a distinction between the concept of determinism we describe here and the usual definition of determinism for time series processes. The term determinism is usually reserved for deterministic processes of finite order, the reason being that without this restriction any observed time series would allow a deterministic explanation. That is, functionally deterministic processes can still be non-deterministic from an empirical point of view, if there is dependence on the infinite past. We will study conditions for the price dynamics to be of finite order later. For the moment, we note only that functional determinism is a necessary, but insufficient condition for the price dynamics to be deterministic.

The discrete choice model can be seen as a special case in which agents can only choose among a finite number of alternative strategies. Provided that the expected future prices of these strategies are all finite, the average expected price is well-defined, so that the strong law of large numbers applies, and the dynamics converges to a deterministic dynamical system with probability one as the number of traders tends to infinity.

Now let us consider the case where Lemma 1 does not apply. If $E_{\phi_t(\theta)}[p^e_{t+1}|\mathcal{F}_t]$ is not finite, two alternative possibilities remain. Either the aggregate expectation $\bar{p}^e_{n,t}$, given $\mathcal{F}_\sqcup$, diverges with non-zero probability, or it tends in distribution to a stochastic variable with a well defined limit distribution. This observation is interesting, as it allows the dynamics to remain random, even in the limit of an infinite number of agents. One can easily find examples where the mean expectation $E_{\phi_t(\theta)}[p^e_{t+1}|\mathcal{F}_t]$ is not finite, but the aggregate expectation $\bar{p}^e_{n,t+1}|\mathcal{F}_t$ has a non-degenerate limit distribution. This happens, for example, when the predictions $p^e_{t+1}|\mathcal{F}_t$ have a Cauchy type distribution. The existence of the mean of the predictions is neither implied by nor implies the existence of the mean of the parameters $\theta_{i,t}$.

4 Dynamic asset pricing

In this section we consider continuous choice dynamics in an asset pricing context, in which agents may employ predictor functions parameterized by θ. In each trading period agents can invest in two different assets: a risky asset and a risk free asset. The risk free asset yields a guaranteed return on investment r. Whereas, investing in the risky asset, quoted at price p_t and paying a stochastic dividend y_{t+1}, yields an uncertain return on investment. In order to find, in their objective, an optimal allocation between the two different assets, the investors have to form expectations about the risk (variance) and return of an investment that includes the risky asset.

It is assumed that each investor is a mean variance optimizer, maximizing the expected risk adjusted return on investment. The information publically available at time t consisting of past prices up to time $t-1$: $\mathcal{F}_t = \{p_{t-1}, p_{t-2}, \dots\}$. Let $E_{\theta,t}[\cdot]$ denote the subjective conditional expectation of investors who use the strategy with parameter θ at time t. It is assumed that the dividend process is IID with a mean $\bar{y}$ known to all traders, so that expected future dividend pay-offs are identical among trader types: $E_{\theta,t}[y_{t+1}] = \bar{y}$. Given the belief type θ, expectations about future price are given by a function of past prices and dividends:

$$E_{\theta,t}[p_{t+1} + y_{t+1}] = f_\theta(p_{t-1}, p_{t-2}, \dots) + \bar{y},$$

where $f_\theta(p_{t-1}, p_{t-2}, \dots)$ is the predictor employed by traders with belief parameter θ at time t.

During period t, ranging from time $t-1$ up to time t, traders are allowed to submit their demand as a function of the new price p_t to be quoted at time t. At that time the market maker will run through the order book in order to obtain the aggregate demand for the risky asset. Consequently, the equilibrium price is quoted, after which a new trading round starts. Without loss of generality we assume that the number of risky assets is constant over time i.e. there is no supply of outside shares. Under the assumption that risk aversion is constant and equal among traders the equilibrium equation for a market with n traders who, at time t, each select a strategy $\theta_{i,t}$ is given by:

$$(1+r)\, p_t = \frac{1}{n} \sum_{i=1}^{n} f_{\theta_{i,t}}(p_{t-1}, p_{t-2}, \dots) + \bar{y}. \tag{3}$$

This market clearing equation thus states that today's price is the sum of the expected future price and dividend, discounted by the risk free interest rate.

4.1 Asset pricing with continuous AR(1) predictors

In this section we discuss a well-known class of belief types, namely that of linear predictors, within the continuous beliefs framework. It will be shown that in such simple cases analytical expressions for the distribution of the belief parameter can easily be obtained. The dynamics of the beliefs distribution can be described by formulating the dynamical system in terms of a finite number of realized prices and the time-varying moments of the distribution of the belief parameter. We will focus on a simple case where all agents have first order linear autoregressive (AR) beliefs. For simplicity the constant term of this first order AR process will be held fixed in this example. Furthermore, we take $\phi(\theta) = 1$, providing a case for which an LTL approach would not be possible.

We follow () and assume that the dividends are generated by the following simple stochastic process:

$$y_t = \bar{y} + \varepsilon_t$$

where ε_t is IID with zero mean. Since all traders are assumed to have rational expectations about future dividend pay-offs, we have $E_t[y_{t+1}] = \bar{y}\ \forall t$.

Let us consider the case of AR(1) predictors now. That is, the traders believe in extrapolating linear trends:

$$E_{\theta,t}[p_{t+1}] = f_t(p_{t-1}\,|\theta) = \theta p_{t-1}.$$

Once today's price is observed, traders evaluate the performances of their predictor by considering squared prediction errors, that is, the utility function associated with belief type θ is given by:

$$U_t(\theta) = \alpha U_{t-1}(\theta) + \pi_t(\theta),$$

where

$$\pi_t(\theta) = -\left(E_{\theta,t-1}[p_t] - p_t\right)^2 = -\left(\theta p_{t-2} - p_t\right)^2, \tag{4}$$

and α denotes the memory parameter, which satisfies $0 \leq \alpha < 1$. The case $\alpha = 0$ corresponds to the case without memory.

After observation of the established price p_t, the distribution of beliefs is updated according to the continuous choice model. As is derived in the previous section, the new distribution characterizing the dispersion of belief types is then given by:

$$\phi_t(\theta) = \frac{\exp[\beta U_t(\theta)]}{Z_t} = \frac{[\phi_{t-1}]^\alpha \exp[\beta\pi_t(\theta)]}{Z_t'}.$$

After substituting the expression in Eqn. 4 for $\pi_t(\theta)$ one obtains:

$$\phi_t(\theta) = \frac{[\phi_{t-1}(\theta)]^\alpha \exp\left[-\beta\left(\theta p_{t-2} - p_t\right)^2\right]}{Z_t'} \tag{5}$$

Since the exponent contains only up to second order forms in θ with a negative coefficient for the quadratic term in θ, the distribution of beliefs in each period can be described by a normal distribution. This shows that the integral $\int \phi_{t-1}(\vartheta)\, \vartheta p_{t-1} \mathrm{d}\vartheta$ exists, and hence that the updated price p_t is a deterministic function of the past price p_{t-1} and the past beliefs distribution ϕ_{t-1}, which is itself determined by the price history. This result depends on the functional form of the predictors (linear) and on the choice of taking squared prediction error as fitness. The latter, however, is equivalent to risk adjusted profits, see the Appendix.

Inserting linear beliefs in terms of prices and rational beliefs about dividends into the equilibrium equation (3) yields:

$$(1+r)p_t = \int \phi_{t-1}(\vartheta)\, \vartheta p_{t-1} \mathrm{d}\vartheta + \bar{y}$$

Rearranging terms gives us the following expression for today's price:

$$p_t = \frac{\mu_{t-1} p_{t-1} + \bar{y}}{1+r} \tag{6}$$

where μ_{t-1} denotes the first moment of the distribution ϕ_{t-1} of the belief parameter θ. Note that the price only depends on the beliefs distribution through its mean value, or the 'average belief'. The other moments of the distribution have no effect on the price dynamics.

If we denote the mean and variance of this distribution by μ_t and σ_t^2 respectively, we have

$$\phi_t(\theta) = \frac{1}{\sqrt{2\pi}\sigma_t} \exp\left[-\frac{1}{2\sigma_t^2}(\theta - \mu_t)^2\right]. \tag{7}$$

If we define

$$m_t = \frac{p_t}{p_{t-2}}$$

and

$$s_t^2 = \frac{1}{2\beta p_{t-2}^2},$$

Eq. (5) gives

$$\phi_t(\theta) \propto \exp\left[-\frac{\alpha}{2\sigma_{t-1}^2}(\theta - \mu_{t-1})^2 - \frac{1}{2s_t^2}(\theta - m_t)^2\right]. \tag{8}$$

By comparing the coefficients of θ^2 and θ in the exponents in Eqs. (7) and (8), the mean μ_t and variance σ_t^2 can be seen to become

$$\mu_t = \alpha \mu_{t-1} \frac{\sigma_t^2}{\sigma_{t-1}^2} + m_t \frac{\sigma_t^2}{s_t^2} \tag{9}$$

$$\frac{1}{\sigma_t^2} = \frac{\alpha}{\sigma_{t-1}^2} + \frac{1}{s_t^2}. \tag{10}$$

The remaining terms independent of θ will be absorbed by the normalization factor Z_t. After introducing the auxiliary state variables $v_t = 1/\sigma_t^2$ and $q_t = p_{t-1}$ the dynamics become (in terms of past state variables):

$$\begin{cases} p_t &= (\mu_{t-1}p_{t-1} + \bar{y})/(1+r) \\ q_t &= p_{t-1} \\ \mu_t &= \alpha\mu_{t-1}\dfrac{v_{t-1}}{\alpha v_{t-1} + 2\beta q_{t-1}^2} + \left(\dfrac{\mu_{t-1}p_{t-1} + \bar{y}}{1+r}\right)\dfrac{2\beta q_{t-1}}{\alpha v_{t-1} + 2\beta q_{t-1}^2} \\ v_t &= \alpha v_{t-1} + 2\beta q_{t-1}^2. \end{cases}$$

Next we would like to examine the stability around the fixed point locally. The solution of $p_t = q_t = p^*$, $\mu_t = \mu^*$ is given by $v_t = v^*$ is $p^* = \bar{y}/r$, $\mu^* = 1$, $v^* = 2\beta(p^*)^2/(1-\alpha)$. Throughout we assume that the parameters satisfy the following conditions: $\alpha \in [0,1)$, $r \in [0,\infty)$ and $\beta \in [0,\infty)$. The following proposition gives the local stability conditions.

Proposition 3 *For $\alpha \in [0,1)$ and $r \in [0,\infty)$ and $\beta \in [0,\infty)$ the system is locally stable around the fixed point if $g(\alpha,r) > 0$, where*

$$g(\alpha,r) = -2 + r^2 + 3\alpha - \alpha^2 + r(2 - 2\alpha + \alpha^2).$$

Proof: The Jacobian matrix evaluated at the equilibrium is

$$J = \begin{pmatrix} \frac{1}{1+r} & 0 & \frac{p^*}{1+r} & 0 \\ 1 & 0 & 0 & 0 \\ \frac{1-\alpha}{(1+r)p^*} & -\frac{1-\alpha}{p^*} & \alpha + \frac{1-\alpha}{1+r} & 0 \\ 0 & 4\beta p^* & 0 & \alpha \end{pmatrix}$$

and the corresponding characteristic equation is

$$\left[(1+r)\lambda^3 - (2+\alpha r)\lambda^2 + \alpha\lambda + 1 - \alpha\right](\lambda - \alpha) = 0.$$

Since α is an eigenvalue, a necessary condition for stability is $\mid \alpha \mid < 1$. This condition is satisfied by the assumption that $\alpha \in [0,1)$. Application of the conditions for stability derived in Jury (1974) for the characteristic equation

$$(1+r)\lambda^3 - (2+\alpha r)\lambda^2 + \alpha\lambda + 1 - \alpha = 0 \tag{11}$$

of the remaining eigenvalues, gives the three conditions

$$\begin{aligned} r(1-\alpha) &> 0 \\ (2+r)(1+\alpha)/(1+r) &> 0 \\ -2 + r^2 + 3\alpha - \alpha^2 + r(2 - 2\alpha + \alpha^2) &> 0. \end{aligned}$$

The first two conditions are satisfied by the assumptions $0 \leq \alpha < 1$ and $r > 0$. The third condition is stated in the proposition. □

Note that the equilibrium price $p^* = \bar{y}/r$ and the switching parameter β do not affect local stability. From the dynamics it can be easily seen that the parameter β determines the scale of v_t, but has no effect on μ_t and hence p_t, since the scale of v_t cancels in the expressions for μ_t and p_t.

The condition $g(\alpha, r) = 0$, for $\alpha \in [0, 1)$ and $r > 0$, gives an explicit solution of the critical value $r_c(\alpha)$ of r given α:

$$r_c(\alpha) = \frac{-(1 + (1-\alpha)^2) + \sqrt{(1 + (1-\alpha)^2)^2 + 4(1-\alpha)(2-\alpha)}}{2}.$$

because the function $r_c(\alpha)$ is monotonically decreasing on $[0, 1)$, and since $r_c(0) = \sqrt{3} - 1$, for $r \in [0, \sqrt{3} - 1]$, this relation can be inverted to give a critical value $\alpha_c(r)$ of α given r:

$$\alpha_c(r) = 1 + \frac{1 - \sqrt{1 + 4r(1 - r^2)}}{2 - 2r}.$$

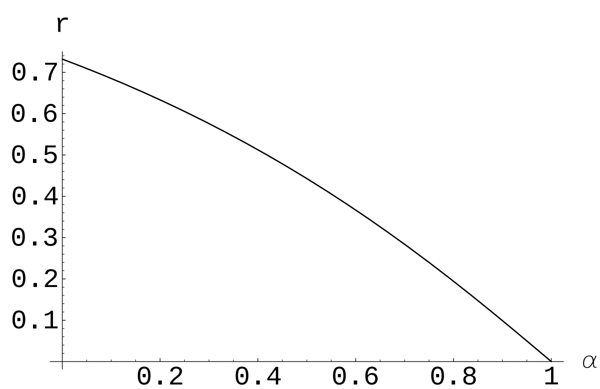

Figure 1: *Bifurcation curve in the (α, r)-plane. The system is locally stable for parameter values above (right) of this curve, and unstable for parameter values below (left) of it.*

Figure 1 shows the line $g(\alpha, r) = 0$ where the bifurcation occurs. The dynamics is locally stable around the fixed point for parameter values above and right of this bifurcation curve, and unstable for parameter values below or left of the bifurcation curve. Figure 2 shows some cross-sections of the basin of attraction of the fixed point in the (p_0, p_{-1})-plane, for parameter values where the dynamics are locally stable around the fixed point. The initial values for μ_t and σ_t^2 were set to their fixed-point values. The figure shows one cross-section for $r = 0.05$ and three cross-sections for $r = 0.10$, with increasing values of α. It can be observed that the fixed point is only locally stable, not globally. It can be observed that the basin of attraction grows when the memory parameter is increased, as one would intuitively expect from the stabilizing effect of the memory parameter.

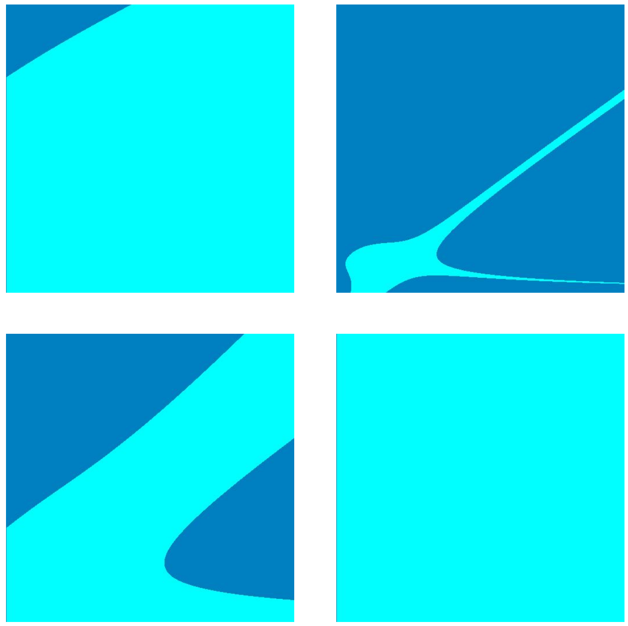

Figure 2: *Cross-sections of the basin of attraction of the fixed point. Shown are the regions* $[0, 100] \times [0, 100]$ *in the* (p_0, q_0)*-plane* $(q_0 = p_{-1})$*, while the initial values* μ_0 *and* σ_0^2 *were set to their fixed-point values. for parameter values (upper left)* $r = 0.05$, $\alpha = 0.98$, *(upper left)* $r = 0.10$, $\alpha = 0.90$, *(lower left)* $r = 0.10$, $\alpha = 0.95$ *and (lower right)* $r = 0.10$, $\alpha = 0.99$. *The fixed points are located at* (p^*, p^*) *where* $p^* = 1/r$ *(we used* $\bar{y} = 1$*).*

Although we derived the local stability conditions, we have not yet examined the type of bifurcation leading to instability. The following proposition states that this is a Neimark-Sacker bifurcation.

Proposition 4 *For* $\alpha \in [0, 1)$ *and* $r \in [0, \infty)$*, the bifurcation curve* $g(\alpha, r) = 0$ *corresponds to a Neimark-Sacker curve.*

Proof: The characteristic equation (Eq. 11) corresponds to a dynamical system at marginal stability if at least one eigenvalue λ passes through the unit circle. This happens when a pair of complex conjugate eigenvalues passes through the unit circle and/or at least one eigenvalue passes through either plus or minus one. To show that all points on the line $g(\alpha, r) = 0$ correspond to a Neimark-Sacker bifurcation, it suffices to show that there no eigenvalues can equal ± 1 on the bifurcation curve. Imposing that $\lambda = -1$ be a solution of Eqn. (11) leads to the (economically irrelevant) condition $r = -1$. Similarly, $\lambda = 1$ can only occur for $r = 0$, or $\alpha = 1$. These conditions nowhere coincide with the bifurcation curve, whence it follows that $g(\alpha, r) = 0$, for $\alpha \in [0, 1)$ and $r \in [0, \infty)$, is a Neimark-Sacker curve. □

Propositions 3 and 4 imply that for fixed values of the memory parameter α, the system is locally stable for sufficiently large values of the interest rate r. When the interest rate decreases the system becomes locally unstable and (quasi-)periodic behavior occurs via a Neimark-Sacker bifurcation.

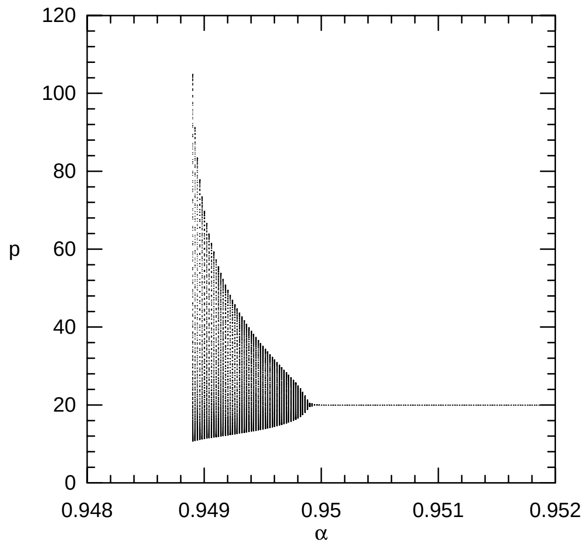

Figure 3: *Bifurcation diagram for* $r = 0.05$. *The critical value for this interest rate is* $\alpha_c(0.05) \simeq 0.949886$.

For a fixed interest rate $r > r_c(0) = \sqrt{3} - 1$ the system is locally stable for all $\alpha \in [0, 1)$. For $0 \leq r < r_c(0)$ the system can be locally stable or unstable depending on the memory parameter α. For sufficiently large values of α the system is locally stable. Lowering the memory parameter leads to (quasi-)periodic behavior via a Neimark-Sacker bifurcation. Figure 3 shows a bifurcation diagram as a function of α for $r = 0.05$. This diagram suggest that after the fixed point solution, also the quasi-periodic solutions become unstable if the memory parameter is decreased. The quasi-periodic solution becomes unstable and only bubble solutions appear to exist. A similar type of behavior was observed for other values of r. Note the increase of the scale just before break-up. The scale of the quasi-periodic solutions appears to tend to infinity just before α reaches the point where they become unstable.

Figure 4 shows some typical time series near the bifurcation curve and close before break-up of the quasi-periodic solutions for $r = 0.05$. Note that some asymmetry can be observed between the speed with which the price increases and decreases, especially just before the quasi-periodic motion becomes unstable ($\alpha = 0.94895$). The decreases are slightly faster. Also the sharp increase of the scale of the fluctuations just before break-up of the quasi-periodic motion is clearly visible. Finally, an increase can be observed in the typical time scale of the fluctuations just before break-up.

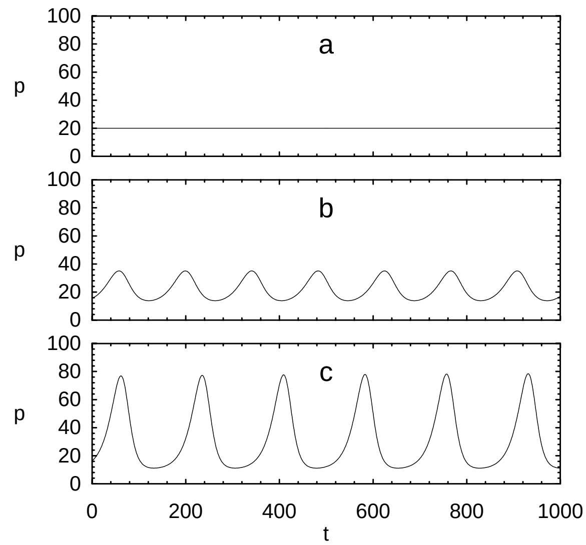

Figure 4: *Time series for* $r = 0.05$, *and* $\alpha = 0.95$ *(a),* $\alpha = 0.9450$ *(b) and* $\alpha = 0.94895$ *(c) respectively. The mean dividend* $\bar{y}$ *is taken to be* 1, *so that the fundamental price is* $p^* = \bar{y}/r = 20$.

5 Endogenous sources of randomness

Traditionally, randomness in asset pricing has always been associated with exogenous shocks, for example due to news that affects a company's future earnings, or fluctuations of the interest rate. The continuous choice framework provides at least three possible endogenous sources of randomness. Firstly, a finite number of traders gives rise to stochasticity, because traders are assigned a belief $\theta_{i,t}$ at random from the beliefs distribution $\phi_t(\theta)$, so that the aggregate expectation $\bar{p}^e_{n,t+1}$ for finite n is a stochastic random variable unless all predictions are identical with probability one. Secondly, for certain combinations of the utility function and the predictor function f_θ, the law of large numbers may not apply because the expectations $p^e_{i,t+1}$ do not have a mean. In those situations, the limiting price dynamics might either become undefined, or the system can tend to a stochastic dynamical system, in the limit where the number of agents tends to infinity. Thirdly, so far, we have implicitly assumed that no agents have a dominant market impact. This might not be the case, for example, if the wealth distribution among agents is fat-tailed. In that case market impact of the wealthiest agents might not disappear when the number of traders tends to infinity. In those cases the law of large numbers can not be applied, providing yet another possible source of endogenous randomness.

In this section we will concentrate on the first two sources of randomness, the finite number of traders, and the inherent randomness due to the nonexistence of the mean of $f_{\theta_{i,t}}(p_{t-1}, p_{t-2}, \dots)$. Although interesting, we consider the study of endogenous wealth effects beyond the scope of this paper and leave it for future research. The reason is that it is far from straightforward to incorporate the effects of the wealth distribution, since it evolves endogenously with the price dynamics, depending on the agent's strategy records. Note that in practice, an agent's wealth evolution will also depend on a number of exogenous factors, such as income and consumption.

5.1 Endogenous noise from a finite number of traders

Theorem 1 implies that when $E[f_{\theta_{i,t}}(p_{t-1}, p_{t-2}, \dots)]$ exists, the market equilibrium equation gives a unique price p_t in the limit of an infinite number of traders. For a finite number of traders, however, the price equation becomes dependent on the choices of individual agents and the price equation based on the mean of the expectations, $E[f_{\theta_{i,t}}(p_{t-1}, p_{t-2}, \dots)]$, can at best be a first order approximation. By examining how the actual mean prediction $\frac{1}{n}\sum_{i=1}^{n} f_{\theta_{i,t}}(p_{t-1}, p_{t-2}, \dots)$ is distributed if the number of traders is finite, we obtain a natural description of this endogenous randomness.

Consider a system with a finite number, n, of agents. The simplest assumption that can be made about the beliefs used by the agents to predict future prices, is that at each time step they independently choose one predictor from $\phi_(\theta)$. This amounts to assuming that Assumption 2 holds. The price equation then becomes

$$(1+r)p_t = \frac{1}{n}\sum_{i=1}^{n} f_{\theta_{i,t}}(p_{t-1}, p_{t-2}, \dots) + \bar{y},$$

where the $f_t(\theta_{i,t})$ are independent and identically distributed. Assuming the conditional mean μ_{t-1} and variance σ_{t-1}^2 of $f_{\theta_{i,t}}(p_{t-1}, p_{t-2}, \dots)$ to exist, the average by the central limit theorem is asymptotically $N(\mu_{t-1}, \sigma_{t-1}^2/n)$ distributed.

For the class of linear AR(1) predictor functions discussed before, expectations were normally distributed so that normality is obtained for finite n. The price equation becomes

$$(1+r)p_t = \mu_{t-1}p_{t-1} + \bar{y} + \frac{\sigma_{t-1}}{\sqrt{n}} p_{t-1}\epsilon_t, \tag{12}$$

with $\epsilon \sim N(0,1)$. In cases without memory $\alpha = 0$, we have $\sigma_{t-1}^2 = 1/p_{t-3}^2$, and the noise term can be written as $p_{t-1}p_{t-3}\epsilon_t/\sqrt{n}$. This example shows that a simple assumption about the distribution of beliefs among a finite number of traders already gives rise to a stylized fact such as conditional heteroskedasticity. In practice the mechanisms by which agents choose their beliefs may be dependent, for example because agents communicate about their choices. Local interaction among agents, for example on a trading floor, can give rise to positive correlation between their strategy choice. This can be dealt with in the model simply by replacing the n agents in the model by a smaller number of independent clusters of agents. This increases the variance of the noise term in Eq. (12). Allowing for temporal dependence on the idiosyncrasies of the agent's preferences (Assumption 1), also leads to the price dynamics as given in Eq. (12). However, the noise terms ϵ_t in that case can become temporally dependent as a result of the slowly changing preferences of agents.

5.2 Inherently random dynamics

Next we consider an example of the dynamics in a case where $f_{\theta_{i,t}}|\mathcal{F}_t$ does not have a mean, but aggregate expectations $\sum_{i=1}^{n} f_{\theta_{i,t}}|\mathcal{F}_t$ has a limit distribution when n tends to infinity.

The starting point is an the asset pricing model with agents that choose among constant predictors: $p^e_{i,t} = \theta_i$. For simplicity we put $\alpha = 0$ (no memory). For the utility function we take

$$U_t(\theta) = -\log(1 + (\theta - p_t)^2)$$

which, like the squared prediction error, is maximal for predictions θ that are equal to realized prices. For the beliefs distribution this gives

$$\phi_t(\theta) = \frac{\Gamma(\beta)}{\Gamma(\frac{1}{2})\Gamma(\beta - \frac{1}{2})}(1 + (\theta - p_{t-1})^2)^{-\beta},$$

from which it follows that,

$$\sqrt{2\beta - 1}\,(\theta_{i,t} - p_t) \sim t(2\beta - 1).$$

Thus, β is distributed symmetrically around $\theta = p_{t-1}$, and for $\beta < \frac{3}{2}$ the mean does not exist.

For $\beta = 1$, $\theta - p_t$ given $\mathcal{F}_t$ is Cauchy(0,1) distributed. Since this distribution is closed under averaging, this gives $p^e_{n,t+1} = \frac{1}{n}\sum_{i=1}^n \theta_{i,t} - p_t \sim$ Cauchy(0,1). The price equation becomes

$$(1 + r)p_t = p_{t-1} + \bar{y} + \eta_t,$$

where η_t is a Cauchy(0,1) distributed random variable, the pdf of which is $f_{\eta_t}(x) = (\pi(1 + x^2))^{-1}$. The price dynamics is stochastic, and the distribution of the noise term is independent of the number of agents.

The contributions from agents with extreme beliefs are not negligible in the limit where the number of traders tends to infinity. The result after aggregation is a price equation with a fat-tailed noise term. The price dynamics is stochastic, and the distribution of the noise term is independent of the number of agents.

6 Concluding remarks

We have proposed a new methodology for modeling the evolution of a heterogeneous beliefs distribution in a dynamic setting. The updating of beliefs takes place according to the continuous choice model, each time new information becomes publically available. Depending on the beliefs space and the beliefs distribution, aggregate expectations are either a deterministic or a stochastic functional of publically available information. This leaves open the possibility of both deterministic and stochastic evolution laws in dynamic feedback settings. The approach is illustrated using asset pricing models as a typical example. Explicit dynamic price equations are derived in this context, and the stability and long-term behavior of the dynamics examined. Examples of stochasticity as a result of the finite number of traders, as well as inherent stochasticity are examined in the asset pricing framework. Conditions are given under which the dynamics are of finite AR order, and it is shown that the model is observable from empirically observed prices.

Some remarks related to the large type limit approach developed in Brock *et al.* (2001) are in order here. There the market is approximated by a large, but finite, number of belief types randomly drawn from a fixed distribution. The authors then give conditions under which the aggregate behavior of all traders tends to a finite dimensional dynamical system. Our approach leads to the same dynamics in the deterministic case, when the opportunity distribution in the model is taken to be the distribution from which the strategies are drawn in the large type limit approach. Although the approaches give the same results in those cases, there are differences that are important from a methodological point of view. Apart from being more suitable for examining the evolution of beliefs distributions, the continuous choice framework is less restrictive than the large type limit approach, since it neither requires the opportunity function to be a probability density function, not does it exclude the possibility of stochastic dynamics.

The methodology developed here is general and can be applied in any context where aggregate expectations determine future states, while expectations themselves are determined by the beliefs distribution, which is adapted when new information becomes available. We consider the framework developed here as a first step towards describing continuous beliefs evolutions, the empirical implications of which are of potential use in a wide range of expectations based feedback models. The continuous choice approach enables one to obtain insights into ways in which endogenous randomness arises in agent models. This aspect is extremely relevant for obtaining meaningful model specifications for empirical data analysis.

References

Brock, W.A. and Hommes, C.H. (1997). Rational routes to randomness. *Econometrica*, **65,** 1059–1095.

Brock, W.A., Hommes, C.H. and Wagener, F.O.O. (2001). Evolutionary dynamics in financial markets with many trader types. Technical Report. CeNDEF Working paper 01-01, University of Amsterdam.

Dagsvik, J. K. (1994). Discrete and continuous choice, max-stable processes, and independence from irrelevant attributes. *Econometrica*, **62,** number 5, 1179–1205.

Hanemann, W. M. (1984). Discrete/continuous models of consumer demands. *Econometrica*, **52,** number 3, 541–561.

Jury, E. I. (1974). *Inners and Stability of Dynamical Systems.* John Wiley & Sons, New York.

McFadden, D. (1973). Conditional logit analysis of qualitative choice behavior. In *Frontiers in Econometrics*, *New York* (ed. P. Zarembka), pp. 105–142. Academic Press, NY.

Resnick, S.I. (1998). *A Probability Path.* Birkhäuser, Boston.