



UvA-DARE (Digital Academic Repository)

The time-variation of volatility and the evolution of expectations

van der Weide, R.

[Link to publication](#)

Citation for published version (APA):

van der Weide, R. (2012). *The time-variation of volatility and the evolution of expectations*. Amsterdam: University of Amsterdam.

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Chapter 5

Heterogeneity as a natural source of randomness

5.1 Introduction

Though it takes little to realize that expectations about future prices are essential for investment decisions made by market participants, it is difficult to study the evolution of expectations explicitly. 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, which, at best, reflect market aggregated expectations rather than individual expectations.

Recent results indicate that, to a large extent, the dispersion of beliefs among market participants bears important consequences to the behavior of the aggregates over time. In the recent literature on heterogeneity several theoretical models have been proposed and a number of relations between heterogeneity and stylized facts were derived and validated empirically. Shalen (1993) and Michaely and Vila (1996), for example, concluded that dispersion of beliefs enhances both trading volume and volatility. This supports the empirical evidence reported by Gallant *et al.* (1992) of a positive correlation between volume and volatility. Recent work also indicates that heterogeneity of expectations might be responsible for several other empirical observations. For example, Diether *et al.* (2002) presented evidence that the dispersion of opinions on future returns affects future earnings negatively, while Ziegler (2002) found that heterogeneity of beliefs can lead to the well-known “smile

effect” in implied volatility derived from option prices.

Apparently, economic observables depend not only on the average belief, but, more generally, on the distribution of beliefs among market participants. Being both omnipresent and hidden, the dynamics of dispersed beliefs is of our key interest in this chapter. The objective is to model the evolution of beliefs explicitly. While doing so, we wish to keep track of the unpredictable nature of individual preferences and relate it to the unpredictable nature of observables. The latter is important for investigating the extent to which economic observables are affected by uncertainty of choice at the agent level. Together, these objectives motivate the introduction of the concept of a *continuous beliefs system* (CBS). By formulating economic dynamic models in terms of a CBS it is possible to address questions directly related to the beliefs distribution and the endogenous randomness in the dynamics, such as: (i) To what extent does the degree of heterogeneity interact with observables such as prices? (ii) Which aspects of the beliefs distribution matter most for the random dynamics? (iii) Do prices and beliefs evolve similarly in small and large markets? (iv) How do differences in market impact and dependence among agents’ choices affect the joint dynamics of prices and beliefs?

Among the first to provide an analytical framework in which agents adapt their beliefs over time are Brock and LeBaron (1996) and Brock and Hommes (1997). The agents’ incentives to switch beliefs are provided by the observed differences in past performance such as realized profits. Although, in this setup agents can choose from various beliefs, or expectation functions, the number of strategies available to the agents is finite. Typical applications are concerned with the simple case in which there are only two different belief types (see e.g. Brock and Hommes, 1998). Evidently a small number of belief types is insufficient to obtain a realistic beliefs distribution. Since each fraction of the population associated with a belief type becomes a state variable, using the Brock and Hommes (1997) framework to build a model with a large number of beliefs will typically lead to analytically intractable dynamics.

Recently, Brock *et al.* (2005) considered the price dynamics of a market with a continuum of traders in a so-called large type limit (LTL) where the number of strategies available to the agents tends to infinity. The belief types are drawn at random from a given distribution of possible types. Under fairly general conditions this (random) sequence of deterministic dynamical systems gives rise to unique deterministic limiting dynamics in the LTL. In this way, an LTL provides insights into the dynamics of a large market in the presence of a large number of strategies.

Simply because the limiting deterministic price dynamics was of their main interest, Brock *et al.* (2005) assigned a merely supporting role to the beliefs distribution. The evolution of the beliefs distribution remains more or less hidden because the state variables describing the beliefs distribution are eliminated in the derivation of the price dynamics. Also, by considering a continuum of agents from the start, their approach implicitly assumes that preferences of individual agents are negligible. As a result the concept of an LTL is not very convenient for our purpose: modeling the evolution of beliefs, while taking into account effects of the unpredictable nature of individual preferences. With this in mind, a CBS follows a different route. In contrast with an LTL, which is based on a continuum of agents who can choose from a finite (but increasing) number of strategies, a CBS is based on a continuum of strategies available to any number of agents. The CBS dynamics of the beliefs distribution and the economic observables can be explicitly formulated for any number of agents with arbitrary market weights. This permits several scenarios which are hard to study using an LTL. For example, since the CBS approach does not assume individual preferences to be negligible, it can be used to examine large market (many agent) limits in which the market is dominated by a small number of agents, or where the agents' expectations are strongly correlated.

A CBS contains the following ingredients. The basic concept underlying a CBS is a *beliefs space*, Ω , in which a class of possible point predictors is represented by a continuous (either scalar or vector valued) parameter θ . At each time t agents form beliefs regarding future economic variables by deciding on a single belief parameter θ in the beliefs space. They do so by evaluating the possible strategies, given the information available prior to time t , using a *performance measure*. This performance measure might be based, for example, on the history of past prediction errors, or on profits a strategy would have realized in the past. In each period individual agents choose the strategy which optimizes their expected subjective utility. Due to differences in tastes, they may differ in the strategies used. The central object in a CBS is a time dependent probability density function $\phi_t(\theta)$ on Ω , called the *beliefs distribution*. The belief, (or predictor, or strategy) $\theta_{i,t}$ actually employed by agent i at time t , from the modeler's point of view, is considered to be a random variable distributed according to the beliefs distribution. This can be thought of as representing the empirical fact that even when agents have identical information, they may still have differences of opinion regarding future revenues (see e.g. Frankel and Froot, 1990; Kandel and Pearson, 1995). The specific form of the beliefs distribution given the history of realized past observables, can be expressed in terms of past performances using

a continuous choice model. Given the individual beliefs of agents just prior to time t , the individual net demand functions collectively determine the next observable (price) to be quoted publically. This involves some *aggregation mechanism*, such as a market maker setting the price at a zero aggregated net demand value. After the observable becomes public, the performances of strategies can be re-evaluated, a new beliefs distribution arises, etc. From a dynamical point of view, the ongoing evaluation of strategies by agents as new information becomes available results in the co-evolution of the distribution of beliefs and observable aggregates.

In a CBS, the beliefs distribution is a deterministic function of past observables. The strategies used by the agents are modeled as random variables, distributed according to the beliefs distribution, which provides a natural mechanism for endogenous randomness. We refer to this as a ‘natural source of randomness’, as it can be associated naturally with the uncertainty of choice. In fact this natural source of randomness can help explain part of the market fluctuations we observe in daily life. For example, the excess volatility as observed in financial markets (see Shiller, 1981) might to some extent be ‘natural’. In some of the applications presented later it will become clear that most of the nontrivial structure arising from heterogeneous beliefs feedback loops resides in the noise rather than in the deterministic part of the dynamics.

One might expect the randomness to disappear from the dynamics as the number of agents tends to infinity. However, it is not hard to find mild conditions under which the choices of individual agents do not average out to give a deterministic law for the aggregate observables in the limit as the number of agents tends to infinity. For example, if the dependence among the agent’s choices is sufficiently large, aggregates may remain random variables even for an infinity of agents. Also, for some combinations of the class of predictors and the performance measure, the conditions for the law of large numbers may not be satisfied, providing another source of randomness which does not vanish if the number of agents tends to infinity.

In general, a CBS gives rise to a Random Dynamical System (RDS) where the randomness carries a natural interpretation. For convenient, but often typical, choices of the performance measure and the class of predictors, the CBS reduces to a low dimensional RDS. The state variables of this RDS are directly related to prices and characteristics of the beliefs distribution, such as the degree of heterogeneity. Since the CBS approach allows stochastic price dynamics to be derived explicitly, it can be used as an analytic alternative to numerical studies based on computationally intensive agent models (for an overview,

see LeBaron, 2000).

The remaining sections are organized as follows. In section 5.2 the concepts of a beliefs distribution and the continuous choice model are introduced. Section 5.3 describes a continuous beliefs system (CBS) in which the co-evolution of beliefs and public information such as market prices is taken into account. As an illustration the implied dynamics in a standard asset pricing context is examined by means of some stylized examples in section 5.4. In section 5.5 the mechanism by which endogenous noise arises from the continuous choice dynamics is described, and the role of the number of agents and their dependence considered. Section 5.6 summarizes and discusses the results.

5.2 Continuous beliefs distributions

In this section we describe how continuous beliefs distributions are obtained with the continuous choice model, which can be seen as a natural generalization of the well-known discrete choice model. The discrete version has among others been employed in an economic dynamic context by Brock and Hommes (1997). For clarity of exposition we will discuss both choice models in an agent based economic dynamic setting, starting with the more familiar discrete choice model.

Agent based models represent market participants as (a typically 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 of strategies, s_1, \dots, s_m say, then agents employing strategy s_i , $i = 1, \dots, m$ are said to be of type i . McFadden (1973) derived an expression for the probability P_i that an individual will select 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 subjective expected utility. Under the assumption that the disturbances of the utility function follow an extreme value distribution, it can be shown that this leads to the *multinomial logit model*:

$$P_i = \frac{e^{\beta U(s_i)}}{\sum_{l=1}^m e^{\beta U(s_l)}}$$

where $U(s_l)$ is the utility associated with alternative l . 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, presuming 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 measurable space Ω , containing all possible strategies that can be selected by the agents. We will refer to Ω as the *beliefs space*. Each possible choice, that is, each element θ in Ω uniquely represents a possible strategy agents can choose from. Note that the choice of the beliefs space is not unique, since any one-to-one transformation of the beliefs space Ω into another space, Ω' , say, will again yield a suitable beliefs representation. To ensure that integrals over the beliefs space can be defined independently of the chosen representation, we explicitly denote the integration measure on Ω by ν .

In analogy with the discrete choice model, we wish to represent the diversity of belief types by a probability measure 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. As in the discrete choice setting, it is convenient to adopt a random utility approach (Hanemann, 1984; Dagsvik, 1994; Resnick and Roy, 1994). The random part of the utility function of an agent affects the strategy a particular agent considers optimal. Therefore, the strategies employed by individual agents in a random utility framework are random variables.¹ The continuous choice analogue of the multinomial logit model is the *continuous logit model* (see e.g. Ben-Akiva and Watanatada 1981; Dagsvik, 1994). The probability that an agent selects a strategy in a subset A of Ω is given by

$$P(A) = Z^{-1} \int_A e^{\beta U(\theta)} \nu(d\theta) \quad (5.1)$$

with $Z = \int_{\Omega} e^{\beta U(\theta)} \nu(d\theta)$. As in the discrete choice setting, β represents the intensity of choice.

Whenever ν is continuous, a corresponding density $v(\theta)$ exists such that $\nu(d\theta) = v(\theta) d\theta$. In that case, assuming that $U(\theta)$ is also continuous, the beliefs distribution P has an associated pdf, denoted by $\phi(\theta)$, given by

$$\phi(\theta) = Z^{-1} e^{\beta U(\theta)} v(\theta), \quad (5.2)$$

¹Note that random utility does not imply that individual agents perceive their own utility functions to be random, only that they are random to the econometrician.

with $Z = \int_{\Omega} e^{\beta U(\vartheta)} v(\vartheta) d\vartheta$. For ease of presentation, in the subsequent sections we only consider cases where ν is continuous, and hence $v(\theta)$ exists.

The function $v(\theta)$ is nonnegative, and can be used to put different weights on different parts of the beliefs space. We refer to $v(\theta)$ as the *opportunity function*. In the case where a particular representation of the beliefs space has been fixed, the opportunity function can be thought of as reflecting the *a priori* faith of individuals in parameters within certain regions of the parameter space. Small values of $v(\theta)$ in a certain region then reflect the agents' tendency of avoiding parameter values in that region. For example, if in a particular representation it is reasonable to assume aversion against using extreme strategies, this can be represented by small values of $v(\theta)$ for these strategies. In that case, such "extreme" strategies need to outperform more common strategies to a large extent before agents are likely to use them.

Although derived differently, in the LTL framework (Brock *et al.*, 2005) extensive use is being made of a probability measure with the functional form given by equation (5.1), all market averages in the large type limit being defined with respect to this measure. For an LTL, the measure ν represents the probability measure on Ω from which an increasing set of strategies is chosen at random. However, as it turns out, in the CBS framework it is not necessary for ν to be normalizable. For example, if $\Omega = \mathbf{R}^m$, $v(\theta) = 1$ on Ω and $U(\theta)$ is quadratic in θ with a single maximum, then $\phi(\theta)$ is a multivariate normal probability density function. In general, the probability measure P given in equation (5.1) is a well-defined pdf if and only if $Z = \int_{\Omega} e^{\beta U(\vartheta)} \nu(d\vartheta)$ is positive and finite, i.e. if $e^{\beta U(\vartheta)}$ is ν -integrable. Clearly, when ν is non-normalizable it can no longer be interpreted as a probability measure. An alternative interpretation is the following. For two disjoint subsets A and B in Ω , with finite integration measures $\nu(A)$ and $\nu(B)$, the ratio $\nu(A)/(\nu(A)+\nu(B))$ can be interpreted as the conditional probability that an agent chooses a parameter value in A , in the absence of any information (U constant), conditional on this choice being either in A or in B .

5.3 Continuous beliefs systems

In this section we discuss the co-evolution of economic observables and the beliefs distribution in a CBS. Next to the beliefs space and the beliefs distribution, two additional ingredients are added. Firstly, we assume that agents evaluate strategies according to some *performance measure*, which might for example be based on past prediction errors,

or profits a strategy would have realized in the past, given the information available now (ex post profits). Secondly, a market mechanism, or more generally, an *aggregation mechanism* is required which translates the individual beliefs of agents into publically available information such as prices. For example, in a cobweb framework the price equation might have the form

$$p_t = D^{-1} \left(\frac{1}{n} \sum_{i=1}^n S(p_i^e(\theta_{i,t-1})) \right),$$

where $S(p_i^e(\theta_{i,t-1}))$ is firm i 's supply and $\frac{1}{n} \sum_{i=1}^n S(p_i^e(\theta_{i,t-1}))$ denotes the average supply, and $D(p)$ demand at a given price p .

For clarity of exposition this section is presented against the background of a standard dynamic asset pricing framework. Today's price will be assumed to equal the market's expected present value of tomorrow's pay-off. Rather than averaging over functions of beliefs, as in the cobweb framework, prices are conveniently obtained by averaging directly over the different beliefs. The modifications required for applications in more general dynamic settings are fairly straightforward. Throughout, time will be considered to be discrete. For a recently proposed generalization of discrete choice models to continuous time, we refer the interested reader to Dagsvik (2002).

The price of the risky asset at time t will be denoted by p_t . For simplicity we discuss the case where the agents are myopic. Prior to time t , they form expectations about the price of the asset at time $t + 1$ (the next time they can sell the asset) including possible dividends paid in the period between time t and $t + 1$. The information available to agents just before time t is denoted by \mathcal{F}_t . In simple cases the information set could consist of a historic record of past prices up to and including p_{t-1} , but in general it might include exogenous variables, such as the interest rate. The possible strategies from which the agent can choose to predict future prices are represented by a function f_θ of the observables in the information set, parameterized by θ . For a given set of information, we can then consider this predictor as a function of θ . In our examples, the available information on which predictions are conditioned is strictly the observed past price history:

$$p_{t+1}^e(\theta) = f_\theta(p_{t-1}, p_{t-2}, \dots).$$

Thus, the prediction of p_{t+1} made by agent i , based on price information up to and including p_{t-1} , using strategy $\theta_{i,t-1}$, is denoted by $p_{t+1}^e(\theta_{i,t-1})$.

To give an example of a beliefs space, the class of d -th order linear predictors consists

of all predictors of the form

$$p_{t+1}^e(\theta) = \theta_0 + \theta_1 p_{t-1} + \dots + \theta_d p_{t-d}.$$

In this case the beliefs are represented in \mathbf{R}^{d+1} and the beliefs distribution is a probability distribution on this space.

The pdf of the time dependent beliefs distribution conditional on \mathcal{F}_t is denoted by $\phi_{t-1}(\theta)$. Schematically, the expectations feedback can then be represented as

$$\dots \rightarrow p_{t-1} \rightarrow \phi_{t-1}(\theta) \rightarrow p_t \rightarrow \phi_t(\theta) \rightarrow \dots$$

Notice that, regardless of the details of the price formation mechanism, it is possible to investigate how a newly established price p_t affects the beliefs distribution. After p_t becomes part of the information set, agents can re-evaluate the available strategies, and at this point the continuous choice model can be invoked to obtain an expression for the new beliefs distribution:

$$\phi_t(\theta) = Z_t^{-1} v(\theta) e^{\beta U_t(\theta)} \quad (5.3)$$

the time dependent analogue of equation (5.2).

Typically, we consider cases where the utility function $U_t(\theta)$ is based on performance or fitness measure of the strategies θ , such as last period's net ex post profits or squared prediction errors. We take into account dependence on further past performances by introducing memory in the model. The evolution of the fitness measure for strategy θ can for example be modeled as:

$$U_t(\theta) = \alpha U_{t-1}(\theta) + (1 - \alpha) \pi_t(\theta), \quad (5.4)$$

where $\pi_t(\theta)$ is the performance in period t , and $\alpha \in [0, 1)$ is a memory parameter.² The utility function then becomes a geometrically weighted sum of ex post performances of strategy θ . An interpretation of the memory parameter α is the following; when analysts or traders consider new strategies it is common to perform back-testing, that is, to test the candidate strategies using historic data. Strategies that would have performed best in the past are more likely to be selected for future trading. The traders are likely to put more

²Alternatively one might define $U_t(\theta) = \alpha U_t(\theta) + \pi_t(\theta)$, but this is in fact equivalent, since for $\alpha \in [0, 1)$ this gives a utility differing by a factor $(1 - \alpha)$ which can be absorbed by β in the beliefs distribution in equation (5.3).

weight on more recent observations, depending on the time scale on which they perceive their world to be approximately stationary.

By substituting equation (5.4) into equation (5.3) the effect of memory can be written as

$$\begin{aligned}\phi_t(\theta) &\propto v(\theta)e^{\alpha\beta U_{t-1}(\theta)+(1-\alpha)\beta\pi_t(\theta)} \\ &\propto [v(\theta)]^{1-\alpha}[\phi_{t-1}(\theta)]^\alpha e^{(1-\alpha)\beta\pi_t(\theta)},\end{aligned}\tag{5.5}$$

which gives an update of the beliefs distribution in terms of the previous beliefs distribution and the last performance measure. Note that in general, since $\phi_t(\theta)$ is an infinite dimensional state variable, this may lead to a very complicated dynamical system. In general the models thus obtained need not be analytically tractable, and to solve the dynamic equations numerical methods may be required. However, in some cases the evolution of $\phi_t(\theta)$ can be completely described by a finite number of variables such as its first k moments, in which case it becomes finite dimensional. This class of models can be used to obtain insights into the interaction of observables and the beliefs distribution analytically. Therefore, in this chapter we will focus on models that are analytically tractable.

Next we consider the formation of the observables p_t , for a finite number, n , of agents, each of which is assigned a strategy according to the previous beliefs distribution $\phi_{t-1}(\theta)$. As mentioned before, some aggregation mechanism is required for determining the next observable from the individual beliefs of the agents. We assume that each predictor implies a unique associated excess demand function, and that given the beliefs of all agents, a unique equilibrium price p_t can be set, such that the market clears. Under the standard assumptions of mean-variance optimization today's market clearing price can be seen to equal the present value of aggregate beliefs concerning future prices and dividends:

$$(1+r)p_t = \frac{1}{n} \sum_{i=1}^n p_{t+1}^e(\theta_{i,t-1}) + \bar{y}\tag{5.6}$$

where n is the number of agents, $p^e(\theta_{i,t-1})$ denotes the expected price of individual i based on past observables available up to and including $t-1$, and \bar{y} expected future dividends. Since our interest in the asset pricing model is strictly concerned with its illustrative use, we refer the reader to e.g. Brock and Hommes (1998) for details on the dynamic asset pricing framework.

In interpreting the price equation it is essential to realize that the $\theta_{i,t-1}$ are random variables, representing the individual strategies chosen by the agents. The dynamical system for a finite number of agents is thus stochastic. In the limit where the number of

agents tends to infinity, the dynamics can, under certain conditions, become deterministic.

Consider the following assumption:

Assumption 5.1 (*Cross-sectional Independence*) *The strategies $\theta_{i,t-1}$ employed by agent i at time t , for each fixed time t are independent random variables, distributed according to the population distribution of beliefs prior to time t , $\phi_{t-1}(\theta)$.*

Assuming independence over agents might seem 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 temporal independence of the idiosyncratic noise terms of each agent over time. This stronger assumption 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 following theorem is concerned with the almost sure behavior of the model in the limit where the number of agents tends to infinity.

Theorem 5.1 (*Law of Large Numbers*) *Given \mathcal{F}_t , under Assumption 5.1, if n tends to infinity, the aggregate*

$$\bar{p}_t^e = \frac{1}{n} \sum_{i=1}^n p_{t+1}^e(\theta_{i,t-1})$$

converges a.s. to

$$E_t [p_{t+1}^e(\theta_{i,t-1})],$$

if and only if $E_t [p_{t+1}^e(\theta_{i,t-1})] < \infty$.

Proof: By Assumption 5.1, the strategies $\theta_{i,t-1}$ conditionally on \mathcal{F}_t are IID random variables in Ω , which implies that the corresponding predictions $p_{t+1}^e(\theta_{i,t-1})$ of agents are also IID. The result is immediate from Kolmogorov's strong law of large numbers for the IID random variables $p_{t+1}^e(\theta_{i,t-1})$, given \mathcal{F}_t (see e.g. Resnick, 1998, p. 220). \square

Note that the mean expected price, whenever it is finite, can be expressed as

$$E_t [p_{t+1}^e(\theta_{i,t-1})] = \int_{\Omega} p_{t+1}^e(\vartheta) \phi_{t-1}(\vartheta) d\vartheta. \quad (5.7)$$

Theorem 5.1 states that, given the information public at time t , a necessary and sufficient condition for the aggregate expectation about p_{t+1} to converge a.s. to the mean expectation

over the beliefs distribution, is the existence of this mean. This result has the following corollary concerning the dynamics.

Corollary 5.1 (*Deterministic Dynamics*) *Under Assumption 5.1 the observable p_t in the limit $n \rightarrow \infty$ almost surely tends to a deterministic function of the information \mathcal{F}_t available at time t if and only if $E_t[p_{t+1}^e(\theta_{i,t-1})] < \infty$.*

This deterministic limiting dynamics can be interpreted as a first order approximation to the price dynamics for a system with a large but finite number of agents. Later it will become clear that the random deviations around the deterministic limit in general can not be represented by a random variable with constant (state-independent) distributional properties.

Notice that discrete choice models can be considered as special cases of a CBS (when allowing for discrete opportunity distributions) in which agents can only choose among a finite number of alternative strategies θ_l , $l = 1, \dots, m$. Provided that the expected future prices of each of those strategies are 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. A discrete choice setting with a continuum of agents, as considered by Brock and Hommes (1997), should thus always lead to functional determinism in the absence of exogenous noise provided that $p_{t+1}^e(\theta_l)$ is finite for each $l = 1, \dots, m$.

At this point let us briefly compare equation (5.7) with its LTL analogue. The population mean of the predictor function $p_{t+1}^e(\theta)$ can be written as

$$E_t[p_{t+1}^e(\theta)] = \int_{\Omega} p_{t+1}^e(\vartheta) \phi_{t-1}(\vartheta) d\vartheta = \frac{\int_{\Omega} e^{\beta U_{t-1}(\theta)} p_{t+1}^e(\vartheta) v(\vartheta) d\vartheta}{\int_{\Omega} e^{\beta U_{t-1}(\theta)} v(\vartheta) d\vartheta}. \quad (5.8)$$

This expression can be seen to be of the same analytic form as the population average given by Brock *et al.* (2005). An important difference, however, is that we do not require $v(\theta)$ to integrate to one. That is, $v(\theta)$ itself is not required to be normalizable, and instead we only require $\phi_{t-1}(\theta)$ to be a well-defined pdf (i.e. $e^{\beta U_{t-1}(\theta)} v(\theta)$ should be integrable). In this sense, equation (5.8) provides a generalization of the analogous expression obtained with the LTL approach. This indicates that an LTL has an associated CBS, but that the converse does not always hold. In fact some of the example CBSes that will be discussed later use a non-normalizable function $v(\theta)$, and hence do not have an LTL representation. Furthermore, even if the beliefs distribution is well-defined, the mean expected price given

in equation (5.7) need not exist. In that case the dynamics does not allow for a deterministic LTL, but as will be shown later, the associated CBS might still be meaningful, with the price being a random variable. While an LTL always gives rise to a deterministic dynamical system, a CBS gives rise to an RDS, containing deterministic dynamics as a special case.

Before moving to a concrete example, note that we only consider point predictors in this chapter. However, some price formation mechanisms can not be formulated in terms of point predictors only. For example, Guesnerie (2002) considers agents who have subset predictors (e.g. an interval) in mind rather than point predictors. Under certain conditions, the agents, by a rationality assumption and a common knowledge argument, can find agreement on a unique trading price. Such a mechanism could be incorporated in a CBS framework, provided that the more general predictors can be represented by a finite number of parameters, and the price agents eventually agree to trade on can be expressed explicitly in terms of the agents' individual beliefs parameters.

5.4 Example: Asset pricing with first order linear beliefs

We consider an example of a CBS in which agents believe that first order linear rules provide adequate predictions. The role of the performance measure is investigated by focusing on two cases: squared prediction errors and squared logarithmic prediction errors, leading to normal and log-normal beliefs distributions respectively.

For illustrational purposes we focus on a simple class of linear predictors. Agents choose a strategy θ that represents their perceived future growth rate. The expected price at time $t + 1$ associated with predictor θ is then

$$p_{t+1}^e(\theta) = \theta p_{t-1}. \quad (5.9)$$

The utility function associated with belief type θ is given by equation (5.4). The price equation, (5.6), in this particular case becomes, in the limit of infinitely many agents:

$$\begin{aligned} (1+r)p_t &= \int \vartheta p_{t-1} \phi_{t-1}(\vartheta) d\vartheta + \bar{y} \\ &= \mu_{t-1} p_{t-1} + \bar{y}, \end{aligned} \quad (5.10)$$

where μ_{t-1} represents the average belief parameter at time $t - 1$. In the next subsections we consider the updating of the beliefs distribution for two different performance measures.

Throughout we assume that the parameters satisfy the following conditions: $\alpha \in [0, 1)$, $r \in [0, \infty)$ and $\beta \in [0, \infty)$.

5.4.1 Updating according to squared prediction errors

First we examine the case where $\Omega = \mathbf{R}$, and the performance measure is minus the squared prediction error:

$$\pi_t(\theta) = -(p_t^\varepsilon(\theta) - p_t)^2 = -(\theta p_{t-2} - p_t)^2. \quad (5.11)$$

After the price p_t has been realized and observed, the distribution of beliefs is updated according to the continuous choice model. The new distribution describing the dispersion of belief types is then given by equation (5.5). Upon substitution of the expression for $\pi_t(\theta)$ into equation (5.5), we obtain:

$$\phi_t(\theta) \propto [\phi_{t-1}(\theta)]^\alpha \exp[-\beta(1 - \alpha)(\theta p_{t-2} - p_t)^2]. \quad (5.12)$$

Note that we have chosen a constant opportunity function, i.e. $v(\theta) = 1$, representing the simplest case where agents hold no aversion against extreme parameter values.³

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. Here we have implicitly assumed that $\phi_{t-1}(\theta)$ is also normal, which can be justified by assuming that the dynamics has been running since the infinite past.

If we denote the mean and variance of $\phi_t(\theta)$ by μ_t and σ_t^2 respectively, we can write

$$\phi_t(\theta) = \frac{1}{\sqrt{2\pi}\sigma_t} \exp\left[-\frac{1}{2\sigma_t^2}(\theta - \mu_t)^2\right]. \quad (5.13)$$

Together with equation (5.12) this gives

$$\phi_t(\theta) \propto \exp\left[-\frac{\alpha}{2\sigma_{t-1}^2}(\theta - \mu_{t-1})^2 - \beta(1 - \alpha)(\theta p_{t-2} - p_t)^2\right]. \quad (5.14)$$

By comparing the coefficients of θ^2 and θ in the exponents in equations (5.13) and (5.14),

³In fact, since $v(\theta) = 1$ is not normalizable, this is an example of a CBS which can not be obtained directly using the LTL approach. However, using the LTL approach, one might arrive at the CBS dynamics by taking $v(\theta)$ to be the pdf of a $N(0, s^2)$ distributed random variable, and taking the limit where s^2 tends to infinity.

the mean belief parameter μ_t , and the variance of the belief parameter, σ_t^2 , can be seen to evolve according to

$$\begin{aligned} \frac{\mu_t}{\sigma_t^2} &= \alpha \frac{\mu_{t-1}}{\sigma_{t-1}^2} + 2\beta(1-\alpha)p_t p_{t-2} \\ \frac{1}{\sigma_t^2} &= \frac{\alpha}{\sigma_{t-1}^2} + 2\beta(1-\alpha)p_{t-2}^2. \end{aligned}$$

The remaining factor, which is independent of θ , is absorbed into the normalization factor Z_t . Notice that the average belief parameter μ_t and belief variance σ_t^2 completely describe the evolution of beliefs and are fully determined by past prices.

The full deterministic CBS, which also includes the price equation given in equation (5.10), becomes:

$$\begin{aligned} (1+r)p_t &= \mu_{t-1}p_{t-1} + \bar{y} \\ \frac{\mu_t}{\sigma_t^2} &= \alpha \frac{\mu_{t-1}}{\sigma_{t-1}^2} + 2\beta(1-\alpha)p_t p_{t-2} \\ \frac{1}{\sigma_t^2} &= \alpha \frac{1}{\sigma_{t-1}^2} + 2\beta(1-\alpha)p_{t-2}^2. \end{aligned} \tag{5.15}$$

The fixed point, which is the solution of $(p_t, \mu_t, \sigma_t^2) = (p^*, \mu^*, \sigma^{2*})$, is given by

$$(p^*, \mu^*, \sigma^{2*}) = \left(\frac{\bar{y}}{r}, 1, \frac{1}{2\beta} \right).$$

In section 5.4.3 the local stability properties of the dynamics near the fixed point are examined.

5.4.2 Updating according to squared logarithmic prediction errors

Next we consider the case with $\Omega = \mathbf{R}^+$, with a performance measure which taken to be minus the squared logarithmic prediction error:

$$\pi_t(\theta) = -(\ln p_t^e(\theta) - \ln p_t)^2 = -(\ln \theta - \ln p_t + \ln p_{t-2})^2. \tag{5.16}$$

An argument for using logarithmic prediction errors rather than just mean squared prediction errors is that this error measure is independent of the price level. This leads to dynamics which scales with the price level and is equivalent before and after possible stock splits.

Using a uniform opportunity function again, i.e. $v(\theta) = 1$, substitution of the expression for $\pi_t(\theta)$ as in equation (5.16) into equation (5.5) gives:

$$\phi_t(\theta) = Z_t'^{-1}[\phi_{t-1}(\theta)]^\alpha \exp[-\beta(1-\alpha)(\ln \theta - \ln p_t + \ln p_{t-2})^2], \quad (5.17)$$

where again Z_t' is a normalization factor independent of θ , but not necessarily equal to Z_t in equation (5.3).

In this case the exponent contains up to second order forms in $\ln \theta$ with a negative coefficient for the quadratic term in $\ln \theta$, which indicates that the distribution of beliefs in each period can be described by a log-normal distribution of the form

$$\phi_t(\theta) = \frac{1}{\sqrt{2\pi}\sigma_t\theta} e^{-\frac{(\ln \theta - \mu_t)^2}{2\sigma_t^2}} = \frac{1}{\sqrt{2\pi}\sigma_t} e^{-\frac{(\ln \theta - \mu_t)^2}{2\sigma_t^2} - \ln \theta}. \quad (5.18)$$

The evolution of the beliefs distribution follows from combining equations (5.17) and (5.18), which results in

$$\phi_t(\theta) \propto \exp\left[-\alpha \ln \theta - \frac{\alpha}{2\sigma_{t-1}^2}(\ln \theta - \mu_{t-1})^2 - \beta(1-\alpha)(\ln \theta - \ln p_t + \ln p_{t-2})^2\right]. \quad (5.19)$$

A comparison of the coefficients of $\ln \theta$ and $(\ln \theta)^2$ in the exponents in equations (5.18) and (5.19), gives the evolution rules for μ_t and σ_t^2 , and the following deterministic CBS is obtained:

$$\begin{aligned} (1+r)p_t &= \exp\left[\mu_{t-1} + \frac{1}{2}\sigma_{t-1}^2\right] p_{t-1} + \bar{y} \\ \frac{\mu_t}{\sigma_t^2} &= \alpha \frac{\mu_{t-1}}{\sigma_{t-1}^2} + (1-\alpha)(2\beta(\ln p_t - \ln p_{t-2}) + 1) \\ \frac{1}{\sigma_t^2} &= \frac{\alpha}{\sigma_{t-1}^2} + 2\beta(1-\alpha). \end{aligned} \quad (5.20)$$

Note that σ_t^2 does not interact with the other variables, and simply tends to its steady state value $1/(2\beta)$ at an exponential rate. Also note that the state-variables $\{\mu_t, \sigma_t^2\}$ no longer denote the mean and variance of the beliefs distribution. In the limit of an infinite number of agents, the average belief becomes

$$\int \vartheta \phi_{t-1}(\vartheta) d\vartheta = \exp\left[\mu_{t-1} + \frac{1}{2}\sigma_{t-1}^2\right],$$

while the degree of heterogeneity as measured by the variance, is given by:

$$\text{Var}_t[\theta] = \exp[2\mu_{t-1} + \sigma_{t-1}^2] (\exp[\sigma_{t-1}^2] - 1). \quad (5.21)$$

It can be verified that the fixed point solution of the system is given by:

$$(p^*, \mu^*, \sigma^{2*}) = \left(\frac{\bar{y}}{1 + r - \exp(\frac{3}{4\beta})}, \frac{1}{2\beta}, \frac{1}{2\beta} \right).$$

Note that for all finite positive values of β the fixed point price p^* is larger than the fundamental value \bar{y}/r . Only when β tends to infinity (i.e. the propensity to err tends to zero) the fixed point price tends to the fundamental price.

Remark: The examples discussed here and in the previous subsection are analytically tractable due to the quadratic terms in the exponents. A closed form analytic derivation of the dynamics might become cumbersome or even impossible if: (i) the beliefs distributions can not be represented within a finite parameter class, closed under updating according to equation (5.5) (it is closed under updating e.g. if $\pi_t(\theta)$ is a finite order polynomial in θ), or (ii) agent's decisions are based on more complicated optimization procedures such as those typical in a dynamic programming contexts. In those cases one might proceed using appropriate generalization of methods known in discrete choice simulation and estimation (see e.g. Keane and Wolpin, 1994). These generalizations are beyond the scope of this chapter and left for future research.

As noted above, we consider models that allow for an analytic derivation of the dynamics, although perhaps exceptional, important from a theoretical point of view, since they may help providing insights into the interaction between prices and the beliefs distribution. In the models derived above one can explicitly see how prices are formed given the beliefs distribution, and how realized prices in turn affect both the census and the dispersion of beliefs.

5.4.3 Local stability

The deterministic dynamical systems derived in the previous subsections, provide analytical descriptions of the co-evolution of both prices and the distribution of beliefs. Although our main motivation for considering these stylized examples is to illustrate how a CBS can be used to obtain insights into the interaction of prices and the beliefs distribution, these simple models also enable one to investigate how the dynamics is affected by the model

parameters. For example, how does the steady state and its stability depend on behavioral parameters such as the intensity of choice β and the degree of memory α ?

We present a brief bifurcation analysis of the deterministic skeletons just derived. However, it should be realized that the value of such an analysis strictly speaking is limited to the deterministic case only. Since we wish to illustrate the CBS methodology conceptually without putting too much stress on particular cases we choose to limit ourselves to a discussion on the local stability around the fixed points. The price equation, (5.10), suggests that increasing the interest rate has a stabilizing effect on the dynamics. Intuitively, also the memory parameter is important for stability, since it has a smoothing effect. The local bifurcation analysis is carried out mainly to check whether this intuition is justified.

The next proposition, which is proved in the Appendix, gives the local stability conditions:

Proposition 5.1 *For $\alpha \in [0, 1)$ and $\beta \in [0, \infty)$ and $r \in [0, \infty)$ the CBSes given in equations (5.15) and (5.20) are locally stable around the fixed point if $g(\alpha, a) > 0$, where*

$$g(\alpha, a) = 1 - 3a^2 - 2a\alpha + 5a^2\alpha + a\alpha^2 - 2a^2\alpha^2,$$

with $a = (1 + r)^{-1}$ for the squared prediction error case (equation 5.15) and $a = (1 + r)^{-1} \exp(\frac{3}{4\beta})$ for the squared logarithmic prediction error case (equation 5.20).

Figure 5.1 shows the line $g(\alpha, a) = 0$ where the bifurcation occurs. The dynamics is locally stable around the fixed point for parameter values below (right) of this bifurcation curve, and unstable for parameter values above (left) of the bifurcation curve. The memory parameter can be seen to have a stabilizing effect, provided that a is not too large. For any value of α in $[0, 1)$ the dynamics can be made locally stable by increasing the interest rate (decreasing a) sufficiently. In contrast with the squared prediction error case, with squared logarithmic prediction errors the intensity of choice, β , also affects stability, since stability is determined, apart from α , by $a = (1 + r)^{-1} \exp(\frac{3}{4\beta})$ rather than r only. In the limit where the memory parameter tends to one, the equilibrium will lose its stability once the interest rate falls below $r = \exp(\frac{3}{4\beta}) - 1$. Thus in the squared logarithmic prediction error case, the interest rate can be very low while sustaining stability, as long as the intensity of choice is correspondingly large.

It is possible to show that the points on the interior of the primary bifurcation curve (i.e. for $\alpha \in (0, 1)$) for both models correspond to Neimark-Sacker bifurcations. On the bifurcation curve at least one eigenvalue crosses the unit circle. Assuming that one eigenvalue




Figure 5.1: Bifurcation curve in the (α, a) -plane. The dynamics of the CBSes described by equations (5.15) and (5.20) are locally stable for parameter values below (right) of this curve, and unstable for parameter values above (left) of it. The parameter a is related to the model parameters through $a = (1 + r)^{-1}$ for the squared prediction error case (equation 5.15) and $a = (1 + r)^{-1} \exp(\frac{3}{4\beta})$ for the squared logarithmic prediction error case (equation 5.20).

crosses the unit circle at a real value (i.e. ± 1) leads to $\alpha = 0$ or $\alpha = 1$, i.e. the boundaries of the allowed values for α , since we required $\alpha \in [0, 1)$. It follows that on points in the interior of the bifurcation line the dynamics become unstable due to a complex conjugated pair of eigenvalues crossing the unit circle.

As a final remark on the bifurcations of the two models discussed above, it should be mentioned that numerical evidence indicates that squared prediction errors give rise to a supercritical bifurcation while squared logarithmic prediction errors lead to a subcritical bifurcation. Using squared prediction errors the stable fixed point breaks up into a quasiperiodic motion on a (topological) circle, which increases in size if the memory parameter is lowered, and eventually breaks up giving way to globally unstable dynamics. In the squared logarithmic prediction error case case, when one crosses from the stable to the unstable region in the parameter space, there is no stable quasiperiodic solution close to the fixed point. In fact, numerically we have only seen cross-overs from locally stable to globally unstable dynamics in the squared logarithmic prediction error case.

5.5 Natural sources of randomness

Traditionally, randomness in economic models is associated with exogenous shocks, for example due to news affecting a company's future earnings, or fluctuations of the interest rate. The CBS framework provides several possible natural sources of randomness, which enter via the market aggregate and can be associated with the uncertainty of choice.

In the previous section we have shown how the CBS approach can be used to derive deterministic dynamical systems in the presence of an infinity of agents. So far, two conceptual differences with the LTL approach (Brock *et al.*, 2005) became apparent: i) the CBS dynamics are expressed not only in terms of prices, but also in terms of variables which fully describe the distribution of beliefs among agents, and ii) the measure ν in the beliefs space need not be normalizable for a CBS. In fact, both examples in the previous section are based on non-normalizable measures ν , and hence cannot be obtained directly as an LTL.

In this section we exploit another important difference between the LTL and CBS approaches. Recall that the LTL approach starts with a continuum of agents, and examines the limiting dynamics as the number of possible strategies tends to infinity. In the CBS approach the starting point is a continuum of strategies. There the behavior of large markets can be studied by examining limits where the number of agents tends to infinity. Because dependence among agents' choices and heterogeneous market impact are not excluded *a priori*, because we avoided the assumption of a continuum of agents, large market limits under various conditions can be considered. As it turns out, for those scenarios the randomness associated with individual choices typically do not average out in the large market limit.

Several possible natural sources of randomness can be identified, which will be considered in the sequel. Firstly, a finite number of traders gives rise to stochasticity, because traders are assigned a belief $\theta_{i,t-1}$ at random from the beliefs distribution $\phi_{t-1}(\theta)$. Secondly, so far, we have implicitly assumed that no agents have a dominant market impact. This need be the case, for example, if wealth is not evenly distributed among agents. In that case the market impact of the wealthiest agents might still be important as the number of traders tends to infinity. Thirdly, dependence among agents can prevent the dynamics from becoming deterministic when the number of agents tends to infinity. If, for example, agents coordinate (at least partly) on a random variable, such as an exogenous variable, or on predictors announced by some 'leading' agents who quote their predictors publicly in an early stage, then the stochastic properties of that variable show up in the price dy-

namics. In the extreme case where all agents have identical idiosyncrasies, they would all use the same (random) strategy. Finally, for certain combinations of the utility function and the predictor function $p_{t+1}^e(\theta_{i,t-1})$, the law of large numbers may not apply because the mean of the aggregate need not exist. In those situations, the price dynamics might still be defined as a stochastic dynamical system, in the limit where the number of agents tends to infinity.

5.5.1 Finite number of agents

In general, for a finite number of agents, the CBS approach leads to a random dynamical system, the only exceptions being cases where all $p_{t+1}^e(\theta_{i,t-1})$ are identical with probability one. In general it might not be straightforward to derive the distributional properties of the aggregate. However, in case the mean and variance of the aggregate are finite, one can appeal to the central limit theorem which states that the distribution of the aggregate is asymptotically normally distributed. The following theorem is a direct application of the central limit theorem.

Theorem 5.2 (*Central Limit Theorem*) *Given \mathcal{F}_t , under Assumption 5.1, if n tends to infinity, the random variable*

$$\sqrt{n} \left(\frac{1}{n} \sum_{i=1}^n p_{t+1}^e(\theta_{i,t-1}) - E_t[p_{t+1}^e(\theta_{i,t-1})] \right)$$

converges in distribution to

$$W \sim N(0, \sigma_t^2)$$

with

$$\sigma_t^2 = \text{Var}_t[p_{t+k}^e(\theta_{i,t-1})]$$

if and only if $E_t[|p_{t+1}^e(\theta_{i,t-1})|] < \infty$ and $\text{Var}_t[p_{t+1}^e(\theta_{i,t-1})] < \infty$.

Example: Logarithmic prediction errors

In this example we consider the asset pricing example discussed in section 5.4.2, with a performance measure based on logarithmic squared prediction errors. We assume that agents make decisions regarding their strategy independently of the other agents (Assumption 1). The predictors $p_{t+1}^e(\theta)$ are linear in the parameter θ , and the price p_t follows from

equations (5.6) and (5.9):

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

The conditional variance of $\theta_{i,t-1}$, given \mathcal{F}_t , is equal to (see equation 5.21):

$$S_t^2 = \text{Var}_t[\theta_{i,t-1}] = e^{2\mu_{t-1} + \sigma_{t-1}^2} (e^{\sigma_{t-1}^2} - 1),$$

so that the conditional variance of $(1+r)p_t = \frac{1}{n} p_{t+1}^e(\theta_{i,t-1}) + \bar{y}$ equals $\frac{1}{n} S_t^2 p_{t-1}^2$. With the asymptotic normal approximation given by Theorem 2, the price equation for a finite number, n , of agents becomes

$$(1+r)p_t \simeq \left(\exp \left[\mu_{t-1} + \frac{1}{2} \sigma_{t-1}^2 \right] + \frac{S_t}{\sqrt{n}} \epsilon_t \right) p_{t-1} + \bar{y},$$

with $\epsilon_t \sim N(0, 1)$. If the number of agents, n , tends to infinity, the noise variance tends to zero and the price equation becomes deterministic. The full dynamics in the large number of agents tends to the ‘deterministic skeleton’ described by equation (5.15). Notice that the analogous asymptotic normal approximation is exact for the squared prediction error case, even for a finite number of agents n , since there the $\theta_{i,t}$ are normally distributed (see section 5.4.1). Also note the emergence of conditional heteroscedasticity. The time-varying variance $\frac{1}{n} S_t^2 p_{t-1}^2$ of the noise term is proportional to the evolving degree of heterogeneity, S_t^2 .

5.5.2 Market impact

Differences in market impact, which may be driven by wealth differences, among agents can lead to the persistence of randomness, even for infinitely many agents. For simplicity we only consider the effect of an unequal market impact without considering the endogenous evolution of wealth, since the latter is beyond the scope of this chapter. Endogenous wealth effects have for example been studied by Cabrales and Hoshi (1996) in a discrete choice framework with two types of agents.

Let us denote the market weight of agent i by w_i , normalized such that $\sum_i w_i = 1$. The

mean expected future price becomes

$$\bar{p}_{t+1}^e = \sum_{i=1}^n w_i p_{i,t+1}^e$$

which has a conditional mean equal to

$$E_t \left[\sum_{i=1}^n w_i p_{i,t+1}^e \right] = E_t [p_{i,t+1}^e]$$

and, assuming conditional independence of $\theta_{i,t-1}$ given \mathcal{F}_t , conditional variance

$$\text{Var}_t \left[\sum_{i=1}^n w_i p_{i,t+1}^e \right] = \left(\sum_{i=1}^n w_i^2 \right) \text{Var}_t [p_{i,t+1}^2].$$

The term between brackets on the right hand side of this equation, is known as the Herfindahl index of concentration. If all agents have equal weights, $w_i = \frac{1}{n}$, the conditional variance is proportional to one over n . For all other weight distributions the conditional variance is larger. A large concentration (large Herfindahl index) thus implies a small effective number of market participants and vice versa. The conditional variance suggests defining an effective number of agents, n_{eff} , say, in terms of the inverse of the Herfindahl index:

$$n_{\text{eff}} = \left(\sum_{i=1}^n w_i^2 \right)^{-1}.$$

The effective number of agents thus defined represents the number of agents that would give rise to the derived conditional variance in case their market impacts would be equal.

5.5.3 Interaction and dependence among agents

As a next source of randomness that can persist if the number of agents tends to infinity, we consider dependence among agents. Several types of dependence can be incorporated in the utility function by adding an interaction term, as in Brock and LeBaron (1996) and Brock and Durlauf (2001). In this approach individual utility depends not only on the past performance of strategies, but also on the choices made by other agents. The joint pdf of

the beliefs parameters of n agents is written as

$$\phi_t(\theta_1, \dots, \theta_n) \sim \exp \left(\sum_{i=1}^n \beta U_t(\theta_i) - \frac{J}{2} \sum_{i,j} h(\theta_i, \theta_j) \right),$$

where $u_t(\theta)$ denotes the performance measure, and $\frac{J}{2}h(\theta_i, \theta_j)$ is a function which captures the additional utility derived from interactions. The degree of dependence is measured by $J > 0$. Although formally the θ_i depend on time, the subscript is dropped here for simplicity.

Following Brock and LeBaron (1996) we consider a utility function which depends on squared prediction errors, and an interaction term that specifies that each agent prefers parameter values close to the consensus. One obtains

$$\phi_t(\theta_1, \dots, \theta_n) \sim \exp \left(- \sum_{i=1}^n \frac{(\theta_i - \mu_t)^2}{2\sigma_t^2} - \frac{J}{2} \sum_{i=1}^n (\theta_i - \bar{\theta})^2 \right), \quad (5.22)$$

where $\bar{\theta} = \frac{1}{n} \sum_{i=1}^n \theta_i$. Equation (5.22) contains linear and quadratic terms in the parameters θ_i and hence represents a multivariate gaussian pdf. Upon comparing this with that of a multivariate normal

$$\frac{1}{(2\pi)^{n/2} \|\Sigma\|^{1/2}} \exp \left(-\frac{1}{2} (\boldsymbol{\theta} - \boldsymbol{\theta}_0)^T \Sigma^{-1} (\boldsymbol{\theta} - \boldsymbol{\theta}_0) \right),$$

where Σ is the variance covariance matrix of the θ_i , the statistical properties of the market average $\bar{\theta} = \frac{1}{n} \sum_{i=1}^n \theta_i$ can be derived. As in the case without interaction the expected value of the market average is given by $E_t[\bar{\theta}] = \theta_0 = \mu_t$. Due to permutation symmetry of the agents all diagonal elements of Σ are identical, as well as the off-diagonal elements. The same holds for Σ^{-1} . This implies that both Σ and Σ^{-1} have $\mathbf{1}$ as an eigenvector. We denote the corresponding eigenvalue of Σ by λ , so that the corresponding eigenvalue of Σ^{-1} is $1/\lambda$. For the elements of Σ^{-1} we find

$$(\Sigma^{-1})_{i,i} = \frac{1}{\sigma_t^2} + 2 \frac{n-1}{n} J,$$

and

$$(\Sigma^{-1})_{i,j} = \frac{J}{n} \quad \text{for } i \neq j.$$

The eigenvalue λ can now be determined from:

$$\Sigma^{-1}\mathbf{1} = \frac{1}{\lambda}\mathbf{1} = \left(\frac{1}{\sigma_t^2} + \frac{(n-1)J}{n} \right) \mathbf{1},$$

giving $\lambda = \frac{1}{\sigma_t^2} + \frac{(n-1)J}{n}$. For the variance of $\bar{\theta}$ one then finds

$$\text{Var}[\bar{\theta}] = \frac{1}{n^2} \text{Var} \left[\sum_{i=1}^n \theta_i \right] = \frac{1}{n^2} \mathbf{1}^T \Sigma \mathbf{1} = \frac{\lambda}{n} = \frac{1}{\frac{n}{\sigma_t^2} + (n-1)J},$$

which indicates that the interaction reduces the variance of the market average $\bar{\theta}$. Thus, when agents coordinate on the population mean, meaning that no agent prefers to be far away from the consensus, the variance of the population mean $\bar{\theta}$ reduces. Notice that the mean market expectation $E_t[\bar{\theta}] = \mu_t$ is identical with or without interaction.

The effective number of agents can be obtained by relating the conditional variance of the aggregate expectation $\bar{\theta}$ to the conditional variance in the absence of interaction. Upon defining n_{eff} through $\sigma_t^2/n_{\text{eff}} = \text{Var}[\bar{\theta}]$, one directly obtains

$$n_{\text{eff}} = n + (n-1)J\sigma_t^2.$$

In the example just considered the interaction resulted in a reduction of the variance of the mean expectation $\bar{\theta}$. An increase in the variance of the market average $\bar{\theta}$ can also occur. A different and perhaps more realistic scenario might be one in which agents coordinate partly on some source of information, being, for example a public exogenous variable. Such a source of information might, but need not necessarily, be related to economic fundamentals. In case the ‘signal’ on which agents coordinate is an exogenous noise source, this can increase the correlation of the agent’s predictors, and hence lead to an increase of the variance of the market average. This scenario can be represented by the following joint pdf:

$$\phi_t(\theta_1, \dots, \theta_n) \sim \exp \left(- \sum_{i=1}^n \frac{(\theta_i - \mu_t)^2}{2\sigma_t^2} - \frac{J}{2} \sum_{i=1}^n (\theta_i - X_t)^2 \right),$$

where X_t represents any random variable on which the agents coordinate. In case X_t is taken to be equal to the parameter θ_i of one of the agents, agent 1, say, so that $X_t = \theta_1$, this covers the situation where one agent announces his strategy publically at an early

stage, after which others coordinate on this (noisy) signal. Clearly, the randomness that enters the joint pdf has a similar effect on all individual choices. If X_t is positive (negative), the expected value of $\bar{\theta}$ increases (decreases). Moreover, this effect does not average out when the number of agents is increased. Coordination on a (possibly endogenous) random variable thus provides an additional source of randomness in the price dynamics.

5.5.4 Inherent randomness

As a final source of randomness we consider an example of the dynamics in a case where $p_{t+1}^e(\theta_{i,t-1})$ does not have a finite conditional mean, as a result of which the law of large numbers does not apply. The aggregate expectation $\frac{1}{n} \sum_{i=1}^n p_{t+k}^e(\theta_{i,t-1})$ does not tend to a constant, but tends in distribution to a well defined random variable when n tends to infinity. We consider an asset pricing model with agents choosing among constant predictors: $p_{t+1}^e(\theta) = \theta$. For simplicity we put $\alpha = 0$ (no memory, only the last observed price is taken into account). A situation in which the mean does not exist arises if for the performance measure one takes

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

The beliefs distribution can now become fat-tailed. For the beliefs distribution one obtains

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

from which it follows that

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

a t -distribution with $2\beta - 1$ degrees of freedom. Thus, the $\theta_{i,t-1}$ are distributed symmetrically around p_{t-1} , and for $\beta < \frac{3}{2}$ the mean does not exist.

For $\beta = 1$, $\theta_{i,t} - p_t$ given \mathcal{F}_t is Cauchy(0,1) distributed. Since this distribution is closed under averaging, this gives $\bar{p}_{n,t+1}^e = \frac{1}{n} \sum_{i=1}^n \theta_{i,t-1} - p_{t-1} \sim \text{Cauchy}(0,1)$. 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 resulting price process is a first order autoregressive process 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.

We conclude that the remark made by Muth (1961) that “allowing for cross-sectional differences in expectations is a simple matter because their aggregate effect is negligible as long as the forecasts are not strongly correlated” only applies if the distribution of predictors is sufficiently well-behaved.

Notice that the non-vanishing randomness of the type described here can only occur with a continuum of strategies. If the number of strategies is finite, provided that each of them predicts a finite price, the mean expectation always exists, in which case, by the law of large numbers, the dynamics become deterministic when the number of traders tends to infinity.

5.6 Concluding remarks

Many will agree that most economic observables, such as realized prices, are to a large extent determined by the expectations of economic agents. In turn, individual expectations are typically shaped by the observed past, as agents will try to incorporate the underlying economic laws at force in their predictions for tomorrow. New observations provide agents with an incentive to update their beliefs whenever their belief appears inconsistent with these observations. One would therefore expect individual beliefs to co-evolve over time with realized prices. Being evidently important but not directly observable, the evolution of the distribution of beliefs among market participants is at the heart of our concept.

To set up a framework for studying the evolution of beliefs distributions, we defined probability density functions on a space of possible strategies, which is called the beliefs space. The continuous choice model is employed to update the beliefs distribution, while the incentives for the agents to switch are provided by the past performances of strategies. Co-evolution of a distribution of beliefs with observed prices thus emerges due to the ongoing evaluation of predictors. This approach to modeling the evolution of the beliefs distribution over time provides several insights into the nature of its feedback with economic observables. The results and implications of our concept are at least two-fold.

Firstly, while being applicable in a wide range of theoretical models where expectations feedback plays a key role, the concept allows one to model the joint dynamics of the beliefs distribution and observables (prices, say) explicitly in terms of a *random dynamical system*. The beliefs distribution becomes a state variable which is endogenously shaped by the assumptions about the functional form of predictors, and the type of performance

measure (squared prediction error or profits etc.) employed. For illustrative purposes, we have included some stylized examples that lead to a normal- and a log-normal distribution for individual beliefs. In these cases, the distribution of beliefs is fully described by the average belief and the dispersion of beliefs, the latter of which can be related to the degree of heterogeneity. The CBS then prescribes explicitly how the dispersion of beliefs affects prices, and in turn how prices affect the dispersion of beliefs.

Secondly, our concept provides what we refer to as ‘natural sources of randomness’. In many economic models, noise, which is required when matching stylized facts is the objective, is often included ad-hoc as additive exogenous shocks. Popular justifications include: exogenous news shocks, model approximation error, and noise induced by trading. Our natural source of randomness introduces a type of endogenous uncertainty that can not be associated with exogenous shocks but rather with the uncertainty of choice deriving from unpredictable aspects of the preferences of economic agents. Moreover, the endogenous noise term is shown to inherit its statistical properties directly from the beliefs distribution. To understand where the stochasticity derives from, note that in a CBS the beliefs distribution describes the likelihood, through the eyes of the econometrician, according to which individual agents select their beliefs. When the number of agents is finite, the aggregate belief by definition is a random variable. From simulation studies (e.g. Lux and Marchesi, 1999) it is known that the stochastic nature of the dynamics in simulated markets vanishes when the number of agents becomes large. This is typically what the CBS framework would predict if agents have comparable market weights and their predictions are well-behaved (have finite mean and variance) and are not strongly dependent across agents. However, the CBS approach shows that there are several scenarios for which endogenous noise remains significant in a large market. This can occur if, for example: (i) wealth is disproportionately distributed (ii) agents’ choices exhibit dependence; and (iii) when the combination of the performance measure and the functional form of beliefs induces nonexistence of the mean (with respect to the beliefs distribution) of the predictor function.

A number of extensions and generalizations might prove worthwhile. For example connecting the concepts of a CBS with that of “rational beliefs”, the latter being recently developed by Kurz (2001). In the work of Kurz, a beliefs distribution describes the individual belief of an agent, and reflects the uncertainty individuals hold regarding future economic variables. This concept is designed to generalize rational expectations theory where all agents have exact knowledge about the future. Although Kurz’s approach is

different from ours, in particular since the beliefs distribution represents a different object, Kurz's definition of 'Endogenous Uncertainty' as 'that component of the volatility of quantities and prices in the economy which is generated by the distributions of beliefs' has similar implications. Both approaches are able to relate their notion of endogenous uncertainty to common financial market "anomalies" such as excess volatility, and ARCH-type structure. When comparing our approach with that by Kurz, a logical extension of the CBS would be to endow agents with notions of uncertainty about future values rather than having point predictors. Then each agent's belief is described by a pdf, so that agents will generally have more dimensions to disagree upon.

One area for future research is the analysis of the random dynamical systems obtained within the CBS framework, using stochastic bifurcation theory. Stochastic bifurcation theory is closely connected to the theory of random dynamical systems, a new field which currently is under rapid development (see e.g. Arnold (1998)). In stochastic bifurcation theory, progress has been relatively fast on continuous time random dynamical systems when compared to discrete time random dynamical systems. The bifurcation theory of discrete time random dynamical systems, which is still in its infancy, would clearly be useful for characterizing the dynamical systems of the type considered here.

5.7 Appendix

Below a proof of Proposition 5.1 is given. The proof uses the fact that the characteristic equations for the local dynamics of the CBS with squared prediction errors (section 5.4.1) and squared logarithmic prediction errors (section 5.4.2) are closely connected. In fact it will be shown that the stability conditions are identical in terms of a if one defines

$$a = (1 + r)^{-1}$$

for the squared prediction error case, and

$$a = (1 + r)^{-1} e^{\frac{3}{4\beta}}$$

for the squared logarithmic prediction error case.

Proof of Proposition 5.1

For the squared prediction error case (section 5.4.1), upon defining $w_t = \frac{1}{\sigma_t^2}$, the dynamics can be casted in the form

$$\begin{aligned} p_t &= (\mu_{t-1} p_{t-1} + \bar{y}) / (1 + r) \\ q_t &= p_{t-1} \\ \mu_t &= (\alpha \mu_{t-1} + 2\beta(1 - \alpha) p_t q_{t-1}) \frac{w_{t-1}}{w_t} \\ w_t &= \alpha w_{t-1} + 2\beta(1 - \alpha) q_{t-1}^2, \end{aligned}$$

where p_t and w_t on the right hand side of the equation for μ_t are functions of the state variables at time $t - 1$, as specified by the first and last equation. The Jacobian matrix evaluated at the equilibrium evaluates to

$$J_{\text{SPE}} = \begin{pmatrix} a & 0 & ap^* & 0 \\ 1 & 0 & 0 & 0 \\ (1 - \alpha) \frac{a}{p^*} & -\frac{1 - \alpha}{p^*} & \alpha + (1 - \alpha)a & 0 \\ 0 & 4\beta p^* & 0 & \alpha \end{pmatrix},$$

with $a = (1 + r)^{-1}$.

Similarly, the CBS with squared logarithmic prediction errors (section 5.4.2) can be

expressed as

$$\begin{aligned} p_t &= \left(\exp \left(\mu_{t-1} + \frac{1}{2w_{t-1}} \right) p_{t-1} + \bar{y} \right) / (1+r) \\ q_t &= p_{t-1} \\ \mu_t &= (\alpha \mu_{t-1} w_{t-1} + (1-\alpha)(2\beta(\ln p_t - \ln q_{t-1}) + 1)) / w_t \\ w_t &= \alpha w_{t-1} + 2\beta(1-\alpha), \end{aligned}$$

the Jacobian at the equilibrium becomes

$$J_{\text{SLPE}} = \begin{pmatrix} a & 0 & ap^* & -\frac{ap^*}{8\beta^2} \\ 1 & 0 & 0 & -\frac{(1-\alpha)a}{8\beta^2} \\ (1-\alpha)\frac{a}{p^*} & -\frac{1-\alpha}{p^*} & \alpha + (1-\alpha)a & 0 \\ 0 & 0 & 0 & \alpha \end{pmatrix},$$

where $a = \frac{e^{\frac{3}{4\beta}}}{1+r}$.

Although the two Jacobian matrices differ, their characteristic equations coincide due to the fact that the upper left 3×3 sub-matrices are identical, and because of the three off-diagonal zeros in the last column/row respectively. In both cases the corresponding characteristic equation can be written as

$$\det [\lambda I - J] = [\lambda^3 - (2a + \alpha(1-a))\lambda^2 + \alpha\lambda a + (1-\alpha)a] (\lambda - \alpha) = 0.$$

Since α is an eigenvalue, a necessary condition for stability is $|\alpha| < 1$. This condition is automatically satisfied under our assumption that $\alpha \in [0, 1)$. It can be readily verified that $(0, 0, 0, 1)'$ is the corresponding eigenvalue in both cases, so that the stability question reduces to that of the upper left 3×3 sub-matrix

$$\tilde{J} = \begin{pmatrix} a & 0 & ap^* \\ 1 & 0 & 0 \\ (1-\alpha)\frac{a}{p^*} & -\frac{1-\alpha}{p^*} & \alpha + (1-\alpha)a \end{pmatrix}.$$

Application of the conditions for stability derived in Jury (1974) to the characteristic equation

$$\det [\lambda I - \tilde{J}] = \lambda^3 - (2a + \alpha(1-a))\lambda^2 + \alpha\lambda a + (1-\alpha)a = 0$$

of the remaining eigenvalues, leads to several conditions on the parameters which, for

$\alpha \in [0, 1)$, can be summarized as

$$1 - 3a^2 - 2a\alpha + 5a^2\alpha + a\alpha^2 - 2a^2\alpha^2 > 0.$$

□