Imperial College London

MEng Final Year Project

Investigating General Equilibria in Artificial Stock Markets

Author:
Daniel Cooke

Supervisor:
Dr. Daniel Kuhn

June 21, 2012

Abstract

We develop a new class of artificial stock market that apply general equilibrium price clearing. Our model is the first multi-asset artificial stock market to use exogenously generated dividends from a geometric vector autoregressive model and is therefore the first truly multi-asset artificial stock market. Our model does not require the concept of a risk free asset and agents are able to form demands for an arbitrary number of assets with arbitrary covariance structure. We derive an optimisation problem for computing approximate equilibrium prices that clear our market and then show our market is able to faithfully reproduce many real stock market dynamics and phenomenon.

Keywords: Applied General Equilibrium; Agent-based Artificial stock-markets; Multi-Asset Market; Market Dynamics.

Acknowledgements

I would like to thank my supervisor, Dr. Daniel Kuhn for kindly agreeing to supervise my project proposal and providing me with support and advice throughout the course of my project. Further, I am grateful to my second marker, Prof. Abbas Edalat for reviewing my draft report. My gratitude goes to CSG for helping with resource requirements, and also to Dr. Anil Seth (University of Sussex) for providing his Granger Causality toolkit. Finally I would like to thank my family and friends for providing much needed support and encouragement throughout my time at Imperial.

Contents

1

Introduction

Economies are complex dynamical systems. Large numbers of self-interested, heterogenous micro-agents compete for limited resources through decentralised global markets. In a system of such apparent chaos, it is remarkable that markets display coherence and order; responding to collective supply and demand fluctuations with such efficiency that many believe price movements are random and cannot be forecasted. Economists have grappled with modelling economic systems for hundreds of years and although the use of financial modelling is questioned by some, [1] there is no doubt it is greatly influential in the policy decisions of many governments, hedge funds and investment banks.[2]

Since the time of Walras, economic modelling has centred around the paradigm of supply and demand equilibrium. The Walrasian economy requires prices be adjusted by a centralised market auctioneer. Utility maximising producers and consumers form supply and demand decisions on scarce resources conditional on the market prices determined by the auctioneer; trading may only take place once an equilibrium price is found. *General Equilibrium Theory* is a branch of microeconomics concerned with exploring the dynamics of equilibria within the Walrasian framework.

Dynamics of equilibria, such has existence and uniqueness, are now well established theory. However, these results only apply to a class of models which conform to a strict set of conditions, most notably that market participants are concave utility maximisers, possess zero degree homogenous demand functions and adhere to rational expectations. Unexpected financial events such as the Black Monday market crash (1987), and the more recent financial crisis have caused many to question the validity of these assumptions. As a result, a new breed of heterogeneous, boundedly-rational agent-based market models are becoming an area of increasing research and development. These simulated artificial stock markets do not in general satisfy equilibria conditions and thus focus is turned to the dynamics of agents operating within the market. Many, such as Gupta et al. (2005), believe artificial markets may be used for time-series forecasting.

Most current research on artificial markets has focused on a single risky asset markets, populated with mean-variance criterion myopic investors. Furthermore, while some artificial markets use equilibrium theory for price clearing algorithms, these are not of the same vein as those established in general equilibrium theory, and there is certainly no requirement for their use, indeed, other pricing algorithms that do not enforce market clearing, such as order-book style matching are also widely implemented. [3]

[1] See for example Taleb (2008)

[2] See for example Holly and Weale (2000, chap. 4) and Pindyck and Rubinfeld (1998)

[3] A survey of artificial markets is given by LeBaron (2006)

While we agree with new bounded market initiative, we do not feel the current class of artificial markets capture the dynamics of real markets in a satisfactory way. Firstly, we do not believe it is possible to fully classify an asset as 'risk-less'; an assumption often made in the current line of artificial markets. Secondly, we feel the abstraction to a single risky asset economy removes some of the most vital competitive characteristics of market dynamics. Both of these observations in our view lead to the same conclusion: the need for a global price adjustment mechanism. For if there is no truly risk-less asset then by definition, for a market to prevail there must exist at-least two unique risky assets. This naturally invokes the idea of general equilibria. Thus, we conclude that removing the assumption of a risk-less asset has dramatic implications for the structure and complexity of fully clearing artificial market models, in that, a general equilibrium schema becomes vital for price formation. To the best of our knowledge, the feasibility of a multi-asset artificial market operating under general equilibrium pricing has not been studied.[4]. Nevertheless, somewhat ironically, we feel it is a natural evolutionary direction for artificial markets to take.

Finding equilibria in a seemingly chaotic market is a challenging mathematical and computation task, and although a small selection algorithms now exist for computing general equilibria, almost all have conditional convergence on general equilibrium theory existence and uniqueness assumptions; there is very little understanding on the applicability of general equilibrium price finding algorithms outside the Walrasian domain. Moreover, many of these algorithms require additional convergence assumptions such as gross substitutes. As described by Noam Nisan and Vazirani (2007a), in some cases it is possible to compute equilibria via convex programming, however, even if the problem domain is convex, many convex solvers implment quasi-Newton methods that require Jacobian and Hessian information. More recently, a new class of tâtonnement style algorithms have been developed that do not require differential information, however, these methods may not converge when the problem domain is non-convex.

Artificial stock market research has focused heavily on recreating real market dynamics such as high volatility, fat-tailed return distributions and high trading volume. These dynamics are absent from standard general equilibrium models as agents are homogenous and non-evolutionary. Our proposal is to combine artificial stock market and general equilibrium models and show this new class of hybrid model is flexible, robust, and is able to produce a wealth of empirically seen market phenomenon.

1.1 Objectives

Our aim is to develop a new class of artificial stock markets that are not restricted to a single risky asset. We hope our new market model will be able to replace, or at least extend the current generation of artificial stock markets and become a valuable research tool for studying economic stock markets, and analysing policy decisions.

We will use methods similar to those employed in computable general equilibrium theory for computing general equilibrium price vectors that clear our market, in particular we will achieve this by deriving an optimisation problem such that optimal solutions characterise approximate equilibrium prices. Following this we will then attempt to vindicate our model by analysing the dynamics of equilibria with relation to general equilibrium theory, and in particular the existence and stability of equilibrium in our market.

Finally, we hope to show our artificial stock market is able to faithfully reproduce the dynamics displayed by real markets. We will do this by empirically testing simulated time-series generated by our market. Additionally, we will analyse if our market is efficient with respect to the efficient market hypothesis.

[4] Another multi-asset artificial stock market has been developed by Westerhoff (2004). However, his model does not enforce market clearing, nor doe it contain exogenously determined assets and thus does not fall within the same class of model we are proposing.

1.2 Report Structure and Contributions

- **Background**: We describe much of the theory required to fully understand the economic setting of our thesis. We show many of key results that will required in the development of our artificial stock market model. [Chapter 2]

- **Artificial Stock Market Model** We develop our artificial stock market model. We begin by analysing real market dividends and develop a model for accurately generating dividend time-series. Following this we derive our agents portfolio choice, demand function, expectations forecasting and learning mechanism. We also discuss the timing of our market and some important implementation concerns. Finally we show a method for calibrating our dividends model. [Chapter 3]

- **Computing General Equilibrium Prices**: We begin by determining if our model satisfies the conditions required for general equilibrium existence. Following this we analyse if existing algorithms for computing general equilibrium prices can be applied in our model, and then develop an optimisation problem for finding equilibrium prices. We also discuss algorithms for solving the optimisation problem and analyse the performance of each method. [Chapter 4]

- **On Equilibria in Artificial Stock Markets**: In this chapter we complete the first phase of our evaluation by examining the dynamics of equilibrium in our market, we analyse if equilibrium prices exist and if so, whether they can be reliably computed. We finish with a brief discussion on other equilibria properties. [Chapter 5]

- **Market Dynamics**: In this chapter we complete the second phase of our evaluation; we asses the time-series generated by our artificial stock market and compare them to real market data we have collected. In particular we analyse if our market is able to to generate volatile prices; if returns distributions have kurtosis; and if our market has high trading volume. Furthermore we briefly examine wealth dynamics in our market. Finally we test if our market is efficient by applying cross-covariance statistics and Granger causality testing. [Chapter 6]

- **Conclusions & Future Work**: We conclude our thesis by drawing together our main finding and presenting a selection of suggested lines of future research. [Chapter 7]

- **Appendix**: In the Appendix we briefly outline some of our key implementation details. We also detail many of the mathematical background material required to understand our model and results, and in particular we show how to construct and analyse Vector Autoregressive models. Furthermore we briefly examine mathematical optimisation. [Chapter 8]

2
Background

In this chapter we review some key literature in economics theory that readers should have some familiarity with before reading the remainder of this thesis. Chronological order is respected where possible. Some references to additional material are provided for readers interested in pursuing related work. Note some mathematical definitions are deferred to the appendix.

2.1 Modern Portfolio Theory

Modern portfolio theory is an area of investment science that formulates statistically optimal portfolio decisions. Investors place money into assets (stocks, bonds etc), in the hope that they will obtain a payoff greater than their initial investment at a later date; "a later date" is assumed to be a single period, this is well defined for assets such as bonds that have fixed maturity dates, but not for assets such as stocks which can theoretically be held forever. For assets such as these we can view a period as an investor's *investment horizon*, that is, the time in the future the investor wishes to recover their funds, for example, a *myopic* investor may only have an investment horizon of a few weeks or months. All assets are inherently risky[1], there is no guarantee of a positive return even for assets deemed 'safe' such as government bonds.[2] Furthermore, investors are assumed to be *risk averse*; they require higher expected returns when taking on additional risk. It is also assumed *no-arbitrage* conditions hold; no completely risk-free profits can be made.

Definition. The *rate of return* on an asset is:

$$r = \frac{X_1 - X_0}{X_0} \tag{2.1.1}$$

where X_1 and X_0 denotes realised funds received and invested respectively.

Definition. The *risk* of an asset is measured by the standard deviation, σ, of the assets return.

We can visualise this definition of risk by looking at Figure 2.1, **HNZ** has a low standard deviation while **NVDA** has a high standard deviation. One would usually describe **HNZ** as the 'safer', of the two assets. [3]

[1] The definition of "risk" differs depending on the setting. For example, in a mean-variance setting risk is measured by standard deviation (σ), where as in the Capital Asset Pricing Model risk is measured by beta (β). It is not necessary to understand the reasoning behind these differences as we will only require σ to quantify risk, however, the reader may refer to Luenberger (1998)

Figure 2.1: Time-series comparison of H.J. Heinz Company common stock (HNZ) and NVIDIA Corporation common stock (NVDA) over a two year period. [Obtained from `http://finance.yahoo.com`]

2.1.1 Mean-Variance Criterion

It is unlikely that a wise investor would hold a single asset, indeed, it can be shown that it is statistically beneficial to hold a collection, or *portfolio* of assets, often called *diversification*. Under the mean-variance criterion we only consider portfolios of assets.

Definition. The *rate of return* on a portfolio of is:

$$r_P = \sum_{i=1}^{N} w_i r_i = \boldsymbol{w} \cdot \boldsymbol{r} \tag{2.1.2}$$

where w_i and r_i is the weight and return of asset i respectively and $\boldsymbol{w}$ & $\boldsymbol{r}$ are the corresponding vectors.[4] Note we must have

$$\sum_{i=1}^{N} w_i = \boldsymbol{w} \cdot \boldsymbol{\iota} = 1$$

where $\boldsymbol{\iota}$ is the vector of ones.[5]

for further details.

[2] Governments have defaulted on bonds in the past, most notably the Russian government in 1998.

[3] For readers wishing to explore reasons why one company may be riskier than another, we recommend Graham and Zweig (2003), which must be regarded as one of the great investment texts.

[4] We use boldface to represent vectors throughout our thesis.

[5] We use notation throughout our thesis.

Definition. The expected return, or mean of a portfolio is

$$\mathbb{E}[r_P] = \bar{r}_P = \mathbb{E}\left[\sum_{i=1}^{N} w_i r_i\right] = \sum_{i=1}^{N} w_i \mathbb{E}[r_i] = \sum_{i=1}^{N} w_i \bar{r}_i = \boldsymbol{w} \cdot \bar{\boldsymbol{r}} \tag{2.1.3}$$

Definition. The variance of a portfolio is

$$\sigma_P^2 = \boldsymbol{w} \cdot (\Sigma \boldsymbol{w}) \tag{2.1.4}$$

where Σ is the covariance matrix of asset returns.

An investor that is risk averse prefers a higher return for increased risk. An investor's taste-for-risk can be quantified with a risk-aversion parameter, α, as α increases the investor is less willing to take risk.

Definition. A mean-variance maximising investor solves

$$\begin{aligned} \underset{w}{\text{maximise}} \quad & \boldsymbol{w} \cdot \bar{\boldsymbol{r}} - \frac{\alpha}{2} \boldsymbol{w} \cdot (\Sigma \boldsymbol{w}) \\ \text{subject to} \quad & \boldsymbol{w} \cdot \boldsymbol{\iota} = 1 \\ & \boldsymbol{w} \geq 0 \end{aligned} \tag{2.1.5}$$

Equivalently, we can infer for a given return an investor would prefer less risk, this leads to the famous Markowitz model

$$\begin{aligned} \underset{w}{\text{minimise}} \quad & \frac{1}{2} \boldsymbol{w} \cdot (\Sigma \boldsymbol{w}) \\ \text{subject to} \quad & \boldsymbol{w} \cdot \bar{\boldsymbol{r}} = \bar{r}_P \\ & \boldsymbol{w} \cdot \boldsymbol{\iota} = 1 \\ & \boldsymbol{w} \geq 0 \end{aligned} \tag{2.1.6}$$

The solution to these problems is a vector, $\boldsymbol{w}$, that characterise the investor's optimal portfolio. We can plot many of these optimal portfolios by solving 2.1.5 for varying α (or $\bar{r}_P$), the resulting solution set, depicted in 2.2 is called the *efficient frontier*.

Without constraints on portfolio weights it is possible an optimal portfolio to contain negative weights, this is called *short-selling*. When an investor sells-short they borrow an asset which they then sell, at a later date the asset must be returned to the original owner (note the original owner also receives any dividend paid during borrowing period). Thus a positive rate of return is made on a short-sale only when the price of the asset decreases over the investment horizon. Many investors are prohibited from short-selling as it is considered very risky for both parties. The effect on the efficient frontier is shown in Figure 2.2 by the red shaded region. Optimisation problems such as 2.1.6 are called *Quadratic Programs*[6] (QP) and can not, in general, be solved analytically. However, there exists a good selection of algorithms for efficiently computing globally optimal solutions for QP's. Details of mathematical optimisation are deferred to the Appendix 8.5.

If we remove the portfolio weight constraint $\boldsymbol{w} \geq 0$ we are left with an unconstrained convex optimisation problem. These types of problem can be solved analytically by solving the KKT optimality conditions which we now demonstrate.

[6] A quadratic program is a special type of optimisation problem, discussed in the Appendix (8.5).

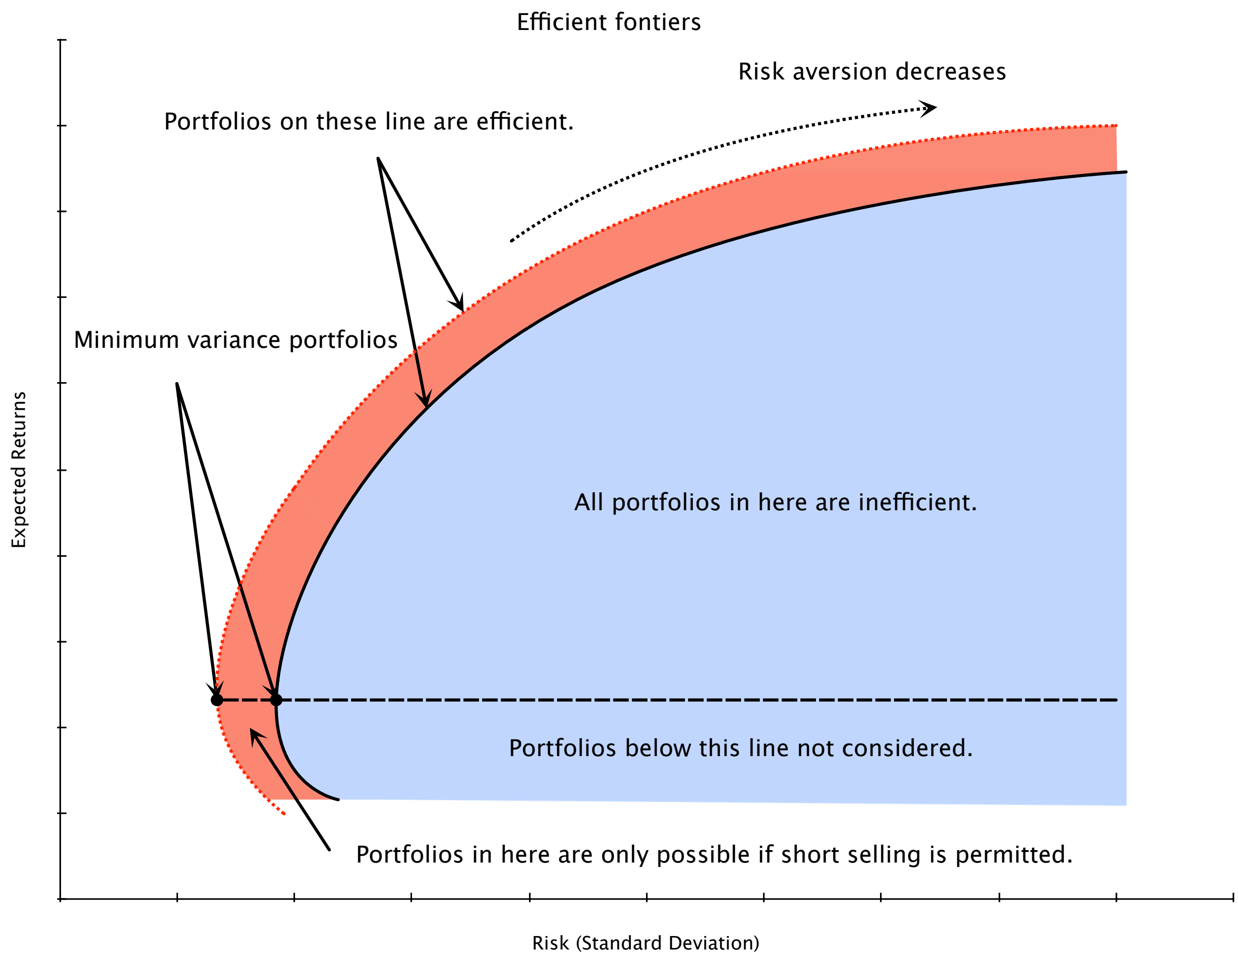

Figure 2.2: Depiction of efficient frontiers with and without short selling.

Example (Mean-Variance solution)**.** By forming the Lagrangian

$$L(\boldsymbol{w}, \lambda) = \boldsymbol{w} \cdot \bar{\boldsymbol{r}} - \frac{\alpha}{2} \boldsymbol{w} \cdot (\Sigma \boldsymbol{w}) - \lambda (\boldsymbol{w} \cdot \boldsymbol{\iota} - 1) \tag{2.1.7}$$

we find the first-order optimality (KKT) conditions by partially differentiating the Lagrangian

$$\begin{aligned} \frac{\partial L(\boldsymbol{w}, \lambda)}{\partial \boldsymbol{w}} &= \bar{\boldsymbol{r}} - \alpha \Sigma \boldsymbol{w} - \lambda \boldsymbol{\iota} = 0 \\ \frac{\partial L(\boldsymbol{w}, \lambda)}{\partial \lambda} &= \boldsymbol{w} \cdot \boldsymbol{\iota} - 1 = 0 \end{aligned} \tag{2.1.8}$$

Therefore

$$\boldsymbol{w}^* = \frac{1}{\alpha} \Sigma^{-1} (\bar{\boldsymbol{r}} - \lambda \boldsymbol{\iota}) \tag{2.1.9}$$

where $\boldsymbol{w}^*$ denotes the optimal portfolio. Noting we have $N + 1$ equations for N unknowns we can can substitute $\boldsymbol{w}^*$ into the second optimality condition

$$\begin{aligned}1 &= \boldsymbol{w} \cdot \boldsymbol{\iota} \\ &= \frac{1}{\alpha}\Sigma^{-1}(\bar{\boldsymbol{r}} - \lambda\boldsymbol{\iota}) \cdot \boldsymbol{\iota} \\ &= \frac{1}{\alpha}\boldsymbol{\iota} \cdot (\Sigma^{-1}\bar{\boldsymbol{r}}) - \left(\frac{1}{\alpha}\boldsymbol{\iota} \cdot (\Sigma^{-1}\boldsymbol{\iota})\right)\lambda\end{aligned} \tag{2.1.10}$$

Therefore

$$\lambda = \frac{\boldsymbol{\iota} \cdot (\Sigma^{-1}\bar{\boldsymbol{r}}) - \alpha}{\boldsymbol{\iota} \cdot (\Sigma^{-1}\boldsymbol{\iota})} \tag{2.1.11}$$

Substituting λ back into 2.1.9 gives the optimal portfolio

$$\boldsymbol{w}^* = \frac{1}{\alpha}\Sigma^{-1}\left(\bar{\boldsymbol{r}} - \frac{\boldsymbol{\iota} \cdot (\Sigma^{-1}\bar{\boldsymbol{r}}) - \alpha}{\boldsymbol{\iota} \cdot (\Sigma^{-1}\boldsymbol{\iota})}\boldsymbol{\iota}\right) \tag{2.1.12}$$

Aside. *Readers familiar with portfolio theory may wonder why we have not considered the concept of a 'risk-free' asset. This is because, as stated in Section 1, we do not consider there to exist a truly riskless asset, thus they do not form part of our model. The impact of a risk-free asset on portfolio theory is significant, and not within the scope of this thesis, we refer those wishing to pursue these concepts to Luenberger (1998).*

2.1.2 Utility Portfolio Choice

We now consider a more abstract notion of portfolio theory which is closely related (and sometimes equivalent) to the mean-variance criterion. It was shown in the previous section that an investor's portfolio decision was dependent on their *risk-aversion*, the following theory formalises this idea.

Definition. A utility function $U(x)$, where $U : \mathbb{R}^n \to \mathbb{R}$, maps a measurable quantity x, such as a vector of real numbers, to a level of satisfaction.

Utility functions are a method of ranking objects, such as good & services, investments, or wealth. Objects with higher utility are more preferable. Some common utility functions are listed in Table 2.1, and shown in Figure 2.2.

Expected Utility

When presented with uncertainty, utility functions represent investor risk preferences. An investor evaluates an investment opportunity by calculating *expected* utility. Investors using *concave* utilities are risk-averse, as the following example demonstrates.

Example. Suppose we decide to use power utility with $\alpha = 0.5$ and are offered an option of either £1 for sure, or £2 if a coin flip lands on heads and nothing if it lands on tails. To decide which option is most preferable we calculate the expected utility of each. The first options utility is $E[U(x_1)] = 1U(1) = \frac{1^{0.5}}{0.5} = 2$. The second options utility is $E[U(x_2)] = \frac{1}{2}U(2) + \frac{1}{2}U(0) = \frac{1}{2}\frac{2^{0.5}}{0.5} + \frac{1}{2}\frac{0^{0.5}}{0.5} \approx 1.41$. The first option has a high utility and is therefore more preferable.

Table 2.1: Common utilities.

Name	$U(x)$
Risk neutral	x
Quadratic	$x - \alpha x^2$
Power	$\frac{x^{1-\alpha}}{1-\alpha}$
Logarithmic	$ln(x)$
Exponential	$-e^{-\alpha x}$

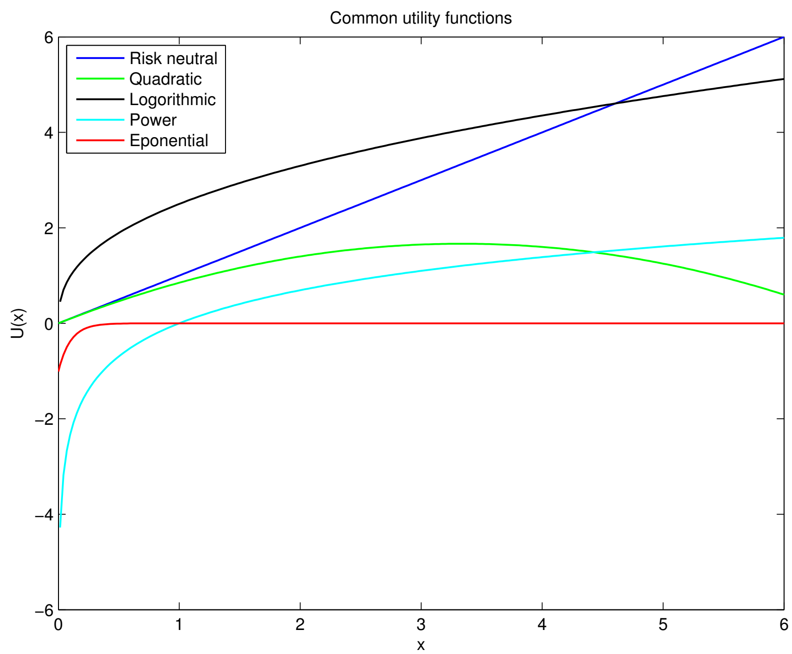

Table 2.2: Some common utility functions.

Expected Utility Maximisation

Suppose an investor has a wealth W_t to invest at time t and uses a utility for wealth U. The investor wishes to form a portfolio such that wealth at time $t+1$, W_{t+1} is maximised. Thus we have the following optimisation problem:

$$\begin{aligned} \underset{w}{\text{maximise}} \quad & \mathbb{E}[U(W_{t+1})] \\ \text{subject to} \quad & W_{t+1} = (1 + r^P_{t+1})W_t \\ & \boldsymbol{w} \cdot \boldsymbol{\iota} = 1 \\ & \boldsymbol{w} \geq 0 \end{aligned} \tag{2.1.13}$$

The solution to this Quadratic Program (or unconstrained convex problem if the portfolio constraint is removed) is the optimal portfolio for this investor.

2.2 General Equilibrium Theory

We know turn our attention to *General Economic Theory*, a key research area in theoretical economics and the foundation of much of our work.

The foundations of general equilibrium theory where introduced by Léon Walras in his 1874 work, *Elements of Pure Economics*.[7] Here, Walras defined an abstract market model. The Walrasian view of the economy is well described by Katzner (1989), which we summarise here. A Walrasian economy is seen as a collection of agents, each of which is a consumer, or a producer. Consumers make buying decisions on

[7] See Walras (2003) for a more-recent republished version of the original work.

commodities produced by the producers, and selling decisions on their own resources (e.g. labor time) while producers decide production quantities. Consumers formulate preferences through utility maximisation, subject to budget constraints; producers maximise profit subject to resources constraints. Consumers and producers interact through perfectly competitive markets, such that markets are fully defined by a pricing system.

Many general equilibrium models still conform to this framework, however, the distinction of a *producer* as defined above is often not clear. For example in a (secondary) stock market, there is no commodity production. Thus we need a slightly modified framework for such purposes.

Definition (Pure Exchange Economy)**.** A pure exchange economy is a Walrasian economy without production, more formally - there exists a bundle of L commodities $x = x^L = (x^1, x^2, \ldots, x^L) \in \mathbb{R}^L$. A price vector for these commodities $p = p_j = (p_1, p_2, \ldots, p_L) \in \mathbb{R}^L$. Thus the wealth can be calculated as $p \cdot x$. Each of the m agents has an initial endowment $e_i \in \mathbb{R}^L$, giving the economy supply vector $\psi = (\psi_1, \psi_2, \ldots, \psi_m)$. Agent i's demands for the assets are given by a demand function $f_i(p, \omega_i)$, where W_i is the wealth of agent i.

Note we have not elaborated a theory of money, however, it is however useful to define a commodity similar to money. We call this special commodity the *numeraire*. It serves as a base reference for all other commodities and is generally used to normalise the price vector.[8]

2.2.1 Walrasian Equilibrium

We now describe the equilibrium argument in the Walrasian setting, this theory applies equally to the exchange economy definition.

Supply & Demand

The equilibrium argument is built upon the notion of *supply & demands*. For a market to exist we must have at-least two commodities (although one could be money), and at-least two market participants. Suppose an economy exists with two commodity where one is money (we make no restriction on market participant and ignore resource requirements). Consumers formulate demands for the commodity at the current price, and simultaneously, producers make supply decisions. There are three possibilities, as shown in Figure 2.3; if the price is such that supply exceeds demands (P_1), then suppliers will be left with a positive excess amount of the commodity. If demand exceeds supply (P_3) then not all consumer demands can be fulfilled leaving a negative excess. In both situations, prices must adjust in a way to balance supply and demand (P_3). This is called the *equilibrium* price. Preferences for an asset may change for reasons other than price change, e.g. a draught would create a greater demand for bottled water, this is called a *shift*. A shift on a supply & demand curve is shown in Figure 2.4.

Partial Equilibrium

In the simple economy just considered we defined equilibrium when supply and demand excess is zero. Real economies contain many assets, each of which will have a supply and demand. *Partial equilibrium* is when, everything else being equal, a single commodity is in equilibrium. For example, consider an economy with fruit commodities. If we fix the price of oranges, pears etc and concentrate on apples, we could find a price that causes supply and demand of apples to be in equilibrium reasonably easily. We then have a partial equilibrium within our fruit economy.

[8] See Barron et al. (2006) for a more formal definition of the numeraire.

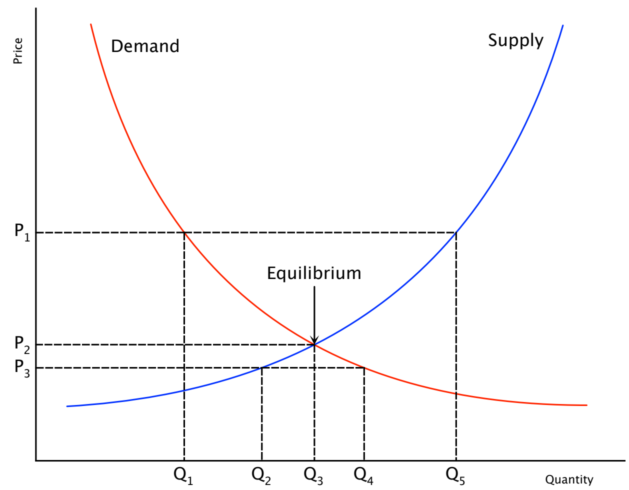

Figure 2.3: A supply and demand curve for a single commodity.

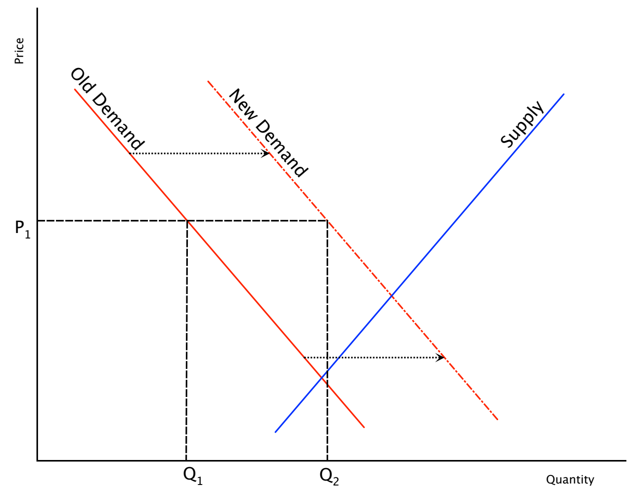

Figure 2.4: A supply and demand curve showing a shift in demand.

General Equilibrium

Partial equilibrium has a fatal limitation: inter-market relationships are not considered (See Starr (1997)). For example, one would expect demand for jet fuel to increase with an increasing demand for flights; partial equilibrium does not account for supply & demand shifts. *General equilibrium theory* looks to find a set of prices such that all markets are simultaneously in equilibrium, it is a beautiful combination of mathematics, economics, and computer science; many consider it akin to finding simplicity within a system of chaos.

A general equilibrium can be shown by an Edgeworth box; a tool visualising efficient allocations in a two good, two agent economy. The indifference curves of both agents are shown for an initial endowment, and by including a price vector we can include a budget line. Figure 2.5 shows an Edgeworth box with two indifference curves passing through the initial endowment vector e. The space shaded between the two curves show allocations which give a *Pareto improvement*.[9] Figure 2.6 shows the indifference curves of the agents under two price vectors, which result in budget constraints 1 & 2. The two dotted line shows each agents demands under all prices p and initial endowment e. We can see that under the first budget constraint, there is dis-equilibrium, where as under the second budget constraint there is global equilibrium. Note the equilibrium point is also Pareto optimal.

We now formally state the definition of general equilibrium using the exchange economy definition. First we define an important property of agent demands functions.

Definition (Walras' Law)**.**

$$p \cdot f_i(p, W_i) = W_i \tag{2.2.1}$$

It is assumed f_i satisfies *Walras' Law* for all i. For any price vector p, total demand in the economy is then given by

$$\sum_{i}^{m} f_i(p, p \cdot \psi_i)$$

[9] *Pareto optimality*, although not directly related to general equilibrium per se, it is intimately related. We say a price vector, $\hat{p}$, is Pareto optimal (or Pareto efficient) if there is no price vector, p, such that $U_i(p) \geq U_i(\hat{p})$ for all i.

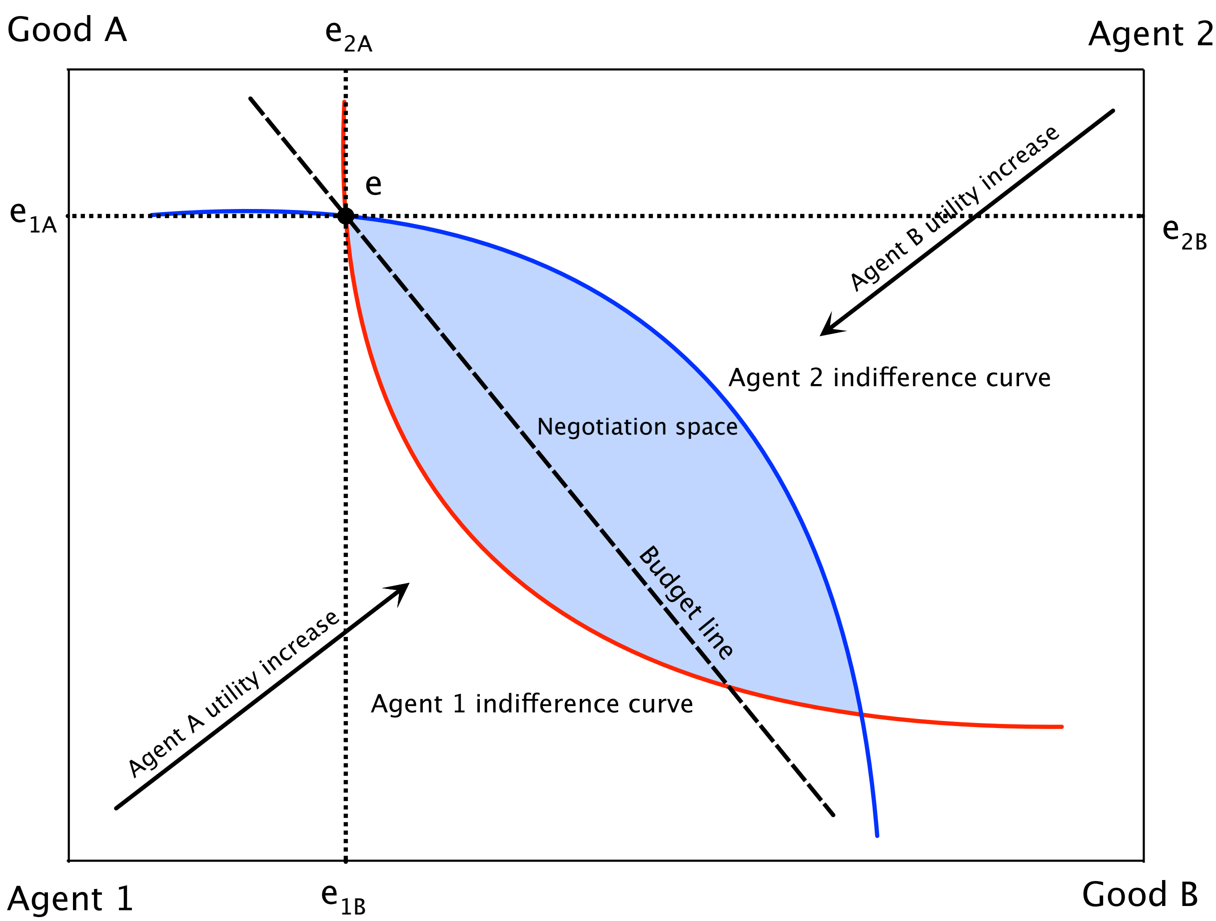

Figure 2.5: Edgeworth box showing efficient allocations. The points in the negation space are Pareto optimal.

Lemma 1. *A price vector p^* is an equilibrium price if*

$$\sum_{i}^{m} f_i(p^*, p^* \cdot \psi_i) = \sum_{i}^{m} \psi_i \tag{2.2.2}$$

Proof. See Appendix 8.2. □

We say an equilibrium is *Walrasian* when each of the agents maximises their utility functions $U_i : \mathbb{R}^L \to \mathbb{R}$, more formally agents solve

$$\begin{aligned} \underset{x_i}{\text{maximise}} \quad & U_i(x_i) \\ \text{subject to} \quad & p \cdot x_i \leq W_i \end{aligned} \tag{2.2.3}$$

The *excess* demand function of each agent i is

$$z_i(p) = f_i(p, p \cdot \psi_i) - \psi_i \tag{2.2.4}$$

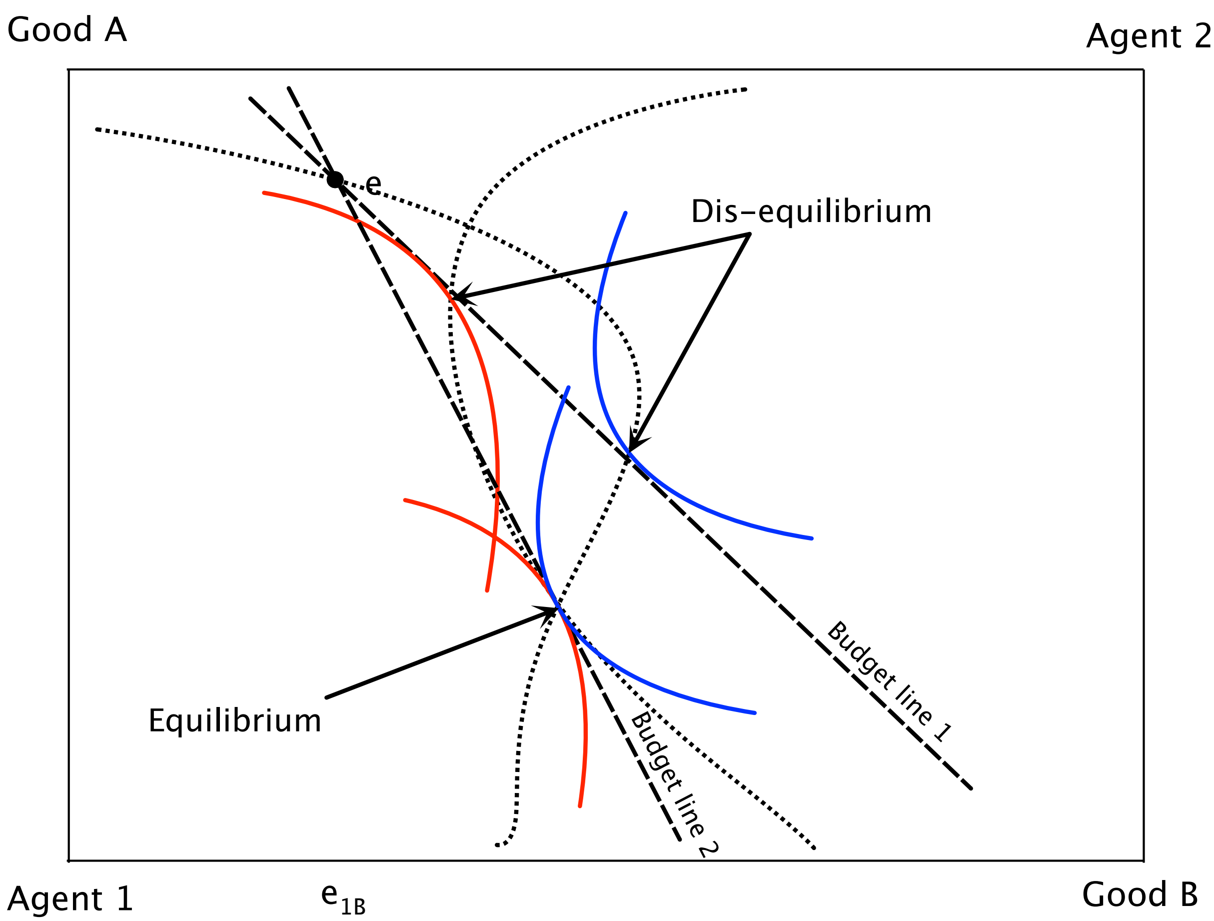

Figure 2.6: Edgeworth box showing dis-equilibrium and equilibrium.

This implies an equivalent definition of Walras' law

$$z_i(p) = 0 \tag{2.2.5}$$

Note, this holds for all prices, not just equilibrium prices. It follows that the aggregate excess demand in the market is given by

$$z(p) = \sum_{i}^{m} (f_i(p, p \cdot \psi_i) - \psi_i)$$

Using this we can obtain the following general equilibrium condition the characterises equilibrium prices

$$z(p^*) = 0 \tag{2.2.6}$$

We can see that a market is in general equilibrium only when each of the assets contained within the market are in partial equilibrium simultaneously.

2.2.2 Questions of General Equilibrium Theory

General Equilibrium Theory raises some important questions about the nature and consequences of equilibrium in an economy. Many of these are addressed by Bryant (2010).

Existence Perhaps the most obvious and important question is whether equilibrium states do indeed exist. This is a surprisingly difficult question to answer, indeed, it was not until 1954 the existence question was proven under the assumption of Arrow-Debreu conditions.

Definition. *Arrow-Debreu* conditions:

I. U_i is continuous.
II. U_i is monotonic.
III. U_i is concave.
IV. $e_i \geq 0$.
V. f_i satisfies Walras' law.
VI. f_i is Marshallian, i.e. is homogenous of degree zero.

Theorem 1. *For any pure exchange economy, if the Arrow-Debreu conditions are satisfied then an equilibrium exists.*

Proof. The basis of the proof comes from *Kakutani's theorem* and *Brouwer's fixed point theorem*. The proof is somewhat involved and is beyond the scope of this thesis. See Arrow and Debreu (1954). Two key observations used in the proof is when conditions 2.2.2 are satisfied, z is continuous and homogenous of degree zero. □

Uniqueness Are equilibrium states unique? That is, is there a single price vector that clears the market or many? It has been proven that under *gross substitutes* there exists a unique equilibrium.

Definition (gross substitutes)**.** A market satisfies the *gross substitutes property* if for any good i, increasing p_i leads to increased demand for all other goods. The market satisfies weak gross substitutes if the demand for every other good increases or stays the same.

Theorem 2. *Under numeraire normalisation, GE markets of gross gross substitutes satisfying the Arrow-Debreu conditions have a unique equilibrium.*

Stability Is there a price adjustment mechanism that always converges to equilibrium prices? This is a key question for *Applied General Equilibrium Theory* which we will discuss next.

Optimality Are equilibrium points Pareto optimal? This property somewhat suggested by the last section. It has been shown that equilibrium states are Pareto optimal (though the converse is not always true). This is know as the *First Fundamental Theorem of Welfare Economics*. It is unlikely we will have much to say on this point.

Sensitivity Do equilibrium states bifurcate under initial conditions? i.e. is the economy a chaotic system? This is a very interesting question and is bound to the realm of higher mathematics.

Empirical congruence Do general equilibrium models give a satisfactorily account of actual economic data? This is perhaps ones of the most important questions in our view. One of our key evaluation points is this very question.

2.2.3 Computing General Equilibria

Having briefly described some of the key points of general equilibrium theory we finish this section with a subject of most interest to us, *how to compute general equilibria*. Much of the work in this area is a result of *Computable General Equilibrium* (CGE) which looks to model empirical data in a general equilibria setting. Many governments use CGE for policy direction (see, for example Shoven and Whalley (1992)).

Walras was the first to introduce a price-adjustment mechanism, which he called *tâtonnement*, defined by the differential equation

$$\dot{p}(t) = \bar{z}(p(t), \omega) \tag{2.2.7}$$

Walras went further with his argument by describing an implementation of tâtonnement: a market 'auctioneer' calls out market prices for each asset, agents within the market formulate demands for each asset given these prices, no trades are completed until equilibrium is achieved.

It is well know that tâtonnement algorithms may fail to converge unless addition assumptions such as gross substitutes are made. The market equilibrium problem can be stated as finding the *fixed points* of 2.2.7. Scarf (1967)[10] was the first to suggest an algorithm for computing approximate fix-points. His algorithm decomposes the *price simplex* into smaller components and uses information about the problem instance to construct a path to an approximate fixed-point. However, the algorithm has exponential running time and requires all of the previous assumptions of general equilibrium existence.

Another class of algorithms developed for computing equilibrium prices apply *global Newton* methods (see, for example Smale (1976), and Kamiya (1990)). These methods do not require gross-substitutes but require Jacobian excess information. Other, auction-style combinational algorithms such as Garg and Kapoor (2004) (surveyed by Noam Nisan and Vazirani (2007b)) have been developed, however, these algorithms require global price information and utility separability over all commodities. More recently, a selection of tâtonnement style algorithm of the form

$$p_i \leftarrow p_i(1 + \lambda_i min\{1, z_i/w_i\}) \tag{2.2.8}$$

have been developed. Cole and Fleischer (2007) argues these algorithms more closely match how prices in real markets are adjusted, however, these algorithm also requires gross substitutes. Finally, there is a class of algorithms that use convex programming (surveyed by Noam Nisan and Vazirani (2007a)) which generally look to maximise utility subject to demands, a hybrid, auction-style convex programming approach using indirect utility has also been developed by Fleischer et al. (2008) which in general requires weak gross substitutes.

2.3 Agent-based Computational Economics

In the Walrasian setting markets operate in an entirely mechanical way, agents form demands through mathematical optimisation with no regard to the environment which they are situated; the current price is the only determining factor in their decision making. The function of the market is thus dependent on the market auctioneer. *Agent-based Computational Economics* removes the dependence of the Walrasian auctioneer, agents are able to act in a operate in a decentralised market and act independently of an equilibrium pricing mechanism. This dramatically changes the analytical and tractable complexity of the Walrasian model as modellers must now consider complexities such as as-symmetric information, strategic interaction, expectation formation, wealth dynamics and social interactions.

[10] See also, Scarf and Hansen (1973) and Scarf (1977).

Tesfatsion (2006) concludes: "The defining characteristic of ACE models is their constructive grounding in the interactions of agents, broadly defined to include economic, social, biological, and physical entities. The state of a modelled system at each point in time is given by the internal data and methods of the agents that currently constitute the system. Starting from an initially specified system state, the motion of the state through time is determined by endogenously generated agent interactions."

Aside. *The term* agent *is often used without caution in economic literature, general equilibrium theory often refers to "agents" in the context of decision making, it is important to recognise that these* representative *agents*[11]*, although of similar vein, represent a different entity to that discussed here.*[12] *Let us be clear that for our purpose unless otherwise stated, agent refers to the bounded rational agents discussed. Further, agents as defined by the Multi-Agent System (MAS) community are seen as autonomous, Russell and Norvig (2010) define an agent as "anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators.". The reader should be made aware that the ACE models described in this thesis are not, in general, the same as those employed by MAS. It is reasonable to say that the "agents" used in ACE models are normal object-oriented programming objects.*[13] *Axelrod (2006) lists some references to agent-based computational literature.*

2.3.1 Motivations

In this section we describe some of the key historical developments that lead to the introduction of ACE. This is not only useful for gaining a better perspective of why ACE is useful, but also because many ACE models use these points as benchmark results. We conclude with some of the key goals of ACE.

Rational Expectations and Bounded Rationality

Rationality is an important concept to to economics, we tend to view humans as rational beings, that is, decisions are made in accordance with reason and logic. This is perhaps unclear, or not strong enough for the purpose of formal reasoning. In general, rationality in economics is concerned with *expectation*, and more specifically *how expectations are formed*. Expectations are said to be *adaptive* when derived from past experience, for example, an agent may calculate expected interest rates based off past interest rates, this is more formally written as,

$$\mathbb{E}[X_t|I_{t-1}] = \mathbb{E}[X_{t-1}] + \lambda(X_{t-1} - \mathbb{E}[X_{t-1}]) \tag{2.3.1}$$

where I_{t-1} is the information available at time $t-1$ and $0 \leq \lambda \leq 1$. In the special case where $\lambda = 1$ this is reduced to $\mathbb{E}[X_t] = X_{t-1}$, known as the Ezekiel model.

Models using the adaptive expectations hypothesis were for some time attacked for information shock susceptibility, as an illustrative example, consider a farmer who's crop plantation decisions are dependent on adaptively expected prices (for when harvest is due), if these expectations are adaptive and there is a severe drought, this will effect yield which will lower price expectations for the following term. It is easy to see how this could lead to dramatic fluctuations in price expectations, and hence price. Secondly, as noted by Gertchev (2007) "When one states that expectations affect reality, if he assumes them to be adaptive, he is ultimately assuming that history, not expectations, affects the future. Adaptive expectations always boil down to a hypothesis of how past variables affect current variables. Therefore, formalising expectations adaptively is contrary to the very purpose of building a theory of expectations.".

[11] Representative agents refer to the homogeneous utility maximisers in Walrasian equilibrium models.

[12] See Kirman (1992) for a critical analysis on the use of representative agents

[13] If the distinction is unclear, see Jennings (1999).

Muth (1961) proposed a new groundbreaking hypothesis of expectations: "I should like to suggest that expectations, since they are informed predictions of future events, are essentially the same as the predictions of the relevant economic theory". This says an individuals expectations are perfectly consistent with realisations; systematic forecasting errors are not made. More formally,

$$X_t = \mathbb{E}_{t-1}[X_{t+k}] = \mathbb{E}[X_{t+k}|I_{t-1}] \tag{2.3.2}$$

Sheffrin (1996), summarises: "Expectations will diverge from actual values only because of some unpredictable uncertainty in the system. If there were no unpredictable uncertainty, expectations would coincide with the actual values - there would be *perfect foresight*.". As noted by Palmer et al. (1994), assumptions are often made when applying rational expectation theory:

- **Complete Information**. Individuals have full knowledge of the problem.
- **Perfect Rationality**. Individuals are able to deduce optimal behaviour or act as if they are.
- **Common Expectations**. Common knowledge[14] of the problem and information is assumed, as well as that all individuals in the system are operating in the same manner.

Rational expectation naturally leads to equilibrium, often refereed to as *rational expectation equilibrium*. Radner (1979) explains: "The possibility for one trader to make inferences from market prices about the information possessed by other traders rests upon his having a "model" or "expectations" of how equilibrium prices are determined, i.e., how equilibrium prices are related to the information initially possessed by the various traders. But this relationship is endogenous to the market system, and if traders have any opportunity to compare the results of the operation of the market with their own models, then a suitable equilibrium concept would require that their models not be obviously controverted by their observations of the market.", in this thesis, he also formalises the definition of a equilibrium under rational expectations. A detailed review is also given by Allen and Jordan (1998).

Definition (Rational Expectations Equilibrium)**.** A rational expectations equilibrium reveals to all traders the information possessed by all of the traders taken together.

Many now believe rational expectations are unrealistic, suggesting the assumptions outlined earlier (complete information etc) are unlikely to be satisfied. The notion of *bounded* rationality, proposed by Simon (1979) suggests man has a limited computational ability in making decisions. In general, bounded rationality is not used for analytical problems, no formal definition has proven tractable, and it evident economists have yet to fully agree on informal characteristics, e.g. Hommes (2006) notes, "A boundedly rational agent forms expectations based upon observable quantities and adapts his forecasting rule as additional observations become available.", where as Palmer et al. (1994) appear to have a different understanding, explaining, "These theories [bounded rationality] impose an intentional limitation on some aspect of an agents task, such as the available knowledge, the computational time or complexity, the memory capacity, the forecasting repertoire, etc. One difficulty is that there are many dimensions in which to bound rationality, and no clear guiding principle for how to set the direction and distance from the zenith of perfect rationality". Gigerenzer and Selten (2002) light-heartedly observes, "bounded rationality has become a fashionable label for almost every model of human behaviour.". We will use the term "bounded" to describe agents not assuming rational expectation.

[14] Common knowledge of some fact expresses that all individuals knows the fact, everyone knows everyone knows the fact. . . etc. See Fagin et al. (2003) for examples.

Heterogeneity and Inductive Reasoning

Under rational expectations assumptions, agents are considered to reason *deductively*, that is, given premises, hypothesis (or expectations) are soundly inferred from a general hypothesis. A well known deductive argument is 'All men are mortal. Socrates is a man. Therefore, Socrates is mortal.'. Reasoning quickly becomes more involved as the problem size increases and one must look to game theory to solve larger scale problems. Arthur (1992) argues humans do not have the capability to reason deductively, proposing we reason *inductively* such that generalisations are formed from past experience. This view is supported by a more formal argument, described in part by Arthur (1995) and Hommes (2001), which we now detail.

Consider a market with unbounded supply of a risk free asset paying fixed return r, and a risky asset pays an exogenously determined stochastic dividends d_t at time $t \in \mathbb{N}$. The risky asset (ex-dividend) price is given by p_t at t. N agents make portfolio adjustments such that z_t shares of the risky asset are owned at t, thus wealth dynamics of agent j are given by

$$W_{i,t+1} = (1+r)W_{i,t} + (p_{t+1} + d_{t+1} - (1+r)p_t)z_{i,t}$$

Publicly available market information ($\{p_i\}$, $\{d_i\}$, economic indicators, news etc) I_t is available at t. Agents are assumed to maximise mean-variance, thus agent j's demand for z solves

$$\max_{z_{i,t}} \mathbb{E}_i[W_{i,t+1}|I_t] - \frac{\alpha}{2} V_i[W_{i,t+1}|I_t]$$

where $\mathbb{E}$ is expectation, V is conditional variance, and α is a risk aversion parameter. Thus

$$\begin{aligned} z_{i,t} &= \frac{\mathbb{E}_i[p_{t+1} + d_{t+1} - (1+r)p_t|I_t]}{\alpha \cdot V_i[p_{t+1} + d_{t+1} - (1+r)p_t|I_t]} \\ &= \frac{\mathbb{E}_i[p_{t+1} + d_{t+1} - (1+r)p_t|I_t]}{\alpha\sigma_i^2} \end{aligned}$$

where $V = \sigma^2 \ \forall i$. Let the fixed supply of the risky asset be given by z^s, then

$$\sum_{i=1}^{N} w_{i,t} \frac{\mathbb{E}_i[p_{t+1} + d_{t+1} - (1+r)p_t|I_t]}{\alpha\sigma_i^2} = z^s$$

where $w_{i,s} = \frac{\sigma_N^2}{\sum_k \sigma_k^2}$ reflecting the weighted average of individuals market expectations. Without loss of generality, let $z^s = 0$. Market equilibrium is then given by

$$p_t = \beta \sum_{i=1}^{N} w_{i,t} \left(\mathbb{E}_i[p_{t+1}|I_t] + \mathbb{E}_i[d_{t+1}|I_t]\right) \tag{2.3.3}$$

where $\beta = \frac{1}{1-r}$. Now assuming rational expectations agents have homogeneous expectations, thus the arbitrage equation 2.3.3 reduces to

$$p_t = \beta(\mathbb{E}[p_{t+1}|I_t] + \mathbb{E}[d_{t+1}|I_t]) \tag{2.3.4}$$

Therefore we must determine $\mathbb{E}[p_{t+1}|I_t]$. Applying 2.3.4 recursively gives

$$\mathbb{E}[p_{t+1}|I_t] = \beta(\mathbb{E}[p_{t+2}|I_t] + \mathbb{E}[d_{t+2}|I_t])$$

But this results in the same dilemma, and we must use the *transversality* condition while applying 2.3.4 repeatedly

Lemma 2. *Transversaity*

$$\lim_{t\to\infty} \beta^k \mathbb{E}[p_{t+k}|I_t] = 0$$

which results in the *fundamental* rational expectations price

$$p_t = \sum_{k=1}^{\infty} \beta^k \mathbb{E}[d_{t+k}|I_t] \tag{2.3.5}$$

Thus shows the fundamental price is a function of the unbiased discounted sum of future dividends, and as such depends on the stochastic dividend process d_t. Note this is the *present value* of rational expectation hypothesis forecasts given by 2.3.2.

Example. If $d_t \sim \mathcal{N}(\mu, \delta^2)$ then,

$$p_t = \sum_{k=1}^{\infty} \beta^k \mu = \frac{\mu}{r}$$

The fundamental price crucially depends on homogeneous conditions; the rational expectation assumptions noted earlier apply. Suppose we assume heterogeneity, the market equilibrium equation 2.3.3 then no longer reduces to 2.3.4. However, we can apply the same logic used to derive 2.3.1 to yield

$$\mathbb{E}_i[p_{t+1}|I_t] = \beta \mathbb{E}_i \left[\sum_{j=1}^{N} w_{j,t} \left(\mathbb{E}_j[p_{t+2}|I_t] + \mathbb{E}_j[d_{t+2}|I_t] \right) \right]$$

This requires agent i to know every other agents expectation of p_{t+1}, this recursive process could be continued such that common knowledge of agent expectation of p is necessary. We conclude under heterogeneous conditions the deductive argument fails for a revision of an agents expectation of others expectations necessarily requires revision for *all* other agents. Such an approach is unstable and we must assume inductive reasoning.

Market Efficiency

The introduction of rational expectations raised many questions about information contained within markets, and in particular, the information revealed by prices. The efficient market hypothesis was formally developed by Fama (1970), three levels of marked efficiency were originally described, but the key insight is information is quickly absorbed by the market; information flow in unimpeded, hence, there is no gain for either the technical or fundamental investor. As described by Dupernex (2007) market efficiency would imply stock prices follow a random walk; price changes are independent, and thus only (presumedly random) new information dictates change. This implies the time series of speculative assets can not be forecasted, past information does not help predict future price movement and investors can not expect above-average returns without accepting above-average risks.

Market efficiency is perhaps one of the most divided topics in economics, many have questioned the empirical testability of the hypothesis and point to stock market events such as the Black Monday crash and the internet stock bubble as evidence against market efficiency. Many of the arguments for and against market efficiency are addressed by Malkiel (2003). Grossman and Stiglitz (1980) found in strong form that market efficiency lead to paradox, explaining, if markets are efficient, and information is costly, then no investor would have incentive to purchase the information, thus full market efficiency is not stable and we must assume information is not costly. Fama (1991) commented “Since there are surely positive information and trading costs, the extreme version of the market efficiency hypothesis is surely false.”. General equilibrium models are of little use exploring market efficiency, agents operate in a static and predictable way thus price dynamics easily predicted.

Trading Volume, Volatility and Fat tails

As detailed by LeBaron (2006), there are some well known empirical markets facts not exhibited by general equilibrium models:

- **High trading volumes**: financial markets generally exhibit large amounts of trading volume. We can see this in Figure 2.7, notice the trading volume is mostly around 4 billion. It is worth noting that trading volumes are also quite stable, which can also be seen in Figure 2.7.
- **Volatility**: market magnitudes are often predictable, moving from periods of relative calm to periods of relative turmoil. We can see this in Figure 2.7, notice the period of high volatility from September to November and relative calm from January onwards. There is also known to be high cross correlation between volatility and trading volume, this can be seen in Figure 2.7 during the period of high volatility (around September); at the same time the trading volume peaks.
- **Fat tails**: first noted by Mandelbrot (1963), over periods of less than a few months returns are not normally distributed, they are usually distributed with too many observations close to the mean, too few in the mid-range, and too many in the tails.[15] This known as a *fat tail* distribution. We can see this in Figure 2.8, here the Normal distribution is shaded in blue and the Fat tail distribution in green, of particular interest is the fat tails which we have labelled.

Figure 2.7: Time series of the Dow-Jones index showing high trading volumes and fluctuating volatility.

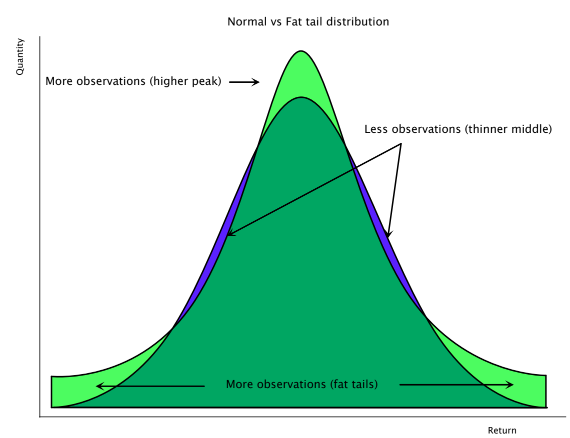

Figure 2.8: Comparison of a Normal distribution and a Fat tail distribution.

[15] Taleb (2010) is a popular text on the distribution of stock-prices, the author argues normal return assumptions is dangerous because any events in the extreme, or what he calls a “black swan” events, have such an impact they should be taken into greater account. He argues in favour of fat tail distributions.

Goals of agent-based computing

Some key goals of ACE are given by Tesfatsion (2006):

- **Empirical understanding**: why have particular global regularities persisted, despite absence of centralised planning or control? Researchers are concerned with building systems populated with agents that realistically represent real investors, they then observe the effects of social interactions and strategic information flow. See for example Chan et al. (1999).
- **Normative understanding**: how can agent-based models be used to discover good economic policy design? Researchers look to build markets that display empirical regularities and design investment policies for a small group of agents, then observe if these agents do better than the market average.
- **Qualitative insight and theory generation**: looks at how agent-based systems may be used to help understand dynamics under alternative initial conditions. Arthur (2006) argues for an agent-based, out-of-equilibrium approach (though he does also suggest a combination of both methods could be employed). This line of research also has links to chaos theory, examining if markets exhibit chaotic behaviour.
- **Methodological advancement**: this area looks to develop theoretical methods that may help ACE researchers study economic systems. Much research in this direction has been undertaken in social interaction theory and machine learning.

2.3.2 Artificial Stock-Market Design

In this section we overview some the key design questions faced by ACE modellers, much of which is described by LeBaron (2001c), LeBaron (2006) and Chang and Harrington (2006).

Pricing Mechanism

Price determination is a fundamental issue for ACE modellers, unlike the Walrasian framework there is no requirement for a centralised auctioneer. The most simple mechanism uses quasi-auctioneer, a market makers calls out prices and orders are calculated. The price then responds using fixed-factor excess adjustment, similar to 2.2.8, of the form

$$p_{t+1} = p_t + \alpha(D(p_t) - S(p_t)) \tag{2.3.6}$$

With this mechanism the market is never in equilibrium. This has may have advantages for adaptive evolving situations but has a big disadvantage in that the market usually does not find clearing prices which leaves the question of how to deal with excesses.

The second approach is to use an equilibrium clearing mechanism. There are three main problems with this method from an ACE perspective. Firstly, there may be too much market clearing; agents are forced to trade on every period. This can be unrealistic if the time periods are taken to be small, especially if trading costs are to be taken into account. Secondly, in real markets agents can trade continuously while this framework requires all agents to trade simultaneously. Thirdly, the computational and mathematical challenges of computing equilibrium prices are very challenging. To the best of our knowledge, beyond very simple cases with a single risky asset, no current artificial stock markets uses this approach.

The trading system which most closely matches real markets is to use an order-book driven system. Here agents place buy or sell orders onto an exchange and a trade takes place when a match is found. The price is then the last strike price. The difficulty with this method is it requires agents to be fully autonomous, it becomes difficult to specify the learning mechanism agents must use. This method requires a truly multi-agent system approach. Some of the leading ACE researchers have suggested this is the most appropriate mechanism.[16]

The last method involves complete decentralisation; agents can trade in an over-the-counter method only e.g. they move around and only trade when they come into contact. This method hasn't been researched in much detail. Like the order-book system, this would require a a truly multi-agent system design.

Strategy and Learning

An important design question is the *type* of agents we want to populate our models. Hommes (2006) distinguishes two types of agent; *fundamentalists* and *chartists*. Fundamentalists base expectations on fundamental economic indicators such as dividend, earnings, macroeconomic growth, unemployment rates etc. They look to buy undervalued stocks and sell overvalued stocks. Chartists, or technical analysts base expectations on speculative strategies such as moving averages, trend following and pattern matching. These types of agent are mostly employed for analytical models that use simple least square learning rules such as those discussed by Bray (1982), who found that under certain initial conditions, the system converges to rational expectations equilibrium.

Most ACE models do not make explicit distinction between agent types, indeed, one of the very goals of ACE is to equip agents with sufficient learning capabilities such that they can autonomously evolve. Genetic algorithms, introduced by Holland (1975) were some of the first evolutionary algorithms to be employed in ACE models. Lettau (1997), the first to use them in an ACE setting found agents took on too much risk, while Duffy (2006) finds that many have found it difficult to interpret the results produced by models using them. The other major criticism of genetic algorithms is they are generally slow to respond to information shocks.

Simpler learning strategies have also been also been suggested, such as "zero intelligence" learning, discussed by Duffy (2006). Agents adopting this strategy behave randomly subject to budget constraints. Zero intelligence learning is generally only used for information gathering and benchmarking, but surprisingly has shown to be very efficient operating in markets with realistic trading environments.

Finally, adaptive learning can be employed. These agents use a variety of techniques from artificial intelligence, there key feature is they must maintain a history of past performance and be able to make inferences from this data. LeBaron (2010) makes use simple Kalman filter style regression forecasts, using a combination of active and passive learning. The model is able to reproduce many of dynamics generated by more complicated learning approaches and is much simpler to analyse. LeBaron (2011) goes on to recommend using a combination of active and passive learning.

Information Representation

Another important design decision is how to agents represent and process market information. Information generally includes price as a minimum, but could include anything from fundamental ratios, other agent performance, news and asset correlations. Classifier systems are often used. A classifier is a bit-string where each bit can take values '1', '0', or '#',[17] each bit represents some market condition e.g. if the stock price went up last period. A classifier is active if it matches the current market condition. Classifier systems have proven very popular. Some well known models using classifiers we were able to re-implement

[16] See for example LeBaron (2006).

[17] '#' means "don't care".

include Palmer et al. (1994), Arthur et al. (1997), LeBaron et al. (1999) and LeBaron (2002). Complexity of information representation varies considerably, for example LeBaron (2001b) and LeBaron (2001a) uses neural networks, while others have used a simple vector of parameters.

Social Interaction

Social interaction is concerned with how agents are able to communicate with each other, if at all. One of the main goals of ACE is to develop a realistic framework that is independent of the pricing system. In real markets agents are able to trade over an exchange or over-the-counter, and although markets now operate on a global scale it is reasonable to suggest not all agents will be in contact with every other agents. Other than having agents interact solely through markets, two alternative approaches have been suggested for modelling this behaviour. The first by Wilhite (2006) suggests agents should operate over fixed networks; agents populate nodes of networks are only able to interact with agents they are directly linked with. Wilhite found changes in dynamics as the underlying network was changed. Vriend (2006) has suggested agents should not only learn to form expectations, but also learn who they wish to interact with, agents may wish only to trade or share information with a group of other agents they find beneficial to do so. Westerhoff (2009) developed a simple market where agents were able to interact and share information. Even in this simple setting, many interesting empirical market dynamics were produced.

Assets

In our view, a crucial design decision is the number of assets that populate the market. This is one of the most underdeveloped areas of ACE, our view is supported by LeBaron (2006) who writes "Another very common and pertinent criticism is that most agent-based financial market models assume a small number of assets. Often agents trade only one risky asset, and one risk-free asset alone, this simplification may eliminate many interesting features.".

Furthermore, it it usual to use assets to introduce some random variation into the market, this generally comes in the form of *dividend* payments. A common technique for dividend modelling is *autoregressive models*,[18] although some models have used simple normally distributed random shocks. Palmer et al. (1994) found that autoregressive models most closely represented real dividend time-series.

2.4 Summary

We began by giving an overview of mean-variance portfolio theory and showed how to calculate optimal portfolios in uncertain markets. We showed how utility functions represent an agents risk-aversion, and how they can be used to construct optimal portfolios.

Next we looked at general equilibrium theory, and in particular, how supply and demand support price movements. We examined how equilibrium prices are formed, and reviewed some of the key questions general equilibrium theory looks to answer. We then reviewed methods for computing general equilibria.

We concluded by examining agent-based computational economics, we gave motivation for using agent-based models and finished by reviewing some key design principles used in current artificial markets.

[18] We discuss autoregressive models in the Appendix (8.4).

3

Artificial Stock Market Model

In this section we develop our artificial stock market model. We give motivations for using certain methods with reference to the background material and in particular, we follow the artificial stock market design principles discussed in Section 2.3.2. We begin with a general overview of the market and follow with detailed mathematical arguments for each part of our design.

3.1 Model Overview

Figure 3.1 demonstrates the high level working of our model. The first key observation is that the dividend generation is set exogenously of the market, the markets sole responsibility is price setting through market clearing. Agents indirectly interact with the market by making portfolio decisions. Agents use market information and feedback from past portfolio decisions as input to a learning mechanism to make adjustments to forecast rules. Forecasts are independent of portfolio selection and learning, this is to allow easier modification of the agents learning mechanism.

As suggested in Section 1, our model will contain multiple risky assets. This raises two fundamental design questions:

- **Asset interactions**: There is strong empirical evidence asset prices are correlated. This is not surprising, for example one would expect the price of stocks in the same sector to closely track each other. We could just ignore this factor, like in Westerhoff (2009) and Palmer et al. (1994), but we feel this removes an important fundamental characteristics of real assets. This raises the question of how we should model asset interactions.
- **Price setting**: As we have already discussed in great detail, we must find a suitable price setting mechanism. this is the most important component in our model and we defer development of our pricing mechanism to the next chapter.

3.2 Assets and Dividends

We define assets in terms of dividends structure and emphasis that asset fundamentals are captured by an underlying probabilistic distributions. The first important observation is that real dividends are always

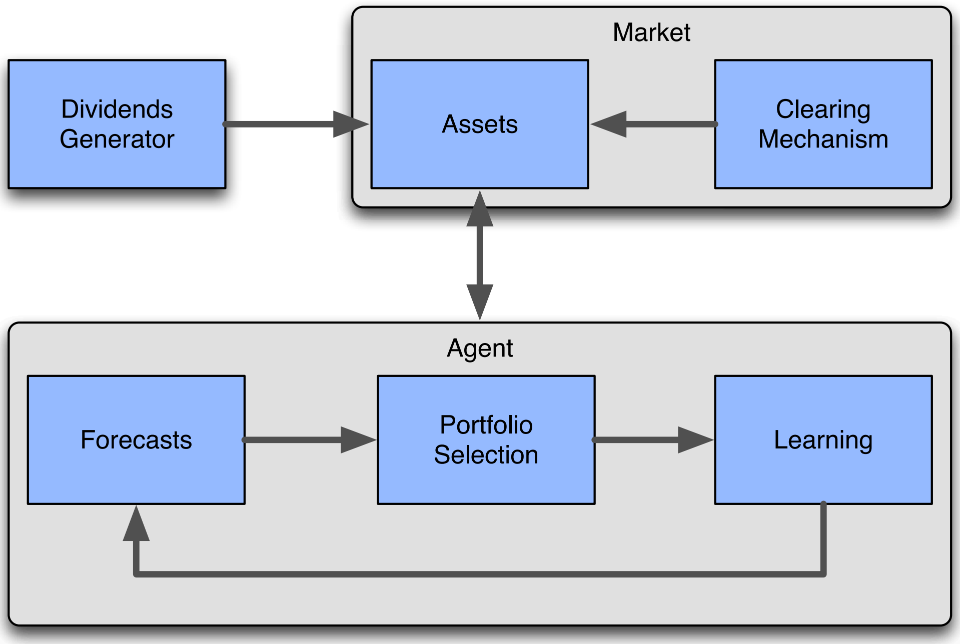

Figure 3.1: Overview of our artificial stock market.

positive. However, time-series generated by Gaussian models are unbounded in the y-plane, thus we must find a way of making samples positive. The first solution would be to simply truncate any negative samples, but although we have seen this employed in some models, we do not feel it is an acceptable approach. In our opinion, a much better solution is to use a log-normally distributed model which is bounded to positive samples. In Section 2.3.1 we showed there is strong evidence returns are log-normally distributed, and although we were unable to find any work relating specifically to dividend return distributions, it seems reasonable to suggest this is also the case. This hypothesis seems to be supported by the samples we observed, shown in Figure 3.2, which show a log-normal shape. The second advantage of a log-normal distribution compared to the truncation approach is statistical moments can be readily calculated, this will be useful for calibration and benchmark analysis.

In Section 2.3.2 we discussed two ways to model dividend structure; Gaussian noise or an autoregressive model. Although a Gaussian noise process would be easier to implement and analyse, it does not have any empirically supported properties such as mean reversion and positive autocorrelation, and is very inflexible. A vector autoregression is essentially a discrete-time version of an Ornstein-Uhlenbeck process, and thus shares many of the same properties such as mean-revision, drift and positive autocorrelation. However we stress that to the best of our knowledge, no study on dividend return distributions or suitable supporting models has been published, therefore we are making entirely debatable choices. Nevertheless, our hypothesis do seem to have some empirical support, which can be seen by looking at Figures 3.3, 3.4 & 3.5. Figure 3.3 shows a real set of dividend time-series, note there is little volatility, there seems to be a high level of autocorrelation and substantial drift. Figure 3.4 shows a geometric[1] multi-variate time-series, calibrated with data from the real dividends. Notice the high volatility and zero drift. Finally, Figure 3.5 shows a geomet-

[1]By geometric, we mean exponentiated, i.e. the log is normally-distributed.

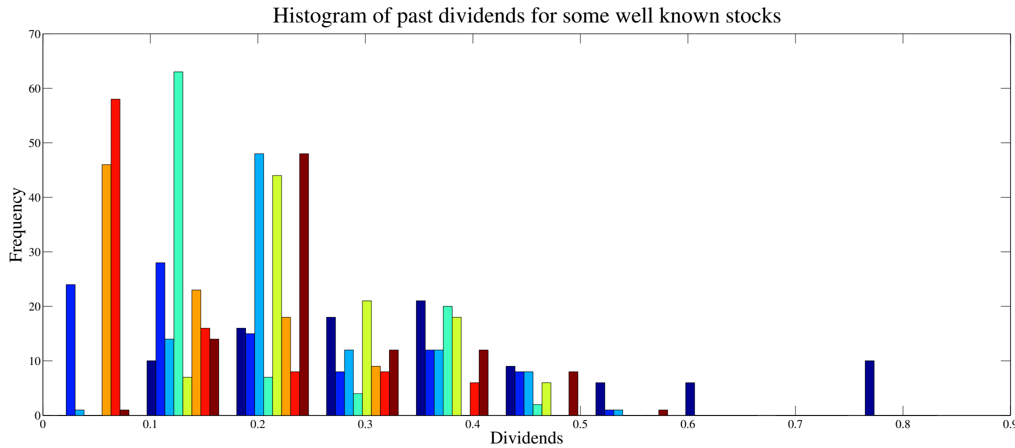

Figure 3.2: Histogram of dividends returns for some well known, dividend-paying stocks.

ric vector autoregression[2] loosely calibrated with the empirical dividends data, observe how the time-series displays many of the same properties the observed dividends does, [3] most notably a low volatility and high autocorrelation. Although there is not as much drift as observed in the empirical data this could be achieved with better calibration. We conclude that geometric vector autoregression's are the most suitable method for modelling dividend time-series.

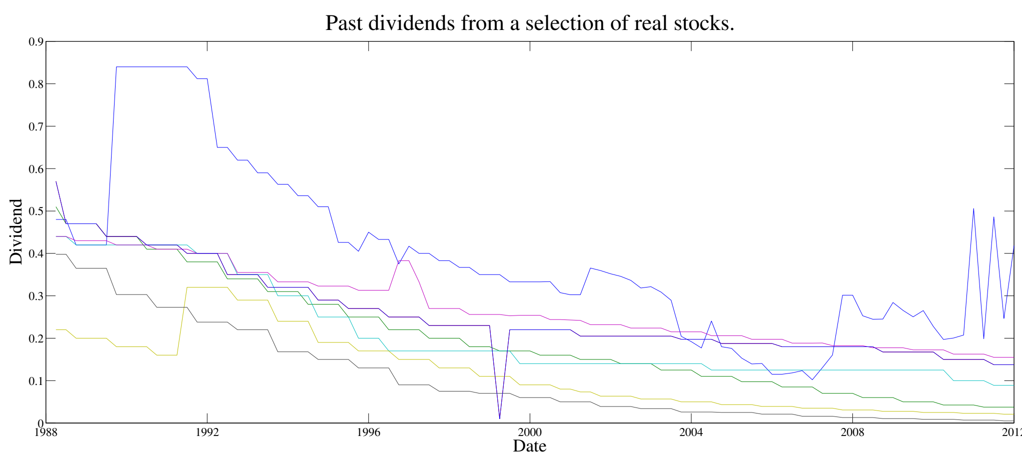

Figure 3.3: Time-series of real dividends taken from some well known stocks.

Given the evidence just considered we decided to use a geometric vector autoregression to model dividend. As described in the overview, this can easily be modified without disrupting the rest of the model if required. We finish our discussion on assets and dividends by formally defining geometric vector autoregressions and deriving mean and covariance statistics which will be required when calibrating and evaluating our model.

[2] An examination of vector auto regressions is given in the appendix (8.4).

[3] We sat "loosely calibrated" as little effort was made to fine tune the model to the empirical data, we just wanted to demonstrate the basic dynamics the model can produce. However, with more thorough calibration we believe the model could reproduce the observed data even more reliably.

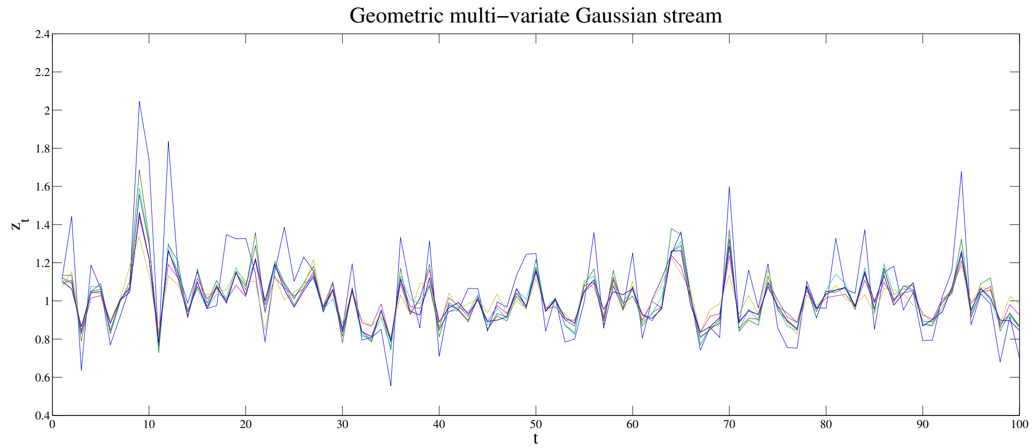

Figure 3.4: A geometric multi-variate time-series.

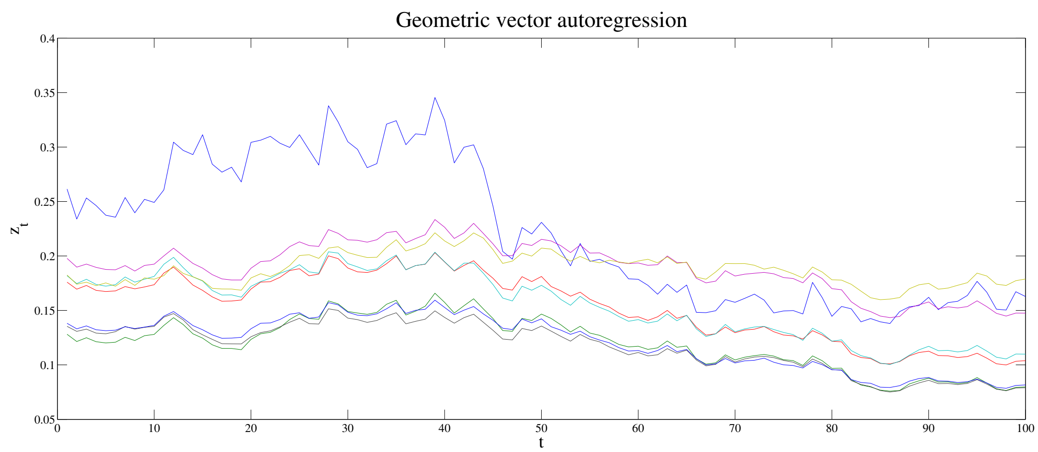

Figure 3.5: A geometric Vector Autoregression.

Definition. Suppose X is a normal vector autoregression of the form $X_t = \boldsymbol{c} + \Phi X_{t-1} + \boldsymbol{\epsilon}_t$. A geometric vector autoregression, Y, is defined as

$$Y_{i,t} = e^{X_{i,t}} \tag{3.2.1}$$

Thus, the dividends, $\pi_{j,t}$ of asset j at time t is given by $Y_{j,t}$.

Proposition 1. *The mean, $\boldsymbol{\mu}$, of a geometric vector autoregression, Y, is given by*

$$\mu_i = e^{((\iota - \Phi)^{-1}\boldsymbol{c})_i + \frac{1}{2}\mathbb{D}\left(\sum_{k=0}^{\infty} \Phi^k \Sigma \Phi^{T^k}\right)_i} \tag{3.2.2}$$

where e is the piecewise, vector function and $\mathbb{D}$ is the diagonal operator.

Proof. Trivial application of Lemmas 14, 15 and 8. ◻

Proposition 2. *The covariance matrix, Γ, of a geometric vector autoregression, Y, is given by*

$$\gamma_{ij} = (e^{\sigma_{ij}} - 1)\left(e^{\hat{\mu}_{ii}+\hat{\mu}_{jj}+\frac{\sigma_{ii}+\sigma_{jj}}{2}}\right) \tag{3.2.3}$$

where $\hat{\mu}_{ij} = (\boldsymbol{\iota} - \Phi)^{-1}\boldsymbol{c}$ and $\sigma_{ij} = \sum_{k=0}^{\infty} \Phi^k \Sigma \Phi^{T^k}$.

Proof. Trivial application of Lemmas 14, 15 and 9. □

3.3 Agents

Stock markets are driven by thousands of self interested heterogenous investors. To precisely model the behaviour of real investor is surly an impossible task, nevertheless, it is vital the agents in our market capture the basic behaviour of real investors in as realistic way as possible; we must model investors *preferences* accurately. In Section 2.2.1 we revealed preferences in a general equilibrium framework are modelled according to a utility maximisation problem which we described in Section 2.1.2. As discussed in Section 2.3.1 these agents conform to the rational expectations hypothesis; agents have perfect knowledge of their environment, however, in Section 2.3.1 we showed in a heterogeneous framework the rational expectations hypothesis fails. In Section 2.3.2 we briefly discussed alternative learning methods for forming expectations. Given the evidence we have considered, in our opinion the learning approach is most suitable. Baring this in mind, along with the other design questions studied in Section 2.3.2 we must answer the following:

- **Portfolio selection**: our model contains multiple risky assets, our heterogeneous agents must choose a portfolio that they deem will benefit them the most.
- **Return expectation**: we must decide how agents form expectations using a some learning method. Agents should use what they have learned to *forecast* future returns.
- **Social interactions**: We must decide if our agents are to communicate with each other, and if so; what information should they share? Who should they share information with such that sharing is mutually beneficial?

3.3.1 Portfolio Selection

Agents portfolio choice is determined by a myopic power utility function in future wealth. This is a reasonable approach given our proposed objectives (computing general equilibrium). Therefore our agents must maximise their utility, which, as explained in Section 2.1.2 requires solving the following maximisation problem:

$$\begin{aligned} \underset{w_t}{\text{maximise}} \quad & \mathbb{E}[U(W_{t+1})] \\ \text{subject to} \quad & \boldsymbol{w}_t \cdot \boldsymbol{\iota} = 1 \\ & W_{t+1} = (1 + r_P)W_t \\ & \boldsymbol{w}_t \geq 0 \end{aligned}$$

We choose a utility function that satisfies the relevant Arrow-Debereu equilibrium conditions 2.2.2. From the utility functions we have seen in Section 2.1.2 this leaves only power utility, exponential utility and log utility; all three are concave and monotonic. We choose the power utility for two reasons: firstly, the power utility exhibits constant relative risk-aversion; secondly, our model has lognormal dividends returns, with lognormal returns and power utility the maximisation problem is simplified. Thus agents must solve the following optimisation problem:

$$\begin{aligned} \underset{w_t}{\text{maximise}} \quad & \mathbb{E}\left[\frac{W_{t+1}^{1-\gamma}}{1-\gamma}\right] \\ \text{subject to} \quad & \boldsymbol{w}_t \cdot \boldsymbol{\iota} = 1 \\ & W_{t+1} = (1+r_P)W_t \\ & \boldsymbol{w} \geq 0 \end{aligned} \tag{3.3.1}$$

Assuming lognormal returns 3.3.1 can be simplified as shown by Campbell and Viceira (2002). We make use of the following result:

Lemma 3. *Maximising the expectation is equivalent to maximising the log of the expectation, that is,*

$$x \in \underset{x}{arg\ max} f(x) \ \ iff \ \ x \in \underset{x}{arg\ max} \log f(x)$$

Proof. log is monotonic. □

The first observation is we can remove the scaling factor $\frac{1}{1-\gamma}$. Thus the objective function problem becomes

$$\underset{w_t}{\text{maximise}} \quad \ln \mathbb{E}\left[W_{t+1}^{1-\gamma}\right]$$

Using Lemma 8, we reform the objective function as follows:

$$\begin{aligned} \ln \mathbb{E}\left[W_{t+1}^{1-\gamma}\right] &= \mathbb{E}\left[\ln W_{t+1}^{1-\gamma}\right] + \frac{1}{2}\text{Var}\left[\ln W_{t+1}^{1-\gamma}\right] \\ &= \mathbb{E}\left[(1-\gamma)\ln W_{t+1}\right] + \frac{1}{2}\text{Var}\left[(1-\gamma)\ln W_{t+1}\right] \\ &= (1-\gamma)\mathbb{E}\left[\ln W_{t+1}\right] + \frac{(1-\gamma)^2}{2}\text{Var}\left[\ln W_{t+1}\right] \\ &= (1-\gamma)\mathbb{E}\left[\ln W_{t+1}\right] + \frac{(1-\gamma)^2}{2}\boldsymbol{w}_t \cdot (\Sigma_{\omega t}\boldsymbol{w}_t) \end{aligned} \tag{3.3.2}$$

where $\Sigma_{\omega,t}$ is the covariance of log-wealth at time t. Taking logs of the budget constraint gives an expression for the current log wealth

$$\begin{aligned} \ln W_{t+1} &= \ln((1+r_{P,t+1})W_t) \\ &= \ln(1+r_{Pt+1}) + \ln W_t \end{aligned}$$

which also gives an expression for the next periods log returns

$$\eta_{t+1} = \omega_{t+1} - \omega_t \tag{3.3.3}$$

Dividing 3.3.2 by $1-\gamma$ and using 3.3.3 gives the agent problem

$$\begin{aligned} \underset{w_t}{\text{maximise}} \quad & \mathbb{E}[\eta_{t+1}] + \frac{1-\gamma}{2}\boldsymbol{w}_t \cdot (\Sigma_t \boldsymbol{w}_t) \\ \text{subject to} \quad & \boldsymbol{w}_t \cdot \boldsymbol{\iota} = 1 \\ & \boldsymbol{w} \geq 0 \end{aligned} \tag{3.3.4}$$

Where Σ_t is the conditional covariance matrix of log portfolio returns. Notice the important role of γ in 3.3.4 - when $\gamma < 1$ agents seek portfolios with more risk; they prefer a portfolio with a higher log-return variance. When $\gamma > 1$ agents are risk-averse and prefer portfolios with a lower log-return variance. In the special case when $\gamma = 1$ agents have logarithmic utility and choose the portfolio with the highest log return. We can now show that our agents are indeed solving the mean-variance criterion.

Proposition 3. *Agents solving the optimisation problem 3.3.6 are mean-variance optimisers.*

Proof. Using Lemma 8 we have

$$\mathbb{E}[\eta_{t+1}] + \frac{1}{2}\boldsymbol{w}_t \cdot (\Sigma_t \boldsymbol{w}_t) = \ln\mathbb{E}[1 + r_{P,t+1}] \tag{3.3.5}$$

using this and the fact that portfolio returns are lognormal we can write 3.3.6 as

$$\begin{aligned} \underset{w_t}{\text{maximise}} \quad & \ln\mathbb{E}[1 + r_{p,t+1}] - \frac{\gamma}{2}\boldsymbol{w}_t \cdot (\Sigma_t \boldsymbol{w}_t) \\ \text{subject to} \quad & \boldsymbol{w}_t \cdot \boldsymbol{\iota} = 1 \\ & \boldsymbol{w} \geq 0 \end{aligned} \tag{3.3.6}$$

Which is a mean-variance optimisation problem. □

To solve the agent problem 3.3.4 we need to specify $\boldsymbol{\eta}_{t+1}$, the portfolio log-return. The return of a portfolio is given in 2.1.2, thus we are required to find $ln(\boldsymbol{w}_t \cdot \boldsymbol{r}_t + 1)$ in terms of log returns on the underlying assets. This is problematic as in general, the log of a linear combination is not the same as a linear combination of logs. However, for short time-periods we can approximate this to an arbitrary degree using a *Mercator series*.

Definition. A Mercator series is a Taylor series of the natural logarithm

$$ln(1 + x) = \sum_{i=1}^{\infty} \frac{(-1)^{n+1}}{n} x^n$$

We use a Mercator approximation to the first degree which is more than adequate for short-period portfolio rebalancing (e.g. weeks), thus we have

$$ln(\boldsymbol{w}_t \cdot \boldsymbol{r}_t + 1) = \boldsymbol{w}_t \cdot \boldsymbol{r}_t$$

Therefore

$$\mathbb{E}[\boldsymbol{\eta}_{t+1}] = \mathbb{E}[\boldsymbol{w}_t \cdot \boldsymbol{r}_t] = \boldsymbol{w}_t \cdot \bar{\boldsymbol{r}}_t$$

Substituting this into 3.3.6 reveals our final agent problem

$$\begin{aligned} \underset{\boldsymbol{w}_t}{\text{maximise}} \quad & \boldsymbol{w}_t \cdot \bar{\boldsymbol{r}}_t + \frac{1-\gamma}{2}\boldsymbol{w}_t \cdot (\Sigma_t \boldsymbol{w}_t) \\ \text{subject to} \quad & \boldsymbol{w}_t \cdot \boldsymbol{\iota} = 1 \\ & \boldsymbol{w}_t \geq 0 \end{aligned} \tag{3.3.7}$$

Throughout this analyse we assumed our agents could not short sell. This is probably sensible restriction to begin with as the model complexity would substantially increase if we were to allow short selling. This is because agents could potentially become bankrupt and there is no clear method for handling bankruptcy (do bankrupt agents simply get removed from the market? What happens to the wealth they have lost?). Also, there are additional implementation considerations when allowing short selling such as how dividends are paid to the original stock owner, this would require some system for handling borrowing, no easy task when assets are infinitely divisible. Thus, at-least to begin with we do not allow short selling. The disadvantage of not allowing short selling is the agent problem is a quadratic program which is generally not easy to solve, running-time then becomes an important issue and it will therefore be important to make sure we use a good quality optimisation algorithm.

For completeness we derive the agents optimal portfolio weights with short selling allowed. First we need to find the Lagrangian of the problem, which is given by

$$L(w_t, \lambda) = \boldsymbol{w}_t \cdot \bar{\boldsymbol{r}}_t + \frac{1-\gamma}{2}\boldsymbol{w}_t \cdot (\Sigma_t \boldsymbol{w}_t) - \lambda(\boldsymbol{w}_t \cdot \boldsymbol{\iota} - 1) \tag{3.3.8}$$

The KKT conditions are then

$$\frac{\partial L(w_t, \lambda)}{\partial \boldsymbol{w}_t} = \bar{\boldsymbol{r}}_t + (1-\gamma)\Sigma_t \boldsymbol{w}_t - \lambda \boldsymbol{\iota} = 0 \tag{3.3.9a}$$

$$\frac{\partial L(w_t, \lambda)}{\partial \boldsymbol{\lambda}} = \boldsymbol{w}_t - 1 = 0 \tag{3.3.9b}$$

We can analytically find the optimal portfolio weights by solving the KKT conditions

$$\begin{aligned} \bar{\boldsymbol{r}}_t + (1-\gamma)\Sigma_t \boldsymbol{w}_t - \lambda \boldsymbol{\iota} &= 0 \\ \Sigma_t \boldsymbol{w}_t &= \frac{1}{1-\gamma}(\lambda \boldsymbol{\iota} - \bar{\boldsymbol{r}}_t) \\ \Rightarrow \boldsymbol{w}_t^* &= \frac{1}{1-\gamma}\Sigma_t^{-1}(\lambda \boldsymbol{\iota} - \bar{\boldsymbol{r}}_t) \end{aligned} \tag{3.3.10}$$

$$\boldsymbol{w}_t \cdot \boldsymbol{\iota} = 1 \tag{3.3.11}$$

Substituting 3.3.10 into 3.3.11

$$\begin{aligned} 1 &= \left(\frac{1}{1-\gamma}\Sigma_t^{-1}(\lambda \boldsymbol{\iota} - \bar{\boldsymbol{r}}_t)\right) \cdot \boldsymbol{\iota} \\ &= \frac{1}{1-\gamma}\Sigma_t^{-1}\lambda \boldsymbol{\iota} \cdot \boldsymbol{\iota} - \frac{1}{1-\gamma}\Sigma_t^{-1}\bar{\boldsymbol{r}}_t \cdot \boldsymbol{\iota} \\ \Rightarrow \lambda &= \frac{1-\gamma + \bar{\boldsymbol{r}}_t \cdot (\Sigma_t^{-1} \cdot \boldsymbol{\iota})}{\boldsymbol{\iota} \cdot (\Sigma_t^{-1} \cdot \boldsymbol{\iota})} \end{aligned} \tag{3.3.12}$$

Substituting 3.3.12 into 3.3.10 gives the optimal weights

$$\boldsymbol{w}_t^* = \frac{1}{1-\gamma}\Sigma_t^{-1}\left(\frac{1-\gamma + \bar{\boldsymbol{r}}_t \cdot (\Sigma_t^{-1} \cdot \boldsymbol{\iota})}{\boldsymbol{\iota} \cdot (\Sigma_t^{-1} \cdot \boldsymbol{\iota})}\boldsymbol{\iota} - \bar{\boldsymbol{r}}_t\right) \tag{3.3.13}$$

3.3.2 Demands

Now that we have decided how our agents form their portfolios, we must determine how they make demands for a given price vector. To do this requires some knowledge of how the market operators, specifically, the

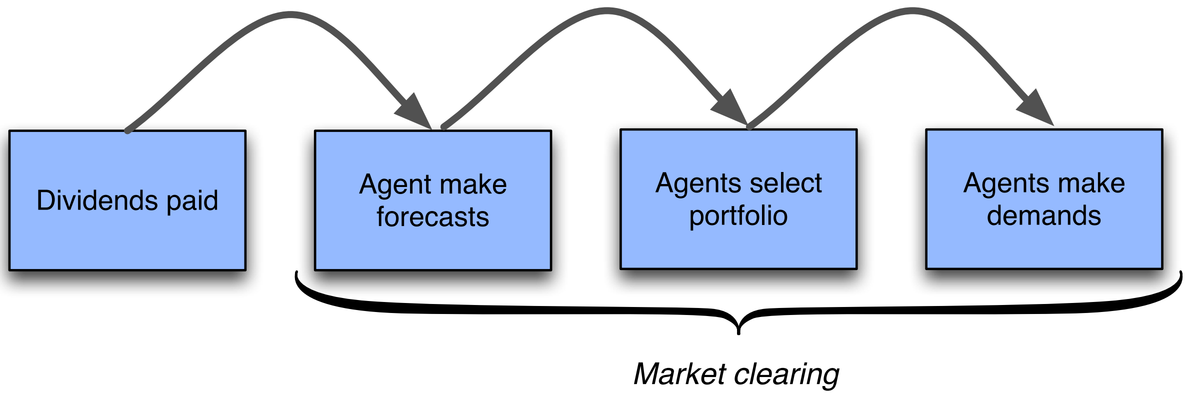

Figure 3.6: Overview of market timing.

market *timing*. An overview of market timing is given in Figure 3.6. The key observation is that the market starts to clear **after** dividends have been paid. In a discrete world market timing is one of the hardest aspects of a real market to mimic, but all things considered this seems like the most sensible approach. Agent i's inter-temporal budget constraint are thus given by

$$W_{i,t+1} = (1 - \lambda_{i,t})(1 + r_{i,P})W_{i,t} \tag{3.3.14}$$

where $\lambda_{i,t}$ is agent i's consumption at time t. We introduce consumption now as a precautionary measure; we do not yet know what impact wealth dynamics will have on market dynamics and we may need to introduce consumption if wealth dynamics behave in an is unexpected way, such as an exponentially increase. Agent i's current period budget constraint is therefore

$$B_{i,t+1} = W_{i,t} + \delta(\boldsymbol{\pi}_{t+1}) \cdot \boldsymbol{\nu}_{i,t} \tag{3.3.15}$$

where $\nu_{i,t+1}$ is agent i's ownership vector at time t (i.e. how many shares of each asset does the agent own) and

$$\delta(\nu_i) = \begin{cases} \boldsymbol{\nu}_i & \text{if } \boldsymbol{\nu}_i > 0 \\ 0 & \text{otherwise} \end{cases}$$

δ stops agents taking dividends for stocks they have borrowed. Note this does not completely solve the problem with short sales previously outlined, we only include this into our model now as it has no effect on our market with short sales not allowed.

Figure 3.6 reveals another crucial part of our model, agent *forecasts*. Up to this point our model has somewhat paralleled a pure exchange economy model defined Section 2.2, however, forecasts do not form part of this definition. This is one of the key difference between our model and a pure exchange economy. One of our main challenges is to investigate what impact this has on general equilibria in our market.

We focus on how agents make forecasts in the next section, but to derive the agent demand function we must introduce some notation for agent forecasts. We define $\boldsymbol{\tau}_{i,t}$ as agent i's forecast vector at time t, and write agent i's optimal portfolio choice, $\boldsymbol{w}^*$, as a function of $\bar{\boldsymbol{r}}$. It follows that agent i's demand for asset j at time t, $\zeta_{i,t}$, is given by

$$\zeta_{i,t}(p_{t,j}) = \frac{w^*_{i,t}(\tau_{i,t}(p_{t,j}, p_{t-1,j}, \pi_{t,j}))_j B_{i,t}}{p_{t,j}} \tag{3.3.16}$$

3.3.3 Expected future returns forecasts

As we have just discussed, agents expectations of future returns are given by a forecast. As described in Section 2.3.1 this is in contrast to general equilibrium models where agents conform to the rational expectations hypothesis. We also looked specifically in Section 2.3.2 at current methods for forming forecasts. Following from this we derive a forecasting scheme with a learning mechanism, depicted in Figure 3.7. The functionality of the two components is as follows:

- **Heterogeneous forecast**: forecasts take the current periods return (variant on the price generated by the pricing mechanism) and use past market data history to generate return expectations for next period. This operation is shown in the diagram on the left hand side; the forecast indirectly receives input from the clearing mechanism, the agent uses the produced forecast to calculate demands, the dotted line indicates the relationship is indirect. Note the dashed line between the clearing process and market history is just there to show the clearing mechanism indirectly updates the market history.

- **Learning mechanism**: learning takes place at the beginning of each period (if at-all). The learning mechanism receives the last clearing price indirectly from the clearing mechanism and uses the market history and the previous forecast to make adjustments to the agents forecasts. Note agent do not *have* to apply learning to their forecasts, we could decide to apply learning probabilistically, hence why the two components are independent.

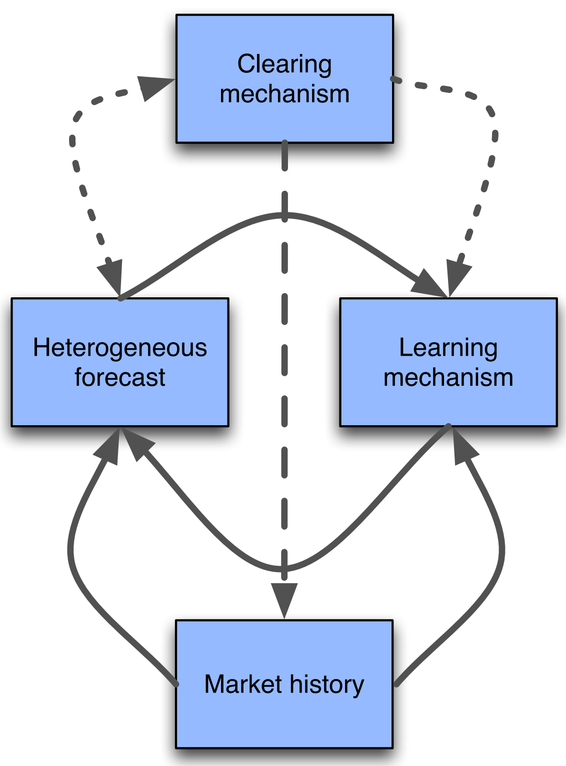

Figure 3.7: Agent forecasting mechanism.

Before continuing we briefly review some current forecasting approaches. In section 2.3.2 we mentioned that the forecast model choice really comes down to the types of investor one is trying to model. Generally analytical models split agents into two categories, fundamentalists and chartists. Each agent type has a single homogeneous forecast rule. For example, in Hommes (2001) both agents future price expectation is as follows:

$$\mathbb{E}[\boldsymbol{p}_{t+1}] = \mathbb{E}[\boldsymbol{p}^*_{t+1}] + f(x_{t-1}, \ldots . x_{t-L})$$

where $x_t = p_t - p^*_t$, representing the deviation from the fundamental price p^* which both types of agent know, L is a lag factor. For a single lag, f is of the form:

$$f = gx_{t-1} + b$$

where g and b determine the trader type. When $g = b = 0$ the rules reduces to $\mathbb{E}[\boldsymbol{p}_{t+1}] = \mathbb{E}[\boldsymbol{p}^*_{t+1}]$ representing fundamentalist expectations. For other values of g and b expectations follow speculative strategies. Georges (2008) uses a similar forecast structure but his model only contains "technical traders". These "few-type" models are easy to analyse but in general do not produce as many market dynamics as more complicated strategies. Another problem we have with few-type models is it is inherently abstract, we have found no evidence that the strategies employed by each "type" of agent actually correspond to the trading strategies employed by the investors they are attempting to represent.

"Many-type" models do not split agents into categories, agents *evolve* over time, this method is thus intimately related to the learning mechanisms employed. We re-implemented the many-type model described by LeBaron et al. (1999) which uses classifiers and a genetic algorithm and found that although the model produced interesting dynamics, it was difficult to analyse. These sentiments have been echoed more recently by LeBaron (2011).

We follow the advice of LeBaron (2011) and use "few-type" set of learning rules with active and passive learning mechanisms, developed in a previous publication by the same author, LeBaron (2010). This should allow easier analysis of our results while still allowing more complex agent dynamics than a standard "few-type" model. Moreover, as suggested at the beginning of this chapter, our model is designed such that it would be easy to switch to a more complex mechanism if required.

Forecasts types

All forecasts use the following long-range statistics:

$$\bar{\boldsymbol{r}}_t = (1 - g_L)\bar{\boldsymbol{r}}_{t-1} + g_L \boldsymbol{r}_t \tag{3.3.17a}$$

$$\bar{\boldsymbol{pd}}_t = (1 - g_L)\bar{\boldsymbol{pd}}_{t-1} + g_L \boldsymbol{pd}_t \tag{3.3.17b}$$

$$\bar{\boldsymbol{\sigma}}^2_{r,t} = (1 - g_L)\bar{\boldsymbol{\sigma}}^2_{r,t-1} + g_L(\boldsymbol{r}_t - \bar{\boldsymbol{r}}_t)^2 \tag{3.3.17c}$$

$$\bar{\boldsymbol{\sigma}}^2_{pd,t} = (1 - g_L)\bar{\boldsymbol{\sigma}}^2_{pd,t-1} + g_L(\boldsymbol{pd}_t - \bar{\boldsymbol{pd}}_t)^2 \tag{3.3.17d}$$

where $\bar{\boldsymbol{r}}$, $\bar{\boldsymbol{pd}}$, ,$\bar{\boldsymbol{\sigma}}^2_r$, $\bar{\boldsymbol{\sigma}}^2_{pd}$ correspond to mean return, price-dividends ratio and variances respectively. g_L is a gain parameter common across all agents.

Forecasts, $\boldsymbol{\tau}_t$, indexed by j are drawn from four well known linear forecast groups, these are:

- **Long-range forecast**: This forecast is a simple passive learning benchmark strategy that uses the long-range statistics given above. Agents using this forecast should generally employ a "buy and hold" strategy.

- **Adaptive linear forecast**: The second forecast rule uses the traditional adaptive linear expectations we described in Section 2.3.1, and in particular is of the same form as 2.3.1, it is given by

$$\boldsymbol{\tau}_t^j = \boldsymbol{\tau}_{t-1}^j + g_j(\boldsymbol{r}_t - \boldsymbol{\tau}_{t-1}^j) \tag{3.3.18}$$

  where g_j is a gain parameter which determines the degree to recent results influence expectations, it is crucial in determining the agents expectations as can be seen by the contour map in Figure 3.8 which shows how $\boldsymbol{\tau}_t^j$ responds to changes in g_j and $\boldsymbol{r}$. Here $\boldsymbol{\tau}_{t-1}$ is set to 0.45, we can see for low values of g_j expectations remain relatively unchanged regardless of the value of $\boldsymbol{r}_t$, but as g_j approaches 1 the $\boldsymbol{\tau}_t^j$ becomes more sensitive to deviation from $\boldsymbol{\tau}_{t-1}$. This passive learning rule has foundations in trend following, technical analysis and Kalman filtering. It therefore acts as our "chartist" strategy.

- **Price-dividends ratio regression forecast**: The third forecast rule uses linear regression of price-dividend ratios. It is given by

$$\boldsymbol{\tau}_t^j = \bar{\boldsymbol{r}}_t + \boldsymbol{\beta}_t^j(\boldsymbol{pd}_t - \bar{\boldsymbol{pd}}_t) \tag{3.3.19}$$

  where β_t^j is the regression parameter that must be learned. This rule acts as our "fundamental" strategy.

- **Lagged return regression forecast**: The fourth, and final rule uses a modified lagged linear regression that attempts to eliminate short range autocorrelations, it is given by

$$\boldsymbol{\tau}_t^j = \bar{\boldsymbol{r}}_t + \sum_{i=0}^{M_{ar}-1} \boldsymbol{\beta}_{t,i}^j(\boldsymbol{r}_{t-i} - \bar{\boldsymbol{r}}_t) \tag{3.3.20}$$

  where M_{ar} is a lag parameter set to a low value and β_t^j is the regression parameter which must be learned. This rule is somewhat harder to classify, it lies somewhere between the two previous strategies.

Regression Updates

We now specify the agents active learning mechanism employed to update the regression parameters, $\boldsymbol{\beta}_t^j$, of the two regression forecast rules. Both rules use a recursive least-squares method which is widely regarded as a very effective learning mechanism. Note we make a slight abuse of notation here by allowing vector division, which should be read in index form.

- **Fundamental regression update**

$$\boldsymbol{\beta}_{t+1}^j = \boldsymbol{\beta}_t^j + \frac{g_j}{\bar{\boldsymbol{\sigma}}_{pd,t}^2}(\boldsymbol{pd}_{t-i}\boldsymbol{u}_{t,j}) \tag{3.3.21a}$$

$$\boldsymbol{u}_{t,j} = \tilde{\boldsymbol{r}}_t - \boldsymbol{\tau}_{j,t-M_{pd}} \tag{3.3.21b}$$

$$\tilde{\boldsymbol{r}}_t = \frac{1}{M_{pd}} \sum_{j=1}^{M_{pd}} \boldsymbol{r}_{t-j+i} \tag{3.3.21c}$$

- **Lagged return regression update**

$$\boldsymbol{\beta}_{t+1}^j = \boldsymbol{\beta}_t^j + \frac{g_j}{\bar{\boldsymbol{\sigma}}_{r,t}^2}(\boldsymbol{r}_{t-i}\boldsymbol{u}_{t,j}) \tag{3.3.22a}$$

$$\boldsymbol{u}_{t,j} = \boldsymbol{r}_t - \boldsymbol{\tau}_t^j \tag{3.3.22b}$$

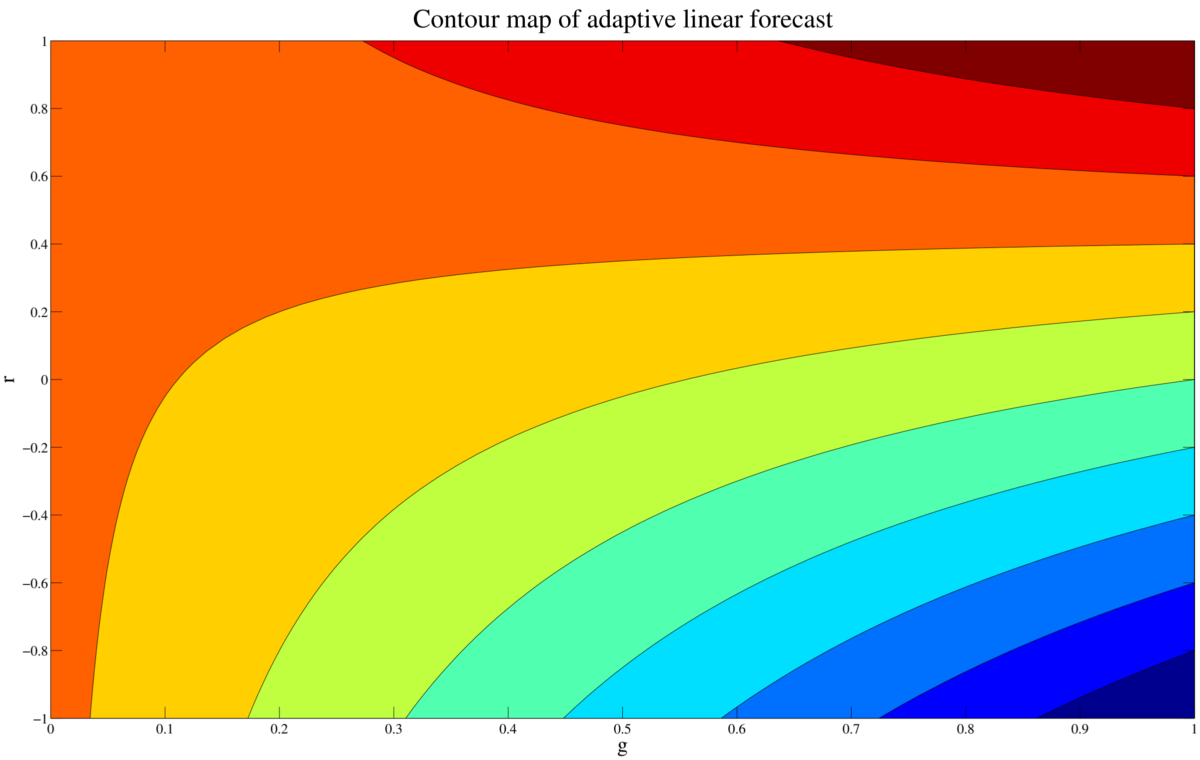

Figure 3.8: Contour map of adaptive linear forecast for varying gain g and current return r_t.

Forecast Utility Updates

The active part of the learning process involves agents having to select which forecast to use. This could be done in a variety of ways but the simplest is select the rule with maximum current utility. This approach is conducive to the agents ultimate goal of maximising its current wealth. Each forecast, τ_t^j, is given a utility that is updated according to the adaptive linear rule

$$\hat{U}_{t,j} = \hat{U}_{t-1,j} + g_u^i(U_{t,j} - \hat{U}_{t-1,j}) \tag{3.3.23}$$

where g_u is a gain parameter, note this is exactly the same rule as the adaptive linear forecast rule, thus g_u is equally important (see Figure 3.8).

Agents use the same utility function for next period portfolio returns as they did for next period wealth, namely the power utility,

$$U_{t,j} = \frac{1}{1-\gamma}(1 + r_{P,t,j})^{1-\gamma} \tag{3.3.24}$$

This concludes our agent definition. Throughout our design we have tried to insure our agents are complicated enough to allow for a good range of market dynamics, and to make the general equilibrium problem interesting but simple enough to be able to analyse effectively. Our learning approach may seem reasonably basic given the range of possibilities, but given our main priority is to understand equilibrium and market dynamics in boundedly rational world, a task that as far as we are aware has never truly been faced, we think it is reasonable to start at this level of simplicity. Once we have gathered our results we should then be in a position judge if more complicated expectation and learning strategies could be employed.

3.4 Social Interactions

We discussed social interaction in Section 2.3.2 and found two possible methods for introducing agent interactions:

- **Fixed networks**: Agents occupy nodes in a fixed network and can only trade with agents they are adjacent to.
- **Intelligent interactions**: Agents learn who the wish to trade and exchange information with in order to maximise mutual utility.

Both of these introduce realisms which, as we discussed, influence market structure and resulting dynamics. The question is should we introduce either of these mechanisms into our model? The first concern we raise is that in our proposed general equilibrium framework the fixed network protocol is a challenging design task. The problem is illustrated in Figure 3.9; in our general equilibrium framework we must have a centralised exchange for price clearing. This is perhaps one of the main arguments against using equilibrium for price finding. To explain the problem, suppose we have a network protocol like the one in Figure 3.9. Here agent 3 can trade with everyone but agents 1, 2, 4 & 5 can only trade with agent 3. The problem is there is no way to form demands for agents 1, 2, 4 & 5 that conforms to our model structure given in the previous section; agents simply receive prices and return a demand vector. Perhaps one possibility would be to introduce a demand 'matrix' that explicitly encodes demands for assets for each trader, but this raises very challenging questions such as how agents should 'split' their demands between rival agents.

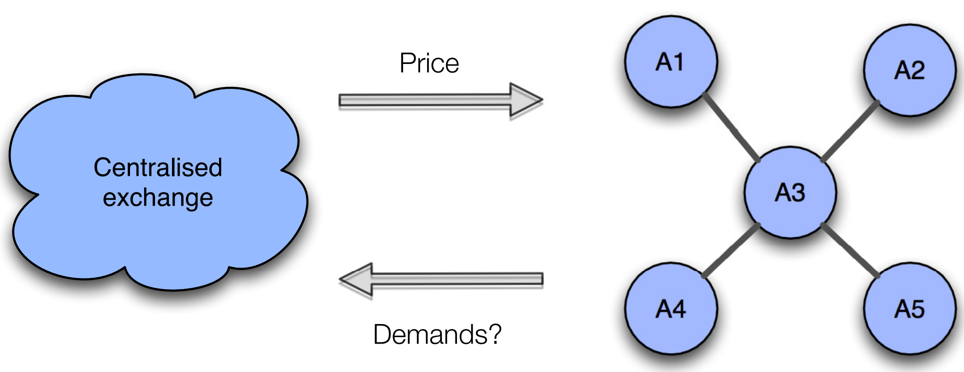

Figure 3.9: The fixed network problem.

We conclude that while the fixed network protocol raises many interesting questions, it is not feasible given the complexity of our task, we defer this to a future research task. Our suggestion would be this protocol should be used with the decentralised "order-book" style pricing protocol discussed in Section 2.3.2.

Having just described why a centralised framework can not support complex demand structures easily, we note that the intelligent interactions protocol is challenging for exactly the same reason. Thus, unfortunately, social interactions can not form part of our artificial market. While this is a little disappointing, we note that little overall progress has been made in studying social interactions in artificial markets, thus we are not alone in finding it difficult to implement such protocols.

3.5 Market Operation

Other than the pricing mechanism, all of our market components are now in place. In this section we summarise market operation by giving a preliminary version of the main market loop, shown in Algorithm 1.

Algorithm 1: Main market loop

Data: Number of periods n, Dividends $\boldsymbol{\pi}$, Agents $\boldsymbol{\alpha}$
Result: Market simulation
Warm up market;
for $t \leftarrow 1$ **to** n **do**
  for $i \in \alpha$ **do**
    $v_{i,t,j} \leftarrow \frac{w_{i,j} W_{i,t-1}}{p_{t-1,j}}$;
    $B_{i,t} \leftarrow W_{i,t-1} + \sigma(\boldsymbol{\pi}_{i,t}) \cdot \boldsymbol{v}_{i,t}$;
  end
  `UpdateLongRangeStatistics(`$\boldsymbol{\pi}_t, \boldsymbol{p}_{t-1}$`);`
  $\boldsymbol{p}_t \leftarrow$ `FindEquilibriumPrices();`
  `UpdateMarket(`$\boldsymbol{p}_t$`);`
  for $i \in \alpha$ **do**
    $W_{i,t} \leftarrow (1 + r^P_{i,t})(1 - \lambda_i) W_{i,t-1}$;
    for $\tau^j \in i$.`Forecasts()` **do**
      τ^j.`Update` ();
      $r^{P_{\tau^j}}_{i,t} \leftarrow \tau^j$.`ExpectedPortfolioReturns();`
      $U_{i,f,t} \leftarrow \frac{1}{1-\gamma_j}(1 + r^{P_{\tau^j}}_{i,t})^{\gamma_i}$;
    end
    $\tau^j_{i,t} \leftarrow$ `MaxUtility(`τ`);`
  end
end

Before market simulation starts we 'warm-up' the market with data generated from our autoregression process. This allows the agents to diverge from their initial conditions, and in particular, allows them to settle on forecasts. Market operation then proceeds as described in the previous sections. First agents update their current period budget constraints using the current period dividend. Then the long range statistics are updated, these are used by agents when calculating their demands during the market clearing phase. The market is then updated with the equilibrium prices and finally agents update their forecasts and actively search for a better forecast. This process is repeated.

The last component of our model to be determined is of course the equilibrium pricing mechanism that clears the market. This will be described in the next section. Although it is worth mentioning now, as we envisioned at the beginning of this chapter (See Figure 3.1), the pricing mechanism is indeed independent of general market operation.

3.6 Calibration

To calibrate our geometric vector autoregression dividends model we observed empirical market data from some well known Dow-Jones index stocks.[4] While this is a good starting platform, there are some important limitations to be aware of. Firstly, it must be noted that not all stocks pay dividend; this is especially apparent for growth companies and in particular technology firms. The abstractions our model makes does not account for growth opportunity, it is therefore vital for market function that assets generate a stable dividend stream. Secondly, companies do not generally pay dividend in uniform, thus statistical analysis on dividend data must be considered to be erroneous due to differing market conditions at the time of respective dividend payments. Lastly, it is important to note there is not sufficient historical data to make certain assertions about the distributions of dividend returns, although, as we have asserted in Section 3.2 we believe there is enough evidence to suggest dividends are log-normally distributed.

In order to calibrate our model we found $\boldsymbol{c}$, Φ and Σ such that the mean 3.2.2 and covariance 3.2.3 evaluated most closely to the sample mean and covariance, $\hat{\boldsymbol{\mu}}$ and $\hat{\boldsymbol{\Sigma}}$. To achieve this we devised the following optimisation problem:

$$\begin{aligned}
\underset{\boldsymbol{c},\,\Phi,\,\Sigma}{\text{minimise}} \quad & |\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}|_2^2 + |\Gamma - \hat{\Sigma}|_2^2 \\
\text{subject to} \quad & \mu_i = \exp\left[\left((\boldsymbol{\iota} - \Phi)^{-1}\boldsymbol{c}\right)_i + \frac{1}{2}\mathbb{D}\left(\sum_{k=0}^{\infty} \Phi^k \Sigma \Phi^{T^k}\right)_i\right] \\
& \gamma_{ij} = (e^{\sigma_{ij}} - 1)\left(e^{\hat{\mu}_{ii} + \hat{\mu}_{jj} + \frac{\sigma_{ii} + \sigma_{jj}}{2}}\right)
\end{aligned}$$

Unfortunately the solution does not guarantee Σ is positive semi-definite as required, thus we use the nearest possible semi-definite matrix.[5] We also added a little random variation to avoid stagnation. While we concede there are probably better methods for solving this problem, unfortunately we have been unable to find any published work on the subject.

3.7 Summary

During this chapter we have built the framework of our artificial stock market. We designed our market in such a way that we could apply and centralised price clearing mechanism at a latter stage. We first gave a general overview of our model and stressed our market would be made of independent components.

We then discussed a model for **dividends generation** assuming our hypothesise that dividends are log-normally distributed. We found the best option was to employ a **geometric vector autoregressive model**.

The bulk of the chapter focused on the agents within our model. First, we derived a **portfolio selection** policy using **power utility maximisation**, we showed how this was given as the solution to a quadratic program in the case of no short selling, and a linear equation when short sales are allowed. We followed with a derivation of the agent **demand** function. Next, we introduced our agents **expected future returns forecasts** using both passive and active learning.

We then explained why **social interaction** protocols are not applicable in our centralised model. Finally we gave an algorithm showing explicitly how our market operates and then showed how we calibrated our geometric vector autoregressive model.

[4] Stocks: BA, BP, IBM, KO, PFE, T, WMT, XOM. Bonds: 3-month TBILL. Data obtained from http://ichart.finance.yahoo.com and http://www.federalreserve.gov

[5] To find the nearest semi-positive definite matrix we used a Matlab function written by Bunchman (2011).

4

Computing General Equilibrium Prices

In this chapter we present an optimisation problem for finding equilibrium prices in our artificial stock market model. As we have discussed throughout this thesis, finding equilibria outside a Walrasian framework poses a challenging problem, indeed, to the best of our knowledge this represents the first artificial stock market model to implement a general equilibrium pricing mechanism.

4.1 Comparison with a Walrasian Framework

We begin by comparing our model with a Walrasian pure exchange economy, and in particular we consider if our model satisfies the conditions for equilibria existence; many of the existence conditions are required for search method convergence. Following this we evaluate if any existing methods for computing equilibria could be employed to solve our problem.

4.1.1 Walras' Law

Recall from Section 2.2 that the most fundamental property of agent demand functions in a Walrasian market is satisfaction of *Walras' law*. We presented two equivalent versions of Walras' law, 2.2.1 and 2.2.5. We begin by verifying our agents demand functions satisfy Walras' law.

Proposition 4. *The agent demand function, ζ, given in 3.3.16 satisfies Walras' law.*

Proof. The first observation is the agents optimal portfolio vector, $\boldsymbol{w}^*$, is always equal to unity independent of the agents forecast return. This is because agents solve 3.3.7 which has constraint $\boldsymbol{w} \cdot \boldsymbol{\iota} = 1$. The second observation is agents expend all current wealth; in a Walrasian setting agents do not receive dividends, thus their wealth is given by W_t, however our agents receive dividends at the beginning of each period thus we must include this into our demand function. It remains to show

$$\zeta(\boldsymbol{p}, \boldsymbol{p} \cdot \boldsymbol{\nu}) - \boldsymbol{\nu} = 0$$

To see this note agent ownership, $\boldsymbol{\nu}_{i,t,j}$, is given by

$$\boldsymbol{\nu}_{i,t,j} = \frac{\boldsymbol{w}_{i,j} W_{i,t-1}}{\boldsymbol{p}_{t-1,j}}$$

thus with a slight abuse of notation by assuming division is in index terms

$$
\begin{aligned}
\zeta(\boldsymbol{p}, \boldsymbol{p}\cdot\boldsymbol{\nu}) - \boldsymbol{\nu} &= \frac{\boldsymbol{w}^* \boldsymbol{p} \cdot \frac{\boldsymbol{w} W_{t-1}}{\boldsymbol{p}}}{\boldsymbol{p}} - \frac{\boldsymbol{w} W_{t-1}}{\boldsymbol{p}} \\
&= \frac{W_{t-1}}{\boldsymbol{p}} - \frac{\boldsymbol{w} W_{t-1}}{\boldsymbol{p}} \\
&= (\boldsymbol{\iota} - \boldsymbol{w})\frac{W_{t-1}}{\boldsymbol{p}} = 0
\end{aligned}
$$

Thus ζ does indeed satisfy Walras' law. □

4.1.2 Homogeneity of ζ

The next important question is do our agents satisfy the Arrow-Debreu conditions 2.2.2? The first three conditions are satisfied by our choice of utility function, and the fourth is easily satisfied; we just need to make sure agents each have a little of each asset initially. Furthermore, we just proved the fifth condition, Walras' law. This leaves the final condition, *zero degree homogeneity of* ζ. More formally, does $\zeta(\boldsymbol{p}, W) \equiv \zeta(\alpha\boldsymbol{p}, \alpha W)$? This says that if we increase prices and wealth by the same constant factor then demand remains the same. In a Walrasian setting this seems perfectly valid; an agents wealth is determined by its current asset bundle so a constant increase in the price of all assets has no effect on demand. This is really a matter of *relative* price; the absolute price of an asset is not of importance, the *relative price* of an asset is what matters, prices are set relative to the *numeraire*.

However in our market demand is not a function of price and wealth, it is a function of price only, wealth is determined by the last periods prices and the current dividend. Furthermore, the future return forecast uses the current market prices and current dividends to calculate expected returns for the next period, but these prices and dividends are normalised in terms of *last* periods prices and dividend, in essence the current price is not scaled by the new price and thus acts as a constant. To see the effect of this see Figure 4.1, this shows demand for a riskier asset in a two asset economy with a benchmark long-range forecast. We can see that when both values are scaled by the same constant factor (points on the diagonal), demand for the riskier asset increases to begin with, and then remains constant. The reason demands are constant beyond this point is the agent has moved all it's wealth into the riskier asset as it's expectation of future returns grew. This is just a simple case, other forecasts may result in different dynamics. The key outcome however is that our agents demand function, ζ, is not homogenous to degree zero. Thus, according to the Theorem 1, our market does not meet the criteria for guaranteed general equilibrium existence.

4.2 Approximate Equilibria via Optimisation

The preceding revelation raises the question of whether equilibrium, or at-least approximate equilibrium, exist in our artificial market. We have shown our market to satisfy all but one of the criteria required, and we have also seen our agents demand functions are homogenous to degree zero past a certain limit. It will be the subject of the next chapter to try and computationally determine whether equilibrium exists in our market. For the remainder of this chapter we assume approximate equilibria do exist and work towards an optimisation algorithm for computing them.

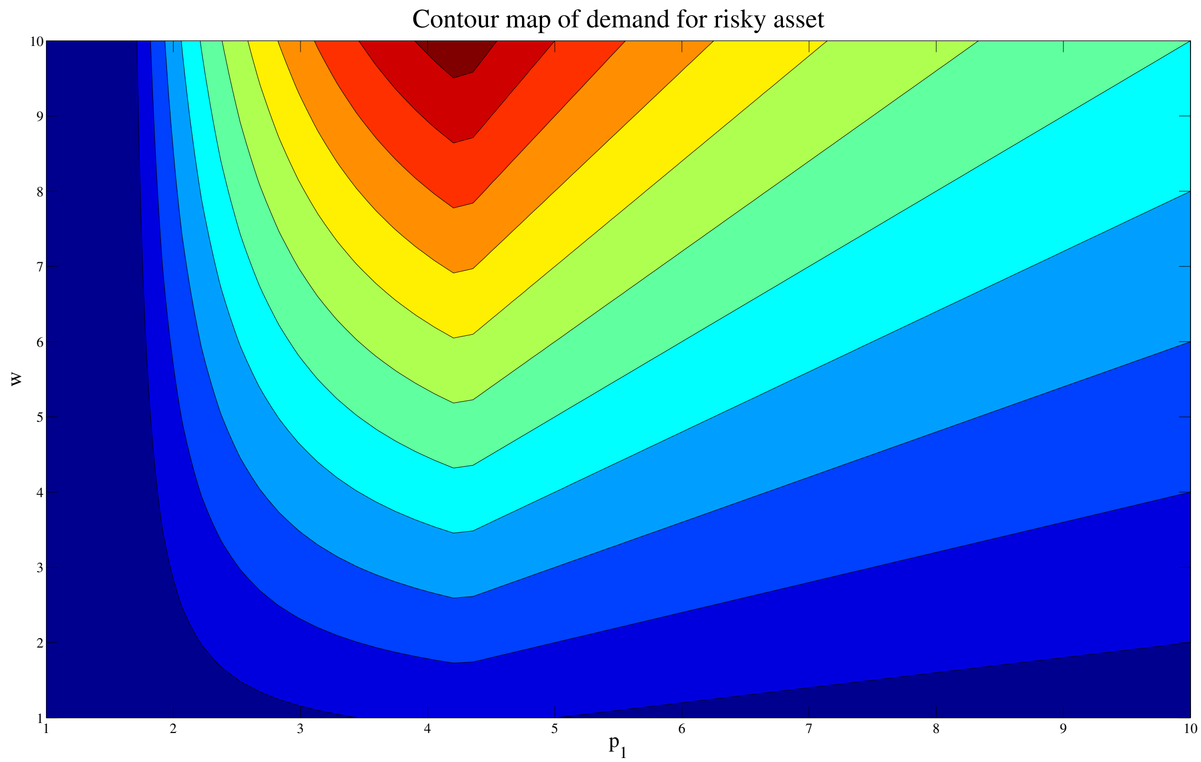

Figure 4.1: Contour map showing demand for riskier asset in a two asset economy.

4.2.1 Gross Substitutes

We briefly discussed methods for computing approximate equilibria in Section 2.2.3 and found many required the assumption of *gross substitutes* or at-least *weak gross substitutes*. Recall the gross substitutes property requires that a price *increase* in one asset leads to a demand *increase* in all other assets; weak gross substitutes only requires that no price decreases occur. If we are to be able to use any of the algorithms presented in Section 2.2.3 we need to verify our market satisfies at-least weak gross substitutes.

To evaluate if our market satisfies gross substitutes we set up a simple market with two assets with differing risk characteristics. We then held the safer asset fixed at price 1 and varied the price of the other asset, the current asset prices are set to 3 and 1 for the riskier and safer assets respectively, dividends is set to zero for both assets and the agents risk aversion was set quite low. The agent was set to use the long-range forecast with a fixed mean and the gain parameter was set as a variable. The resulting contour map is shown in Figure 4.2. We can see that other than for very low values of g_L, demand for the safer asset changes dependent on the new price of the riskier asset. The reason for this is as g_L approaches 0 the agent ignores any new information; its forecast becomes homogenous. Given this, and the fact that agent forecasts are set randomly, we can not conclude that our market satisfies gross substitutes (or even weak gross substitutes).

4.2.2 Convexity of z

We can now see why the homogeneity and gross substitutes properties of ζ are important for proving existence, and efficient computing equilibria. Prices can be adjusted with full knowledge of the possible domain of outcomes; the excess function z is *convex*. Convex optimisation problems can in general be solved very

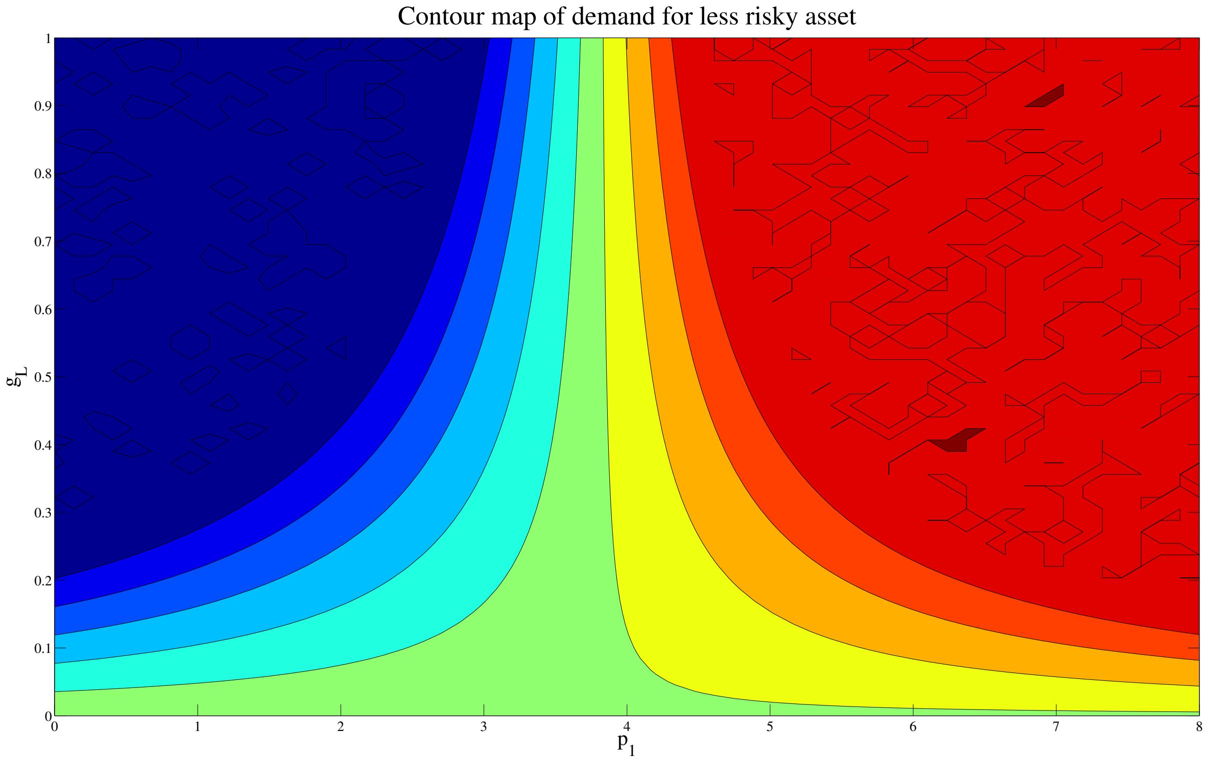

Figure 4.2: Contour map showing demand for safer asset in a two asset economy using long range forecasts with gain parameter g_L. The price of the risky asset is p_1.

efficiently due to well known differentiability conditions. Our market does not satisfy either of these two conditions sufficient for convexity of z.

As can be explicitly seen in Figure 4.3 which shows excess for the riskier asset in our simple market described earlier, the excess function z in our market is non-convex. Aggregate demand has phases of 'dips' and 'humps' as expectations change. We conclude that to find an equilibrium price we must form a non-convex solution method. To the best of our knowledge, this type of problem has not been addressed by either applied general equilibrium or agent-based computing.

4.2.3 Differentiability of z

The next property we must consider is differentiability of ζ, and hence of z. Although our problem is non-convex we may still be able compute local optimums in reasonable time using Newton methods such as SQP if we can calculate the Jacobian and Hessian matrix of z. Unfortunately this is not the case, for two reasons:

- **Agent portfolio selection**: If we disallow short selling agent preferences are found using a maximisation problem. This is not differentiable analytically.
- **Agent forecasts**: Agents use forecasts in their demand functions that is only known at run-time and hence cannot be differentiated at compile time.

We conclude that our algorithm must either use a numerically computed Hessian approximation, or a differential free method.

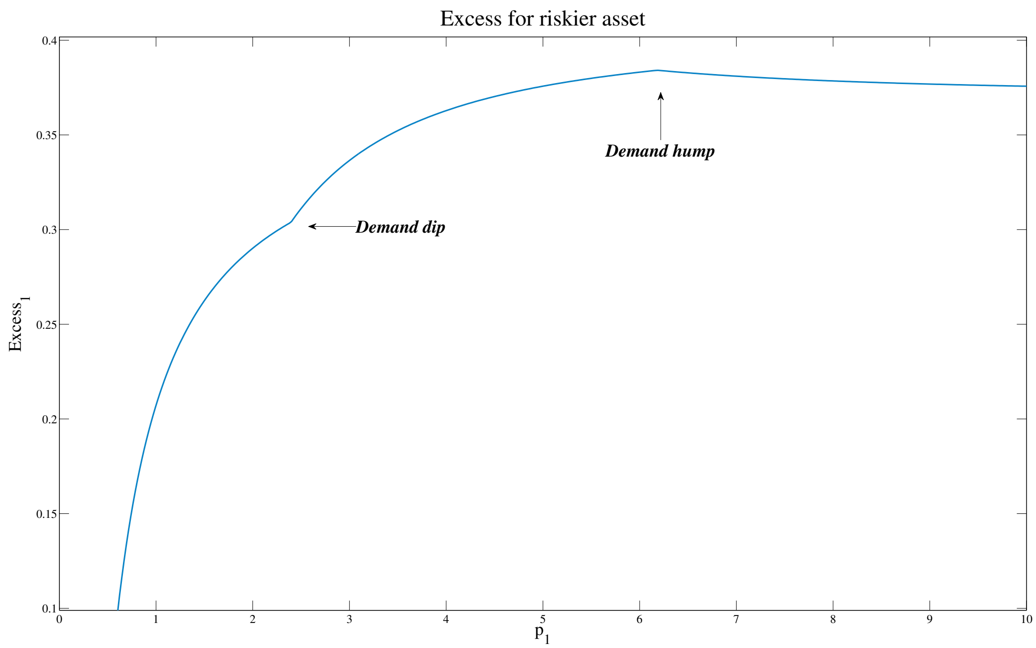

Figure 4.3: Excess of riskier asset in simple two-asset market.

4.2.4 Problem Constraints

The final component of our problem we must consider is *constraints*. Unconstrained optimisation problems are much easier to solve than constrained equivalents, indeed, as we have seen they can sometimes be solved analytically. If our problem was unconstrained we could apply either a Trust-Region Line Search, or a Conjugate Gradient algorithm. Unfortunately this is not the case, we must limit prices to being non-negative, thus we must use a constrained optimisation algorithm.

Normalisation of z

To compute equilibrium prices our algorithm must find a price vector that satisfies the equilibrium condition 2.2.6. An important observation is z can be either positive or negative, this is the case as it is the aggregation of individual asset excesses that can themselves be positive or negative. As we are using optimisation techniques for computing equilibrium prices we must constrain the value of z (the objective function). The simplest method of achieving this is to normalise z using the Euclidean norm. Note this does not effect the equilibrium condition as $z(\boldsymbol{p}) = 0$ iff $\|z(\boldsymbol{p})\|_2^2 = 0$.

4.2.5 Model Modifications

There is one further modification we must make to our original model, wealth growth must be constrained. In contrast to a Walrasian model our market has exogenously included wealth; relative wealth *increases* over time. Furthermore, if approximate equilibria are computed there will be some excess ownership due to the market being over-cleared. Without any intervention wealth levels would exponentially increase leading to exponential price increases. This is the reason we introduced a *consumption*, λ_i, for each agent.

4.3 Solution Methods

As discussed in the last section to compute approximate equilibrium prices we must solve a non-linear, non-convex constrained optimisation problem. More formally we must solve the following problem:

$$\begin{aligned} \underset{\boldsymbol{p}_t}{\text{minimise}} \quad & \|z(\boldsymbol{p}_t)\|_2^2 \\ \text{subject to} \quad & \boldsymbol{p}_t > 0 \end{aligned} \tag{4.3.1}$$

In addition we must either use a derivative free, or quasi-Newton with Hessian approximation method. These types of problem are not easy to solve and in general solution methods are only guaranteed to find local optimums. Nevertheless, we can take some solace in knowing our problem is not highly-nonconvex, indeed, with a sensible initial value we believe approximate equilibrium prices can be found in reasonable time.

We now briefly comment each of the three solution methods we attempted and then finish the section by presenting error testing results for each method. We then summarise the findings and select an algorithm to proceed with.

4.3.1 Nelder-Mead Simplex

The first algorithm we tested was an adaption of the well known derivative free Nelder-Mead simplex method developed by Nelder and Mead (1965). Although Nelder-Mead is an unconstrained method, we reasoned that as it is derivative free we might be able to add a linear constraint by including an ad hoc barrier into the objective function. Abstractly, the algorithm works by forming an $n + 1$ point simplex in $\mathbb{R}^n$. On each iteration the vertex with the worst objective value is chosen and the simplex is either reflected, expanded or contracted along the the line joining the selected vertex and the centroid of the remaining vertices. If no better point can be found then only the vertex with the best objective value is retained and all other vertices are shrunk towards that vertex. The process then repeats.

The workings of the Nelder-Mead algorithm seem to suggest our barrier method might be feasible, especially for for simpler search spaces such as those with fewer agents and assets. We reason that if the algorithm selects a point within our barrier then that point will just be ignored, so assuming we set the barrier large enough and start the search in a point outside the barrier the algorithm should work as expected.

4.3.2 COBYLA

The second algorithm we tested was the COBYLA constrained optimiser developed by Powell (2007). The method works by forming linear linear approximations to the objective and constraint functions by forming linear interpolations at $n + 1$ points in the search space. A series of simplex are then formed using the interpolations points which is then reduced until a vector of points with sufficiently optimal objective value is achieved.

The COBYLA method represents a direct competitor to our Nelder-Mead adaption; it will be interesting to see if a simplex method designed with constraints in mind can do better than an adapted unconstrained simplex method. Unfortunately we could not find much material explaining the full workings, or performance of COBYLA.

4.3.3 Limited-memory BFGS-B

The final algorithm we tested was the popular quasi-Newton L-BFGS-B method developed by Byrd et al. (1995). This is not a derivative-free method, but only requires Jacobian information to update a Hessian approximation on each iteration, however the Jacobian can itself be approximated using the techniques detailed by Nocedal and Wright (1999). To avoid memory cost problems we used the limited-memory verison of BFGS-B. This version of BFGS-B only stores an approximation to the computed Hessian matrix, thus saving having to store a dense n^2 matrix.

The performance of L-BFGS-B is considered one of the best of its kind. Given our problem is relatively simple compared to many non-linear problems we feel this algorithm is most likely to be the most successful. However we do have concerns that the Jacobian approximation have a critical impact on it's performance.

4.4 Results

In this section we present the results of our initial equilibrium algorithm experiments. These tests are not designed to evaluate equilibrium dynamics, but to test the performance and relative accuracies achieved by each algorithm. This will allow us to choose the most appropriate algorithm moving forward. We then end with a brief discussion on the general run-time characteristics of each algorithm.

4.4.1 Approximation Errors

To test the accuracy of the algorithms we ran independent simulations of our market for varying amounts of agents and assets. These tests were kept very general and we stress the errors seen here might not be a good representation of equilibrium approximations, however the results will be useful in determining the relative performance of each algorithm. We kept the test sizes reasonably small as the tests take a while to run and we feel the number of test cases we have considered should be adequate to extrapolate for larger setups. Supplies for each asset where set to 1. Throughout the tests we kept the number of periods fixed at 50 which should be more than adequate to get good sample base. We ran 10 separate tests for each test setup and took the 5 worst from each to calculate the mean normalised error, $\bar{\mathcal{E}} = ||z||_2^2$, of each algorithm. The results are shown below in Table 4.1:

Table 4.1: Mean normalised equilibrium approximation errors.

No. of agents	No. of assets	Nelder-Mead $\bar{\mathcal{E}}$	COBYLA $\bar{\mathcal{E}}$	L-BFGS-B $\bar{\mathcal{E}}$
5	2	0.4786	0.4159	0.3756
5	3	∞	0.3398	0.3734
5	5	∞	0.5328	0.4139
10	2	0.5231	0.3921	0.3572
10	3	∞	0.4995	0.3904
10	5	∞	0.5577	0.4933
30	2	0.5423	0.4380	0.3333
30	3	∞	0.4322	0.4478
30	5	∞	0.5941	0.5010
50	2	0.6059	0.4792	0.5102
50	3	∞	0.5053	0.4329
50	5	∞	0.6012	0.5398

The first and most obvious of the results seen are the ∞ mean norms given by the Nelder-Mead algorithm. While this result is worrying, it is perhaps not as bad as first appears; further analysis showed the algorithms performance to vary drastically, in some cases the algorithm performed very well (sometimes better than the other two), but on other occasions the algorithm diverged dramatically and terminated only when reaching its objective function valuation limit. We could find no definitive reason why this was the case, but we suspect it might be when the prices are close to zero; this would almost surly be a result of the naive adaption of the algorithm for purposes it was clearly not designed.

Our results show COBYLA and L-BFGS-B fair reasonably similarly for fewer assets, neither seem to be particularly effected by the number of agents in the market, but a larger number of assets seems to have an impact on the accuracy of both. However, these results do seem to indicate L-BFGS-B scales slightly better than COBYLA. This may be because of the increases non-linearity of the ζ as more assets are added to the market; the linear approximations used by COBYLA may not be as accurate as non-linearity increases.

More generally, the results given thus far do seem a little worrying. A mean approximation error percentage of $> 30\%$ may not seem sufficient. While this will be analysed in more detail in the next chapter we show some initial analysis now to give some preliminary insight into the cause of these high figures. On further analysis we found that like for Nelder-Mead, while the norms where in general much closer to zero than the figures presented here (sometimes reaching as low as 1×10^{-10}), anomalies such as those seen in Figure 4.4 were somewhat frequent, and distorted the norm excess means. These types of disturbances seemed to effect all algorithms (though the Nelder-Mead disturbances were of a much greater magnitude), in the next chapter we will try and determine if this is a consequence of our model, or the solvers we have chosen to use.

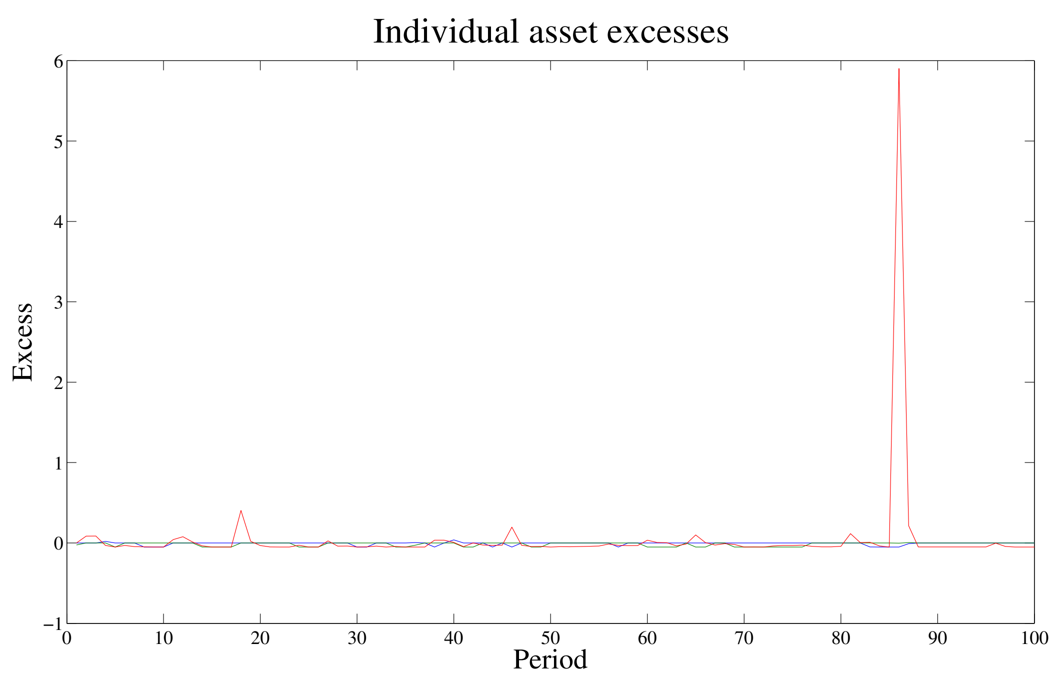

Figure 4.4: Individual asset excesses for a test simulation.

4.4.2 Run-time and Memory Costs

Calculating an estimated run-time and memory costs is important as it gives us a practical limit for our market setup. While we could not derive analytical performance measures we were able to gain a reasonable understanding of the systems performance by analysing the main bottlenecks in the system, these are:

- **Portfolio selection**: Agent portfolio selection under no short selling requires solving a quadratic program. Quadratic programs can be solved in polynomial time for positive-definite Q-matrices. Our Q-matrix will be at-least semi-positive definite, but there is no guarantee it will be positive-definite, also, for large number of assets the problem becomes increasingly harder to solve, especially as our covariance matrix in unlikely to be sparse. Furthermore, we do not know how many times agents will need to calculate their optimal portfolios each period, this is dependent on the convergence rate of the equilibrium price finding optimisation routine. Thus the number of assets will be the first main variable in determining run-time and memory costs.
- **Price finding**: As we have discussed in this chapter the equilibrium price finding algorithm is non-linear, non-convex and constrained. Our problem is not particularly non-convex for small numbers of agents and assets but we can not assume this is the case for larger markets. Thus it is highly likely that the number of agents will also be a limiting factor in run-time and memory performance.

We were also able to capture some more solid figures for run-time and memory usage while running simulations on the departmental servers. The run-time and memory usage seemed to be quite stable among similar setups, Figures 4.5 and 4.6 shows some results we obtained from various simulations. We can see run-time increases linearly with both agents and assets. Figures 4.7 and 4.8 show memory usage over 10 periods for varying numbers of assets and agents. We can see these too, are both linear.

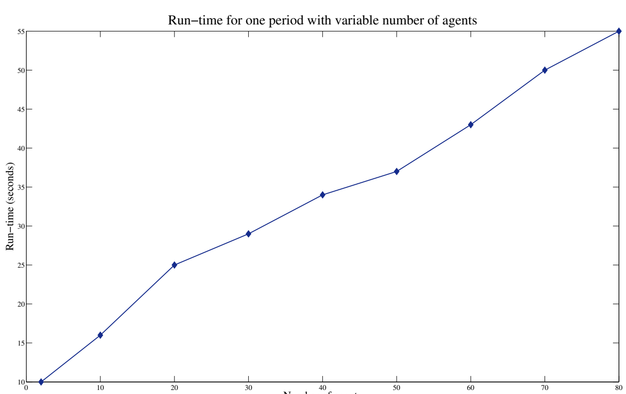

Figure 4.5: Run-time (in seconds) for one period with varying number of agents.

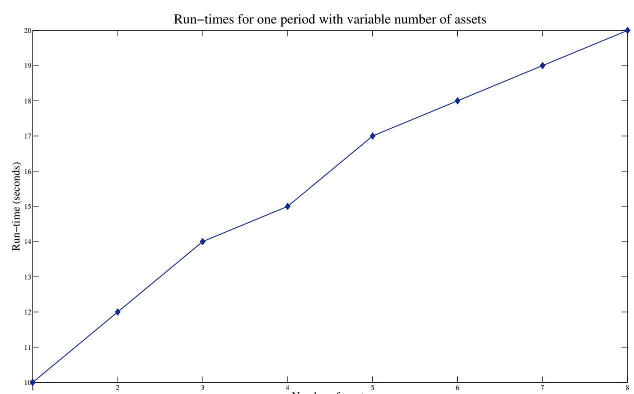

Figure 4.6: Run-time (in seconds) for one period with varying number of assets.

Although our memory usage tests show only linear expansion, this is still quite a concern. Even over just 10 periods memory usage rose rapidly with increased number of agents, this may partly be due to the number forecast histories each agent must store. Given the results presented here, and our experience working with our simulation,[1] we estimate the maximum number of agents and assets over 100 periods we can simulate in reasonable time is 60 agents and 8 assets.

[1] We ran simulations on a machine with 24GB of memory available and found with figures over these amounts the entire memory was quickly consumed and virtual memory initiated, once this occurs the simulation becomes painfully slow.

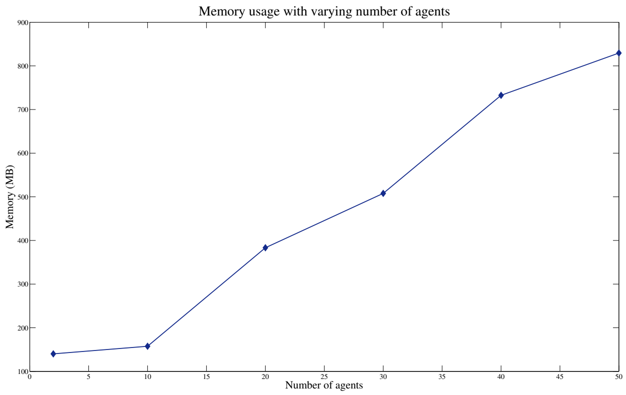

Figure 4.7: Memory usage (in MB) over 10 periods with varying number of agents.

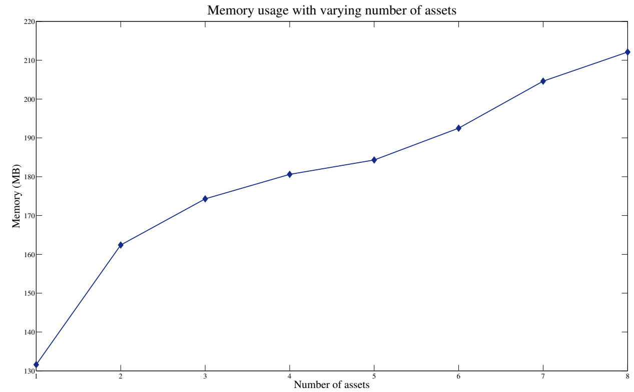

Figure 4.8: Memory usage (in MB) over 10 periods with varying number of assets.

4.5 Summary

In this section we designed an algorithm to **compute general equilibria prices** for our artificial market. We showed that our market satisfies all but one of the conditions sufficient for equilibria existence, the agent demand function ζ is **not homogeneous to degree zero**. We found this was a result of the dividends paid in terms of prices normalised in the last period, and agents forecasts.

Next we showed that our market does not satisfy the **gross substitutes** property required by almost all currently available algorithms for computing general equilibria. We followed by showing the excess demand function z is **not convex** and **non-differentiable** due to agent portfolio selection and forecast structure. We conclude general equilibrium prices in our market must be solved by a **non-linear, non-convex constrained minimisation problem**.

We then presented three algorithms for computing the price optimisation problem: **Nelder-Mead**; **COBYLA**; and **L-BFGS-B**. We explained why these methods are applicable and then showed performance analysis results for each algorithm. We found that the Nelder-Mead method was not acceptable due to instability of convergence, and while the other two methods gave better overall results, they too showed signs of **unstable convergence**.

Finally we found two potential bottlenecks in our market and showed run-time performance figures from simulations of our market using the L-BFGS-B method under differing initial conditions. We concluded that while the running time figures were not overly concerning, memory requirements severely limited the possible size of our market.

5

On Equilibria in Artificial Stock Markets

In this chapter we focus on the dynamics of equilibria in our market using the results from the last chapter. In particular we analyse the existence of equilibria from a computational perspective and evaluate more generally if general equilibrium pricing is a appropriate for artificial stock markets.

5.1 Existence

While it was never the goal of this thesis to find a mathematical argument for existence of general equilibria in our artificial stock market, indeed, doing so would be a challenge for even the most accomplished mathematician, we would like to try and determine computationally if equilibria existence is likely. Before we begin our analysis it is worth commenting on the consequences of equilibrium *not* being certain. While it would seem counter-intuitive to apply general equilibrium pricing to a market that does not guarantee existence, we stress that even applied general equilibrium models, those that apply empirical data to a Walrasian framework, are only certain to find approximate equilibrium. Furthermore, we make note that some artificial markets do not use an equilibrium pricing method at-all, recall from Section2.3.2 that some artificial markets are purposefully designed to be out of equilibrium. In reality, it is unlikely that while globally optimal solutions non-linear, non-convex optimisation problems remain NP-Hard to compute, non-approximate general equilibrium pricing will be applicable in artificial markets. Nevertheless, we believe that an approximate equilibrium market would still be an extremely useful research tool in working towards the goals of agent-based computing.

5.1.1 Rational Expectations Equilibrium

First we briefly evaluated equilibria existence under rational expectations assuming agents all use long-range linear forecasts, we gave the agents knowledge of the dividends structure (expected value and mean) and all risk-aversions were kept at the same level to enforce homogeneity. For this part of the evaluation we only used the L-BFGS-B algorithm.

Under rational expectations we found the equilibrium points were in general much easier to compute. Indeed, when in rational expectations equilibrium the agent demand function becomes convex, this is because the previous periods price and dividends does not form part of the conditional information set used to form the new expectations (see 2.3.2). An example of a rational expectations excess function is shown in Figure 5.1. This market consists of three assets of varying risk, we set one of the prices fixed to just above the optimum to give a better picture of the search space. Notice how the there is a clear zone of prices that form a convex basin, the equilibrium price lies within this basin and can be quickly found by the L-BFGS-B method.

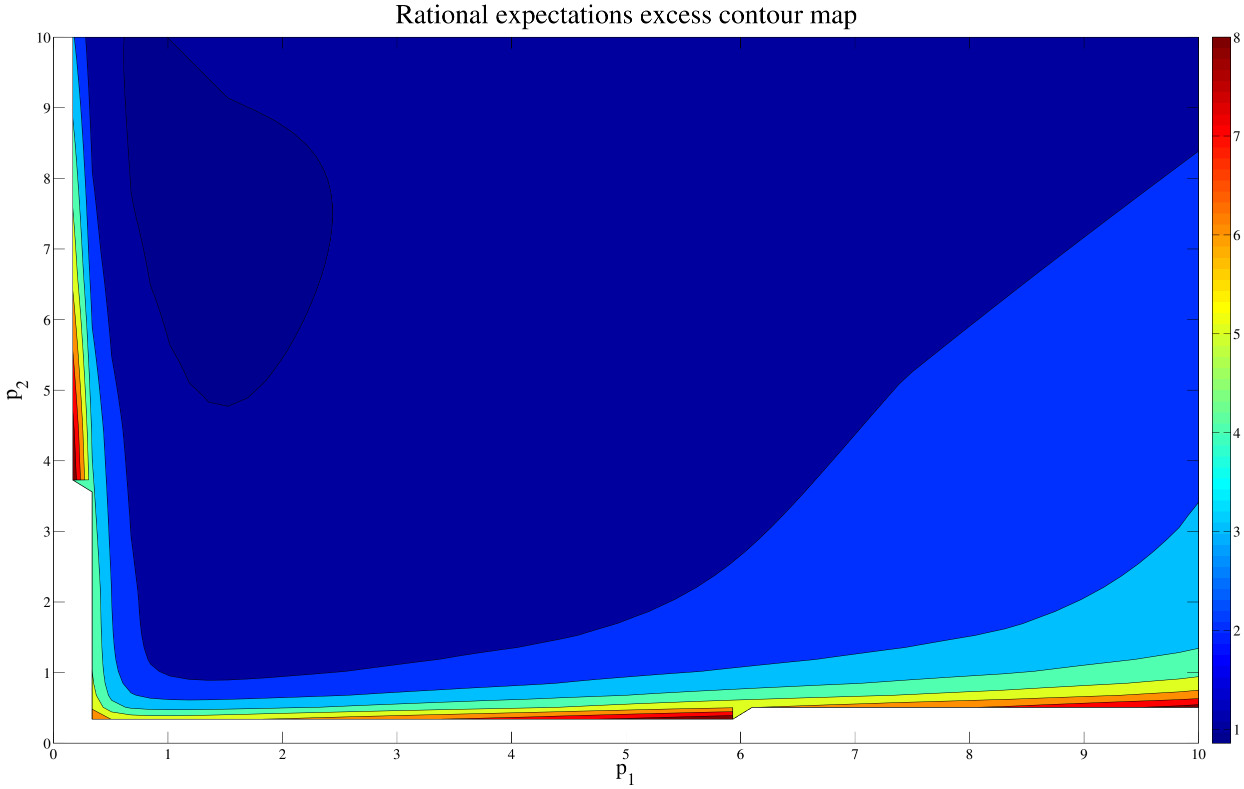

Figure 5.1: Contour map of normalised excess, $|z|_2^2$, under rational expectations.

Setting up our market under rational expectations gave us the opportunity to find for a benchmark for further equilibrium approximations. In theory, under rational expectations our market should fully clear on every period as, by definition, the rational expectations equilibrium occurs when traders have the same expectations, and hence form portfolios in equal proportions. To examine if this was indeed the case we set up our market with 50 agents all operating under rational expectations forecasts, we included 3 assets into the market with all supplies set to 1. We then simulated 50 periods of trading. The resulting approximation errors can be seen in Figure 5.2. We can immediately see that the errors have extremely small variance (this is why all indices on the y-axis appear as 1). However, what is more surprising is that the error is not 0 or even close to 0, indeed, the error is almost consistently 1. This was most troublesome, it appeared that all agents in our market where completely ignoring one of the three assets and the market could only achieve a partial equilibrium with the other two assets.

We tried many different market set-ups with varying amounts of assets and assets but the problem seemed to always remained. The answer inadvertently emerged while we were attempting to add mechanisms into our market for allowing short sales. We found that by allowing agents to short sale, the problem disappeared. This led to our first key insight into market equilibrium existence.

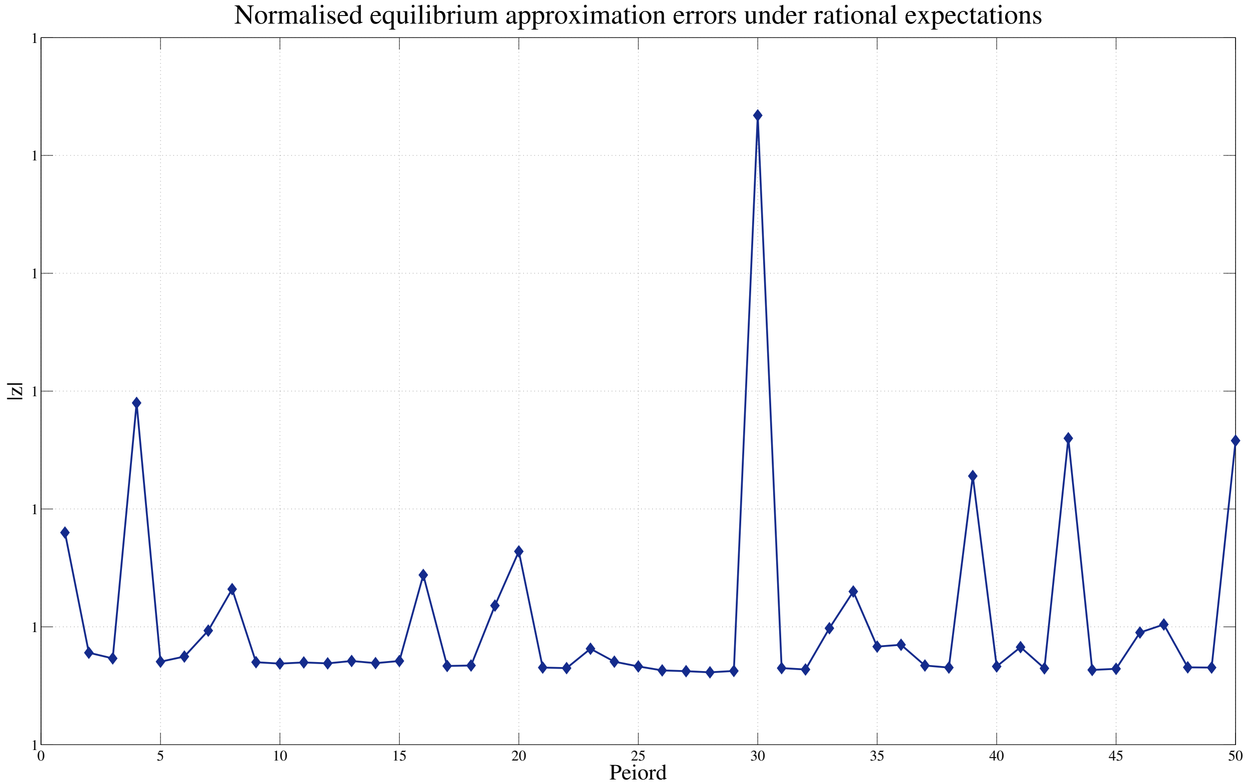

Figure 5.2: Normalised equilibrium approximation errors under rational expectations.

5.1.2 Consequences of No Shorting

As suggested above, it appeared that by allowing agents to short sale our market could be fully cleared. In Section 3.3.1 we made the decision to disallow short sales to simplify the design of our market, it seems this is a route taken by many artificial stock-market designers, in some cases we found that the approach taken to shorting was far from adequate, for example Palmer et al. (1994) and LeBaron (2010) both use a portfolio selection policy that *does* allow shorting (i.e by deriving the agent portfolio problem with no portfolio constraint), but then goes on to say that agents will *not* be permitted to short sale. We can only assume that portfolios that would usually have negative weights are being truncated.

For the current class of artificial markets this does not seem to be a problem, even for those that use equilibrium conditions; as we have seen, these markets are only required to find a equilibrium price that clears the market for a single asset, indeed, in LeBaron (2010) equilibrium prices are found by way of a simple binary search. However, for our market with a minimum of two assets being involved in equilibrium it appears the consequences of disallowing short sales have a direct impact on the existence of equilibrium prices. To see if this was the case we set up the market described above in the same manner, but this time we permitted agents to short sale by changing the agents portfolio selection from the quadratic program 3.3.7, to the analytical portfolio choice 3.3.13. The first difference we found was a dramatic decrease in computation time and memory requirements, though this is perhaps not surprising; agents are no longer having to invoke a quadratic program solver to find optimal portfolios. Of much greater interest was the change in equilibrium search space, recall that from Figure 5.1 that although the search space was convex, it was still somewhat irregular, and more importantly did not fall to below 1.

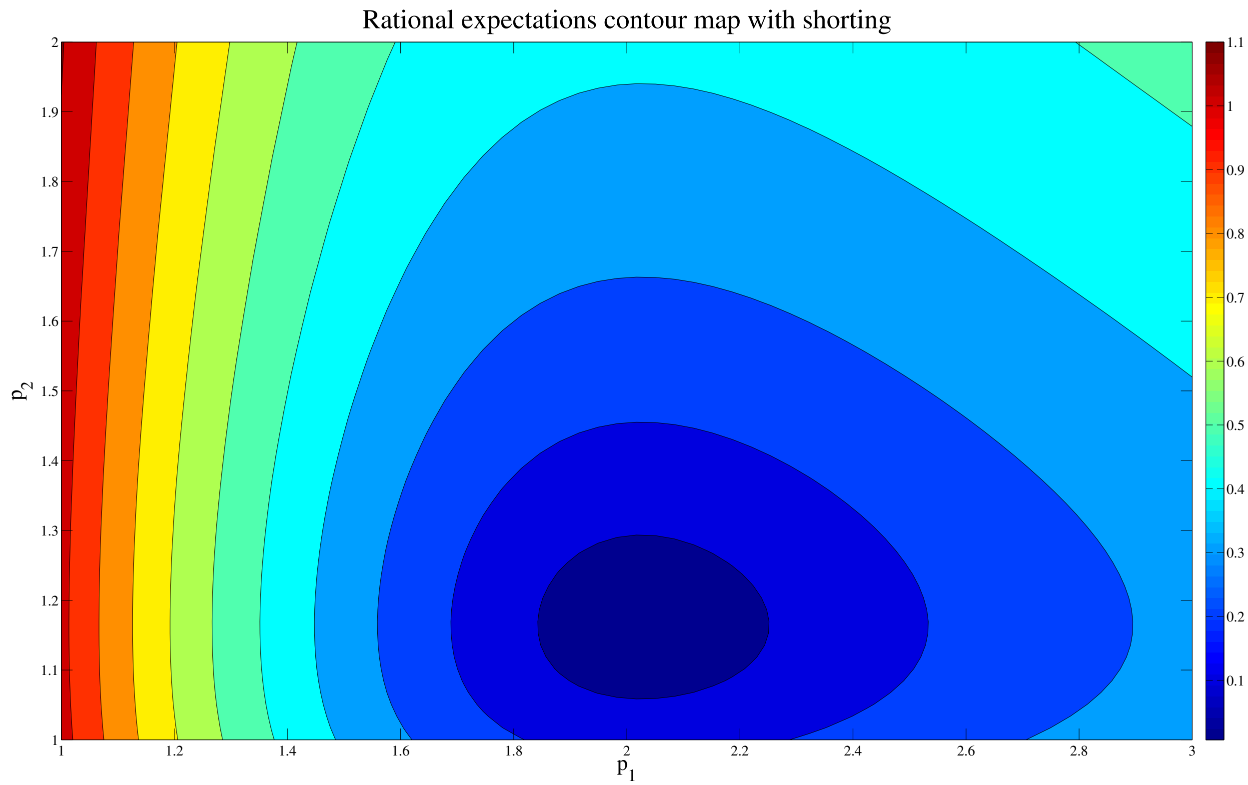

Figure 5.3: The equilibrium search space under rational expectations with shorting allowed.

A contour map of the new setup with shorting is shown in Figure 5.3. There are two key insights this gives us, firstly the search space is much more regular, this may be because agents are no longer hitting a minimum value in their portfolio selection. More importantly we can see the search space is now contains points that are close to zero. We conclude that in artificial stock markets using general equilibrium pricing, short selling has a direct impact on the existence of equilibrium and indeed, the stability of finding those equilibrium.

Unfortunately, we were not able to adapt our market to process short sales in full. Like many others, we have found this to be one of the most challenging design questions; it seems there is no clear method for settling short trades within a centralised environment. The problem is somewhat similar to the those faced by introducing social interactions described in Section 3.4. Even with the precautionary measures we took by disallowing agents to claim dividends earned on borrowed assets, we could find no way of actually paying the dividends to the agent that lent the asset, it is not even clear *which* agent lent the asset in our model. It would appear that once again, a decentralised market would solve this problem. We envision that in this setting it would be much easier for tracking exactly which stocks are owned by which agents.

With no clear method for clearing assets that have been shorted, our market was soon overwhelmed by short sales that far outreached the endowments of each asset. The resulting wealth dynamics are illustrated in Figure 5.4, agents with low risk aversions start shorting the safer assets and find that they soon become bankrupt, once this occurs the agent demand formula, which as we showed must satisfy Walras' law effectively inverses agents demands and soon there is no wealth left in the market. Once this occurs there is no longer equilibrium, the search for equilibrium prices soon exponentiates and the market descends into a state of stagnation.

Figure 5.4: Wealth dynamics for a selection of agents in a market with shorting.

5.1.3 Existence in Non-rational Markets

Having learned that our market is incapable of processing borrowing, and that short-sales are key for rational expectations equilibrium the most important question is now "*does equilibrium exist in markets without rational expectations and shorting disallowed?*". We find ourselves in a somewhat troublesome position; it could be equilibriums do exist but our search method cannot find them. To try and determine if this was indeed the case we ran small simulations with all three algorithms described in Section 4.3 in parallel and observed if they converged and diverged from equilibriums in unison. If just one algorithm could find an equilibrium point while the others failed, we could assert with greater confidence that the divergence errors described in Section 4.4.1 are more likely to be a result local optimums rather than equilibria non-existence.

The results for one simulation, shown below in Figure 5.5 are highly revealing. We can clearly see that there are periods for which one algorithm has found an approximate equilibrium point close to 0 which the others have only managed to find an approximate equilibrium point that is nearer 1. We found this phenomenon occurred during almost every simulation in this parallel set-up we did, indeed in some 44% of periods over a series of 200 we found that the Nelder-Mead and COBYLA were able to find a better approximate equilibrium point than L-BFGS-B. The reasons behind this are difficult to conclude, it would appear beyond doubt there is a fragility with each of the algorithms; as we noted in the previous chapter the Nelder-Mead algorithm is highly volatile in convergence, this can also be seen in Figure 5.5 where in periods 0, 4, 6, 9 and 16 it was able to find an approximate equilibrium close to zero while both L-BFGS-B and COBYLA were not. Nevertheless, it cannot be ignored that there are periods such as 3, 17 & 18 where no equilibrium close to 0 could be found.

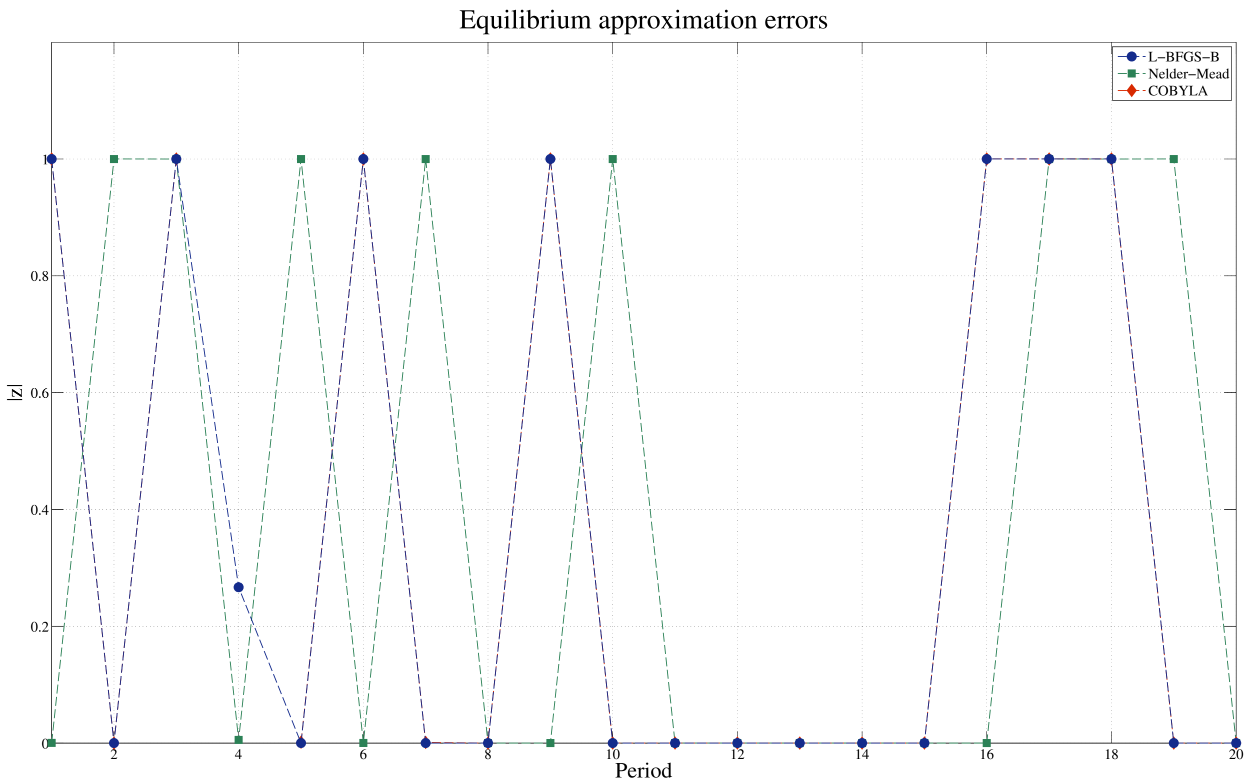

Figure 5.5: Equilibrium approximation errors for L-BFGS-B, Nelder-Mead and COBYLA.

Unfortunately, for the same reasons we were unable to find a mathematical argument for existence, we have no method of determining whether equilibriums actually exist for periods where none of the algorithms could find any. We are thus reliant on the accuracy of the available algorithms which as we have seen, are not guaranteed to succeed. However, one observation that is quite telling is that the error seems to be bounded by the initial supply of the asset (1 for the simulation in Figure 5.5). If no equilibrium exists for these periods it is almost certainly for the same reasons given in the previous section, namely restriction of short selling. We can try and gain some further insight into this by looking at the number of *partial equilibriums* that were found for each asset during computation of general equilibrium.[1] If it is found that one particular asset has a significantly low partial equilibrium count then it would seem reasonable to conclude that the short sale restriction is the indeed the problem. This is because if there is just a handful of assets that are constantly out of equilibrium it suggests that agents are hitting a boundary on their demands for those assets, this is exactly the restriction short sale restriction puts on agent portfolios.

We simulated our market with the same set-up used to compute the results in Figure 5.5 but also recorded the number of partial equilibriums that were found during each iteration of the L-BFGS-B search, hence it is the relative number of partial equilibriums that is important, not the actual number. Figure 5.6 shows the results. We can see that asset 6 is in partial equilibrium significantly less than the other assets, this is very strong evidence for our hypothesis that short sale restriction is causing general equilibrium non-existence.

[1] Recall from Section 2.2.1 that an asset is in partial equilibrium when when it is in equilibrium independently of the other assets. For example we could have 7 out of 8 assets in partial equilibrium. We need all assets to be in partial equilibrium to have a global equilibrium.

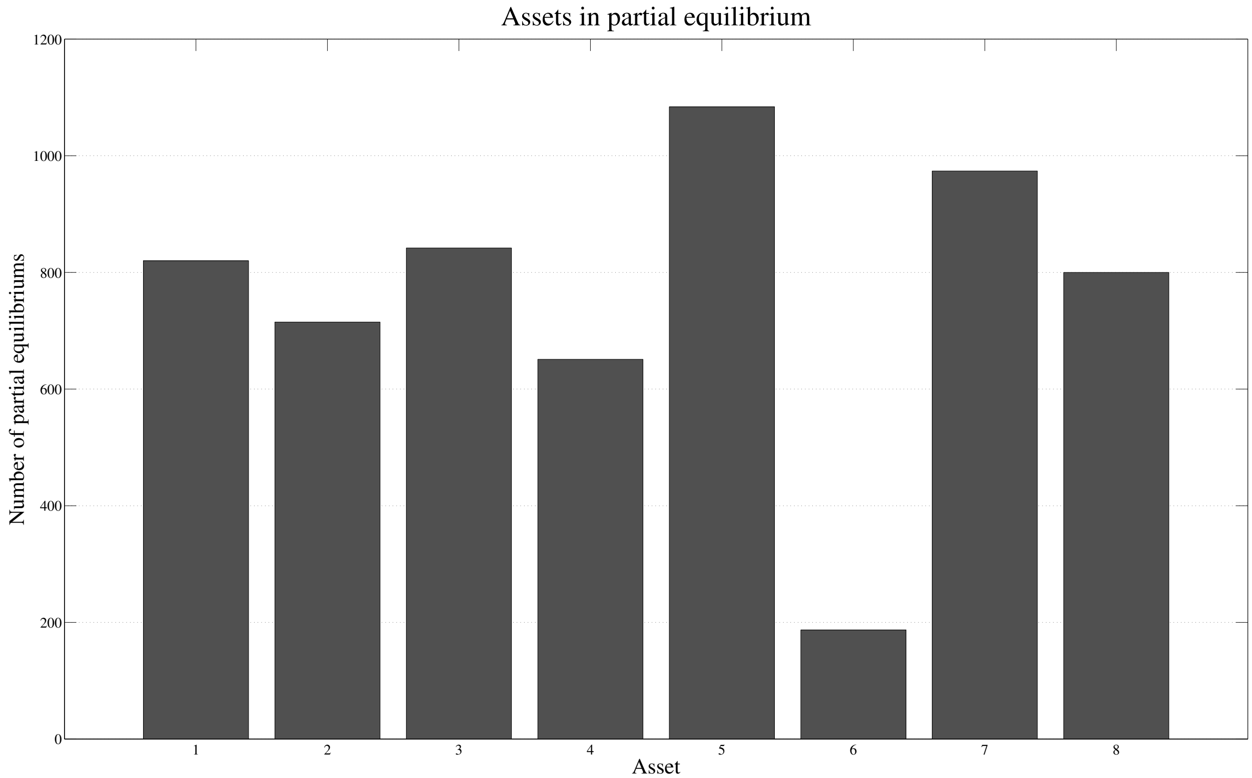

Figure 5.6: Partial equilibriums found during computation of general equilibrium using L-BFGS-B.

5.2 Stability

In Section 2.2.2 we discussed some of the key questions faced by general equilibrium theory. One of these, existence, has been already been discussed. Another important question was that of *stability*; how well could the equilibrium search mechanism actually find equilibriums? This is an import question for our model.

To test for stability in our artificial market we needed to establish a bound for equilibrium approximation errors (any equilibrium prices found by the search mechanism would need to have a normalisation error below this bound to be recorded as an approximate equilibrium). We felt the most natural choice given the findings we have just made would be a normalisation error of 1, this accounts for the short sale problem with a single asset. We also had to decide which algorithm to chose for running stability tests. From our initial analysis in Section 4.4.1 we stated the L-BFGS-B algorithm would be the best choice. Considering this, and the analysis on existence we have completed in this chapter we think this is probably the best choice; although Nelder-Mead does seem to find equilibrium points where the others don't, as can be seen from Figure 5.5, L-BFGS-B finds equilibriums more often, indeed, in our parallel existence testing we found L-BFGS-B found equilibriums better, or equivalent equilibriums in 56% of periods.

For this test we simulated our market under different initial settings, in particular with a variable number of assets and agents. The number of periods was kept fixed at 52. We ran each test 5 times, and used the average of the worst 4 to calculate the percentage of periods where approximate equilibriums were found. The results, shown in Table 5.1 indicate that even with just L-BFGS-B our pricing mechanism was able to find an acceptable equilibrium in the majority of periods.

Table 5.1: Percentage of approximate equilibriums found for different market set-ups.

No. of agents	No. of assets	% of approximate equilibriums found
2	2	99.5
2	3	99.0
2	8	97.6
20	2	99.5
20	3	97.6
20	8	96.6
30	2	98.6
30	3	96.2
30	8	90.0
40	2	98.6
40	3	95.7
40	8	88.5

Having said this it does seem clear that the accuracy of the approximations decreases with the number of assets, and to a small degree the number of agents. The increases error with increasing number of agents can be safely attributed to the accuracy of the l-BFGS-B method, indeed, it is quite likely that that a better Jacobian approximation method, or a larger number Hessian approximation vectors would eliminate this error. The increased error with number of assets is almost certainly due to more assets being included into the "short sale problem", there is little we can do to counter this error other than the modifications we suggested in Section 4.2.5. Our advice to those wishing to design their own general equilibrium artificial stock markets would be do introduce a robust short-sale mechanism, although as already discussed, this is not an easy task.

5.3 Uniqueness, Sensitivity & Optimality

The remaining questions of general equilibrium theory outlined in Section 2.2.2, while interesting, were not addressable in the time we had. However, we can offer some of our thoughts on these issues given our experience with our market.

- **Uniqueness**: We are almost certain that equilibrium points in our market are unique. While we have found no evidence to confirm this claim we have found non to the contrary either. It is also worth mentioning that our market is 'close' to satisfying the conditions for uniqueness given in Section 2.2.2.

- **Sensitivity**: Our market seems does seem to be very sensitive to initial conditions. We often found that on some simulations equilibrium points would be relatively close, while on other occasions they would oscillate wildly. We feel this is perhaps the most interesting questions to be answered, and would certainly require much mathematical analysis such as that undertaken by Brock and Hommes (1998).

- **Optimality**: Somewhat similar to uniqueness, we would be very surprised if equilibrium prices were not optimal, or almost optimal. Our market is very close to satisfying the *First Fundamental Theorem of Welfare Economics* conditions which would guarantee this is indeed the case.

5.4 Conclusions

We conclude with strong confidence that general equilibria do exist in our market when short sales are permitted, however it is not so clear when short sales are disallowed. It is unfortunate we can not complete similar partial equilibria analysis on our market with short sales permitted. Nevertheless, we feel the approximation error caused by a small amount of assets being out-of-equilibrium during the course of trading is not of great significance, indeed, we can easily neutralise the error by adjusting wealth in our market through consumption.

It is interesting to hypothesise if the short sale problem we have discovered in our market has any parallels to reality. Could a real market function properly without borrowing? Almost certainly not, for one there is the obvious problem with investors simply not having adequate funds to purchase a more risky investment such as a homeowner buying a mortgage. Moreover, we postulate that we would see similar patterns to what we have seen in our market, namely that if investors were not able to short then demand for certain types of asset would be such that supply and demand would almost never be in equilibrium.

We are confident our pricing mechanism is stable enough to act as a price clearing mechanism in our artificial stock market. However, we are weary of the limitations given the problems we have described. We repeat that those wishing to pursue this line of work should think vary carefully about a robust borrowing mechanism to avoid the short-sale problems we have uncounted.

Finally, the assertions we have made on uniqueness, sensitivity and optimality are largely hypothetical, each would almost certainly be its own avenue of research. We feel the most important question that for the future of artificial stock markets is sensitivity, we say this because of all the properties discussed this is the one that would almost certainly have the most impact on market dynamics.

6

Market Dynamics

In this chapter we evaluate our artificial stock markets market dynamics. Having analysed the equilibria in our market in the last chapter and concluded that our approximate equilibrium pricing mechanism is robust enough to clear out market, we now turn to the consequences on the actual market of general equilibrium pricing. In particular we will try and determine if out market has any of the properties exhibited by real markets outlined in Section 2.3.1. These were **high volatility**, **high trading volume** and **fat-tailed returns**. Another interesting property we wish to determine is if our market satisfies the **efficient market hypothesis**; as discussed in Section 2.3.1 there is much debate as to whether real markets are efficient, thus our discussion here will be more hypothetical, it will however allow us to say with confidence whether our market converges towards a rational expectations equilibrium. We finish we a brief analysis on learning dynamics in our market.

6.1 Overview of Evaluation Procedure

Before beginning our evaluation we briefly describe our evaluation procedure, and in particular the set-up of our market during evaluation. Note our set-up for this series of tests is slightly different from those carried out during equilibrium analysis, for example during equilibrium analysis we carefully controlled certain parameters such as agent risk aversion. This was because these parameters had little effect on equilibrium conditions and therefore were not important for the results gathered.

6.1.1 Market Setup

Like most artificial stock markets, our model has many adjustable parameters that influence market dynamics. To evaluate every variable of our market is unrealistic, indeed, one of the main criticisms of artificial stock markets has been there are too many parameters.[1] Thus we focus on the parameter that are of most relevance to our task, namely the number of assets. All other parameters will be set as follows:

- **Number of periods**: Unless stated otherwise we keep the number of periods set to 52 to represent 1 year of market data (recall from Section 3.3.1 that our myopic agents are assumed to have an investment horizon of one week.). This is a reasonable test period given the computational requirements of our market and should allow us to properly analyse the data.

[1] See for example LeBaron (2006).

- **Number of agents**: The number of agents will be set within the range 40-60 depending on the number of assets. We choose this range based off the results of our computational requirement analysis conducted in Section 4.4.2. In terms of general project analysis this is perhaps our biggest disappointment, we have been unable to reach anywhere near the amount of agents we would have liked to include in our market. As discussed in Section 4.4.2 the memory requirements of out model went far beyond what we had anticipated, and in hindsight we would have almost certainly chosen a language that has explicit memory control such as C.[2]. Implementation details are almost always completely ignored by the agent-based economics society, but from our experience working with what is perhaps one of the most computationally demanding markets developed, it seems clear that if artificial stock market research is to continue in this direction then researchers should have memory requirements in mind prior to market development.

  It is difficult to asses the consequences of this flaw in our market, given, this is the first artificial stock market to use general equilibrium pricing we have no other markets for direct comparison. Further, we estimate we would need a machine with upwards of 1000GB of memory to reach the same numbers of agents as current artificial stock markets.[3]

- **Endowments**: Agents are endowed with equal amounts of each asset, this is useful as it gives agents equal wealth before trading begins. We can then asses more carefully how different agents fared in the market without having to consider any agents having had a 'head start'.

- **Forecasts**: Agents are given 30 of each type of the four forecasts given in Section 3.3.3. We follow LeBaron (2010) and set forecast gains according to the half-life rule

$$\frac{1}{2} = (1 - g_j)^{m_h}$$

  where m_h for each forecast is set as follows:

  - **Long-range gain**: 10, corresponding to a decay of 10 years. This may seem long given our simulation length, but we stress the long-range forecast is a benchmark for other forecasts, and should not be drastically influenced by short-run volatility.
  - **Other forecasts**: For all other expectation forecasts, and the active learning utility maximisation forecast are chosen uniformly from [0.5,1,2,5,10] corresponding to 0.25, 0.5, 1, 2.5 & 5 years respectively.

  We provide agents the covariance matrix of dividends returns calculated in 2 for their Σ matrix. This is so agents don't have to learn the asset relationship structure. Perhaps an interesting extension of our market would be to introduce another learning mechanisms for Σ, however, we feel there is already enough complexities within our market and we would like to see if agents can take advantage of this provided information.

- **Consumption**: The agent consumption parameter λ_i is adjusted dynamically to compensate for any approximation errors given by the clearing mechanism as described in Section 4.2.5, it initially set to 0.

[2] Implementation details are discussed in the Appendix 8.1.

[3] The number of agents vary greatly between markets usually depending on the type of forecasts used, for example those markets that use genetic algorithms usually contain around 500-1000 agents, but those that use "few-type" forecasting can have anything up to 100,000 agents.

- **Risk Aversions**: Agent risk aversions are drawn uniformly from the range [1,5], this allows for a large variety of agent risk characteristic in our market.
- **Clearing algorithm**: As suggested in the previous two chapters we will use the L-BFGS-B algorithm for all market dynamic simulations.

6.1.2 Simulations

As suggested in 3.5 we "warm-up" our market with data generated exogenously from our geometric vector autoregression. We set the warm-up phase to 1000 periods, this gives the agent forecasts a chance adapt to the dividends data.

Where possible we ran 5 simulations of our market and used averages for the following results. While this is satisfactory for return distributions and mean returns it would not be helpful for more in-depth analysis such as Granger causality testing. For these types of tests we generally analysed the simulation we deemed 'most interesting'.

6.2 Comparison with Real Market Dynamics

As stated in the introduction to this chapter we will be analysing if out market displays any of the characteristics of real markets outlined in Section 2.3.1. As a further evaluation point we will briefly look at wealth dynamics in our market, it will be interesting to see if traders generally do equally well or some become more wealthy than others.

6.2.1 Volatility

We analyse volatility from two perspectives. Firstly we measure volatility using standard mathematical calculation, then we visually examine some of our generated time series for signs of bubbles and crashes.

Log-return Volatility

We measure volatility as follows:

$$\sigma_T = \sigma\sqrt{T}$$

where *sigma* is the standard deviation of logarithmic yearly returns and T is the time horizon in years. Given we have one years worth of market data we use $T = \frac{1}{52}$.[4]

We obtained data[5] for a small selection of well known assets and the S&P 500 index.[6] The volatilities of these assets are shown below in Table 6.1. Other than the S&P 500 these assets were selected somewhat at random and represent benchmarks for assets in our market.

Following this we then simulated our market with varying number of assets and calculated volatilities for asset that were 'high risk', 'medium risk' and 'high risk'.[7] Note that all prices are given relative to the numeraire which clearly always has unity price. Note for markets with two assets there is no 'low' or

[4]It is common to use $T = \frac{1}{256}$ as there are only 256 trading days in the year, but our model does not consider "non-working days", and our model has weekly periods, thus this is not the appropriate measure.

[5]Data obtained from `http://ichart.finance.yahoo.com`.

[6]The S&P 500 is a collective index of 500 American companies and is generally considered the best reference for the market as a whole.

[7]These 'risks' were taken by looking at the dividends process covariance matrix underlying each asset in the market.

Table 6.1: Volatilities for some well known assets.

Asset	S&P 500	Apple Inc.	NVIDEA	Goldman Sachs	British Petroleum
σ_T	0.09697	0.14183	0.19752	0.16956	0.11849

'medium' risk assets, similarly for a three asset market there is no 'medium' risk asset. The results are shown below in Table 6.2

Table 6.2: Sample volatilities of assets in our market for varying number of assets.

	Volatilities, σ			
Number of assets	High risk asset	Medium risk asset	Low risk asset	Average
2	0.70539	N/A	N/A	0.70539
3	0.99472	N/A	0.49813	0.7464
4	1.03125	0.731082	0.50293	0.7551
8	1.63498	1.104892	0.69768	1.1459

Comparing the volatilities from our market with those from the real assets shown in Table 6.1 we can see that our market is over-volatile. Some possible reasons for this phenomenon are:

- **Number of agents**: One likely cause is the low number of agents in our market; if there were more agents then each agent would have less impact on equilibrium prices, we often found that demands by a few agents would dramatically effect the equilibrium price. To conform this we ran smaller tests with more agents and indeed found that the asset volatilities fell; with ~ 150 agents we saw an approximate 5% decrease in volatility.

- **Equilibrium approximation**: We discussed equilibrium approximation at some length in the last chapter. This could be a result of the approximation error caused by the short-sale problem; if an asset is out-of-equilibrium then it could cause a sharp increase in all of the markets asset prices.

- **Dividends calibration**: We found that the market dynamics to be sensitive to the underlying dividends structure. For example, consider Figure 6.1 where we have simulated our market with highly correlated dividends. We can see that the resulting time-series are also highly correlated. This suggests that the exogenous dividends has a strong influence on the dynamics of the market.

  The consequences of this are very interesting. Do dividends have a similar impact on real markets? This is a highly debatable subject and really comes down to whether one believes in the 'fundamentalist' view or the 'technical' view. Indeed, as we have seen, asset correlation plays a fundamental role in portfolio theory, and as asset prices are given as the discounted sum of future dividends (see 2.3.5) portfolio theory would seem to suggest that dividends correlations do have an important role in price dynamics. Those who believe in the market efficiency hypothesis would almost certainly take the opposite view.

It is quite possible that all these factors have contributed to the increased levels of volatility in our artificial market. It is unfortunate that we are not able to fully test the first two postulates more fully, for reasons discussed in this section and the last. Each of these postulates are quite possibly research questions in their own right. However, we can go some way in answering the 'excess volatility' phenomenon by examining the generated price time-series visually.

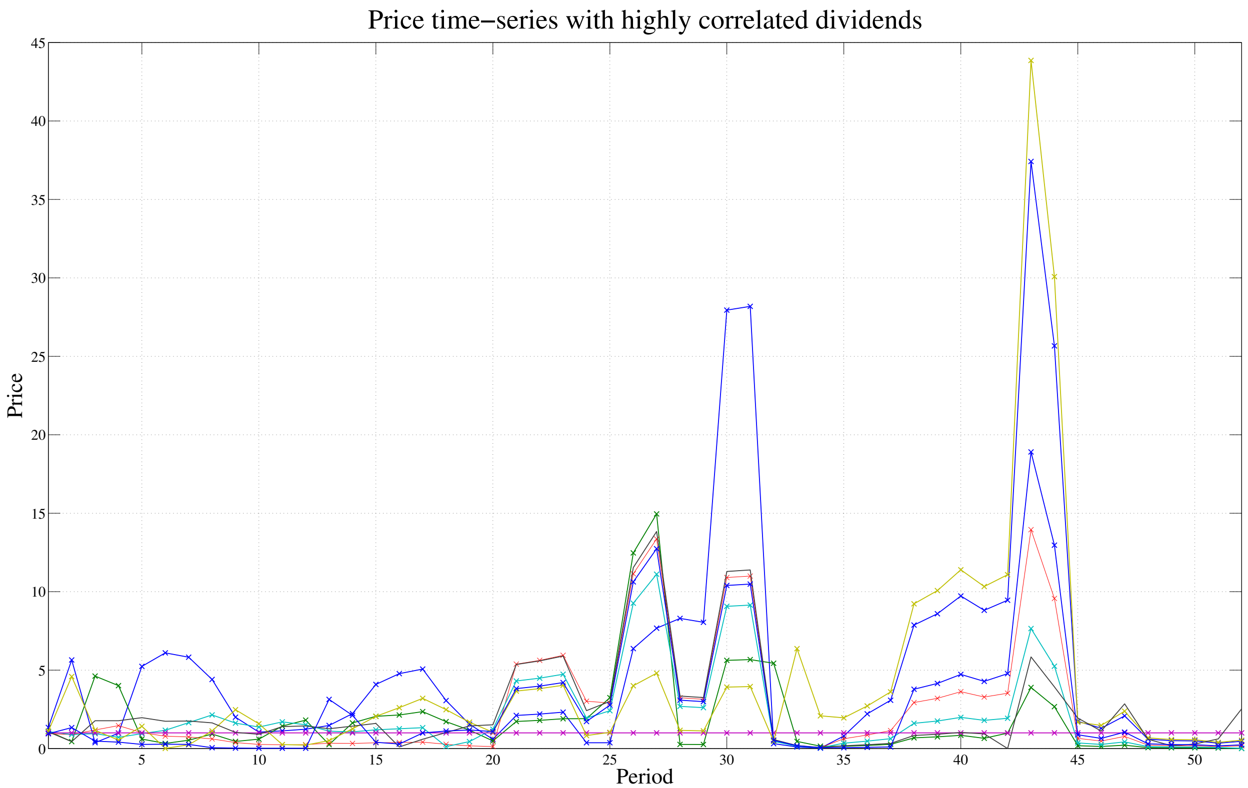

Figure 6.1: Price time-series with highly correlated dividends.

Bubbles & Crashes

Bubbles & crashes are a phenomenon seen often in real markets, and indeed many artificial stock market researches, have seen similar phenomenon in their markets. One market developed by Day and Huang (1990) is of particular interest. Their market uses out-of-equilibrium pricing, and they found that their market was particularly susceptible to bubbles & crashes in certain setups.

We found the "bubble & crash" phenomenon to be particularly prevalent in our market. An example is shown in Figure 6.2 where we can see a significant bubble. These types of events are deemed to be 'speculative' phenomenon which seems to contradict somewhat the type of dynamics we saw in Figure 6.1 where prices were highly correlated indicating high 'fundamental' presence.

This evidence seems to support the postulations we made previously; it is quite probable that with the small amount of agent we are restricted to there could be a higher proportion of one type of investor in the market, this would lead to very different types of volatility we have seen. Furthermore, it also seems to support our hypothesis that the equilibrium approximation error is causing such bubbles & crashes, as seen by Day and Huang (1990); this would also lead to a significant increase in volatility.

Another explanation for the higher than expected volatility is we have underestimated the time-horizon of our myopic agents. Originally we assumed time-horizons were one week which fell within the range of "myopic" investment horizons, however this choice was not determined mathematically. We found that by recalculating volatility in our market assuming period lengths of ~ 3 months the results were much closer to empirical results than before, for example in the 8 asset market shown in Figure 6.1 we found the volatility of the market average under a 4 month time-horizon was near 0.3, much closer to the observed figures.

Figure 6.2: Price time-series for a three asset market showing a ‘bubble’.

Furthermore, as found by LeBaron (2001a), the time-horizons of agent population has a dramatic impact on market volatility, this almost certainly agrees with our findings. Whether a long time-horizon is a problem or not really depends on the type of analysis required by the modeller. It was not within the scope of this project to analyse such details, however, we are confident that our model could easily be adjusted to accommodate different time-horizons if required. For example, one could adjust the number of terms used from the Mercator series to calculate the agents portfolio return, or, adjust the agents learning function parameters. However, we warn that by decreasing the time-horizon the computational requirements needed to simulate a satisfactory length of time becomes increasingly demanding, and also, the modeller must also be aware that trading costs should be taken into account when modelling smaller time-horizons which present another challenging design issue.

6.2.2 Return Distribution

The next important market property we consider is “fat-tails” or more formally, *kurtosis*. As we have discussed, real markets returns display kurtosis over certain periods, a property not present in existing general equilibrium models. Having discovered in the last section that our initial time-horizon estimate was too short to agree with empirical data (with a time-horizon of 1-3 months seeming more likely), evaluating our markets return structure against empirical data is a good further opportunity to add evidence to our new hypothesis; if empirical evidence supports our simulated market data over similar period lengths we can say with greater confidence our model is accurate over periods of 1-3 months. We begin by analysing kurtosis mathematically.

Definition. The kurtosis of a distribution, X, is defined as follows:

$$k = \frac{\mathbb{E}(X-\mu)^4}{\sigma^4}$$

A normal distribution has a kurtosis of 3, thus distributions with more outliers (or "fatter tails") have a kurtosis greater than 3. In order to evaluate if our artificial market gave returns with kurtosis we first gathered empirical return data from some well known assets for period lengths of 1 week (52 weeks total), and 1 month (52 months total) and calculated the kurtosis for each, the results are shown in below Table 6.3.

Table 6.3: Sample kurtosis for a selection of well known assets.

	Asset				
Period	S&P 500	Apple Inc.	NVIDEA	Goldman Sachs	British Petroleum
1 week (52 total)	3.0250	2.4166	2.0636	2.7565	3.2819
1 month (52 total)	3.9053	12.0580	4.5777	3.3339	12.6757

The first observation is that returns for a period length of 1 month have higher kurtosis that for period lengths 1 week, this agrees with the findings of Mandelbrot (1963) who discovered returns showed highest kurtosis for periods of approximately 3 months. The next observation is there is a large range of values, especially for monthly returns. Finally, we can see that a high kurtosis does not imply a high volatility.

We simulated our market with an 8 asset setup, however to for this simulation we made some of the adjustments suggested in the previous section, namely, increasing the number of agents (although only by a small amount), and recalibrating the exogenous dividends with a smaller covariance's. The results are shown below in Table 6.4

Table 6.4: Sample kurtosis and volatility for each asset in our 8 asset market.

Asset	1	2	3	4	5	6	7	8
k	12.4175	21.3300	6.7893	23.1179	7.9126	11.3376	7.3240	3.0724
σ_T	0.58172	0.626735	0.709685	0.58403	0.62947	0.66717	0.64546	0.21329

We can see that these results support our new hypothesis of longer time-horizons, indeed, this is strong evidence to suggest our model supports a time-horizon of approximately 3 months, this is highly interesting given our exogenous dividends is calibrated to empirical dividends that is paid quarterly. We can also see that our model also displays a high range of kurtosis values, and, that like the empirical data, a high kurtosis does not imply a high volatility.

It is also useful to observe the return distributions visually. Figure 6.3 shows returns for a simulation of our market with just two assets, we can see this is log normally shaped like the monthly returns for Apple shown in Figure 6.4. Notice also they have a similar outliers on the far rights, this is exactly the type of "Black swan event" Taleb (2010) warns of and is strong evidence to suggest our model could potentially be a safer choice for policy evaluation. Figures 6.5 and 6.6 show the distributions used to calculate the kurtosis results. We can see that although there are some differences such as our model showing less kurtosis on the left tail, they are both fat-tailed. This can be seen more clearly by looking at the logs of the two distributions, shown in Figures 6.7 and 6.6. We can see here that both are normal, and reassuringly similar. This gives us great confidence our model is faithfully producing real market dynamics, even when limited by a small number of agents. We believe our model could give even more impressive results if this limitation was not present.

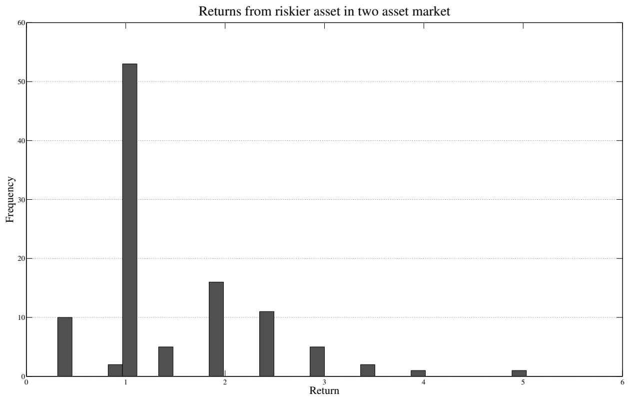

Figure 6.3: Period returns from a simulation of our market with 2 assets.

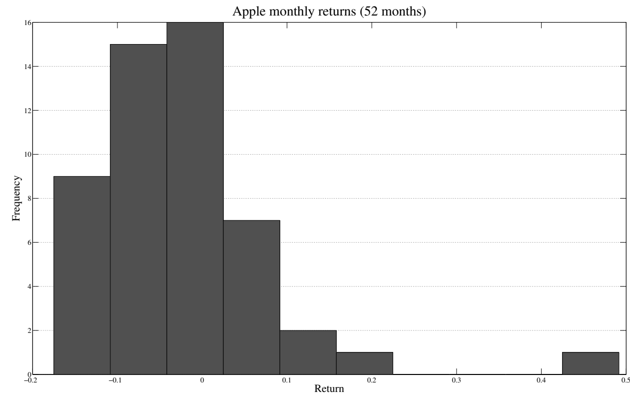

Figure 6.4: Monthly returns for Apple over 52 trading weeks.

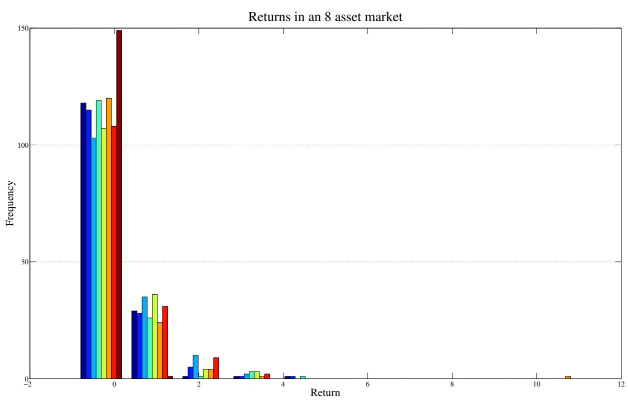

Figure 6.5: Period returns for a simulation of our market with 8 assets.

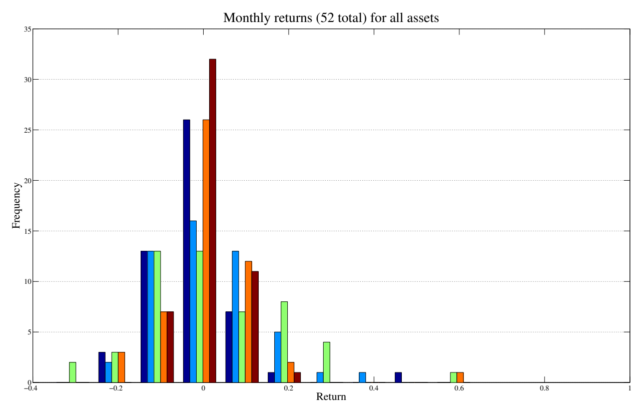

Figure 6.6: Monthly returns for the assets in Table 6.3 over 52 trading weeks.

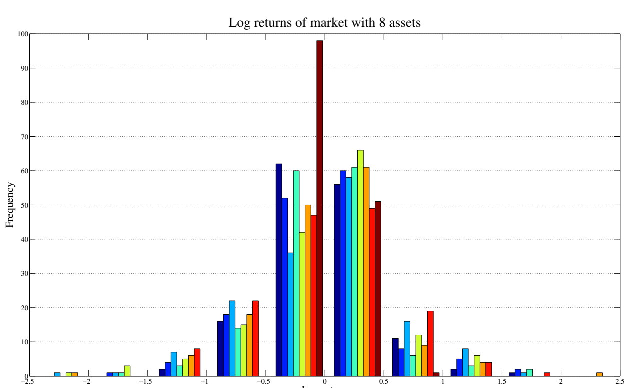

Figure 6.7: Log of Figure 6.5.

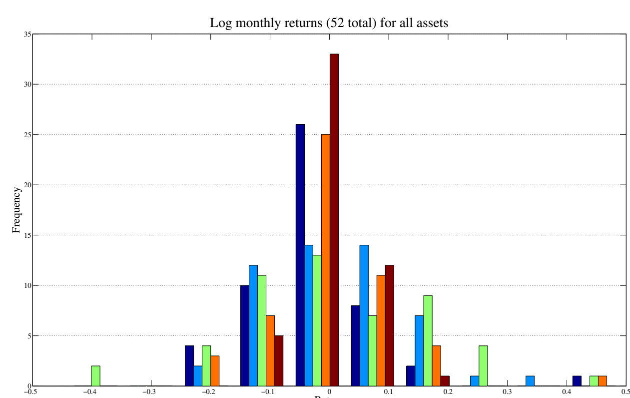

Figure 6.8: Log of Figure 6.6.

6.2.3 Trading Volumes

The next market property we examine is trading volume. We have seen that real markets display high levels of trading volume. It is quite challenging to record actual trading volumes in our market for the same reasons why a robust short selling mechanism is challenging, there is no arbitrator that records asset movement. The only role of the centralised exchange is to clear the market through prices. Another problem is that due to the limited number of agents in our market, it is difficult to say what constitutes a 'high' trading volume; or put another way, how many traders actually trade on real markets?

One way we can try and gain some insight into trading volumes is by examining the normalised portfolio changes made by our agents each period. We can conclude that if agents are changing their portfolios often then our market displays signs of high trading volume. Figure 6.9 shows the numeraire normalised price time-series of the 3 asset market simulation used for this test and Figure 6.10 shows the mean normalised portfolio changes for the entire market population (i.e the mean of $|\boldsymbol{w}_{i,t} - \boldsymbol{w}_{i,t-1}|$ for all agents i).

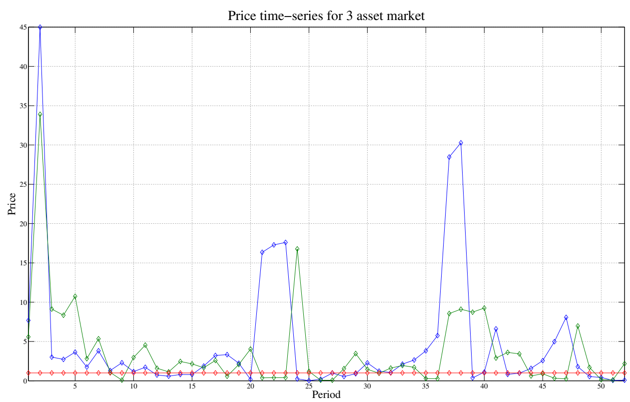

Figure 6.9: Price time-series for portfolio test.

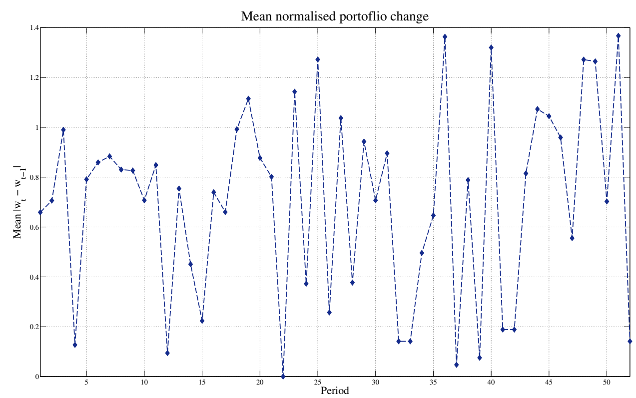

Figure 6.10: Mean normalised portfolio changes.

We can see that agents change their portfolios quite readily on average. This could suggest agents in our population are too speculative, this would also go some way to explaining the bubbles & crashes observed earlier. Further insight can be made by looking at the histogram of portfolio differences as shown in Figure 6.11. This, together with what we have seen in Figure 6.10 would suggest that there are actually two groups of agents, one group that change their portfolios often and another group that make few adjustments. This would actually suggest that some agents are becoming 'too' fundamental and implementing 'buy and hold strategies' while the market is being driven by speculative agents. Whether this is the case in real markets is a very open and interesting question. It is frustrating we are unable to carry out experiments with more agents in our market to see if this phenomenon is always the case, however, we hope that this work will motivate others to continue in the direction we have initiated.

One final observation we can make is that like in real markets, trading volume and market volatility have strong cross-correlation (see Section 2.3.1), for example we can see that when trading volume begins to increase in period 15, soon after in period 20 the market sees a sharp rise in volatility, and then in period 22 when trading volume sharply falls, the market price soon follows. This is repeated near period 40. It seems in our market that trading volume is dictating volatility, whether this is the case in real markets is an open question, but in our view this is more likely than the contrapositive.

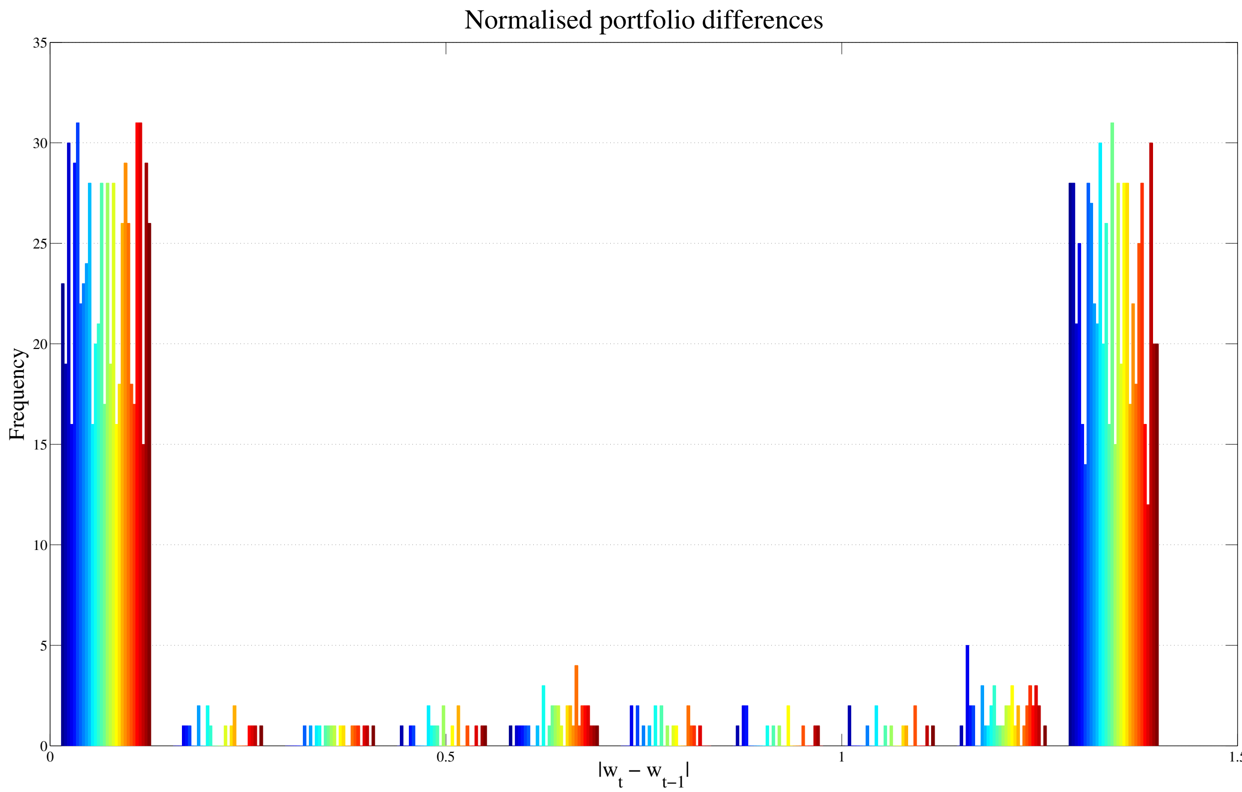

Figure 6.11: Histogram of agent normalised portfolio change.

6.2.4 Wealth Dynamics

We end this section with a brief examination of wealth dynamics in our market. Wealth dynamics is not a market property that is often discussed in current agent-based computing literature, however we feel it too has an important role in real market dynamics. Real markets are notorious for being populated with investors who experience highs and lows; this year top hedge fund manager is all to often next years market sheep. There are a very select few, such as Warren Buffet, who manage to continually make high returns on the stock market. We were curious to see how agent wealth dynamics changed in our market, for example, could agents sustain high levels of wealth?

The results are somewhat surprising, Figure 6.12 shows agent wealth dynamics as a time-series while Figure 6.13 shows the corresponding return series for the same market simulation. Firstly, we can see that agent wealth is volatile, this is perhaps not quite the level of fluctuation we would witness from a lay investor, but it is quite possible this degree of activity would be seen by firms with a large market presence and high risk strategies such as hedge funds. On the other hand there do seem to be periods where agents have sustained success (blue line), and one agent (green line) never seems to succeed.

Another interesting observation is overall wealth is not always correlated with market returns, indeed, we can see periods where market returns are high, but only a single agent has managed to take advantage of the opportunity. During other periods, the market has not done particularly well yet some agents seemingly see an increase in wealth. This kind of phenomenon is well known in stock markets were a small group of investors have managed to beat the market by some degree. This is a highly interesting subject and we feel much further development could be made from larger, more sophisticated simulations of our market.

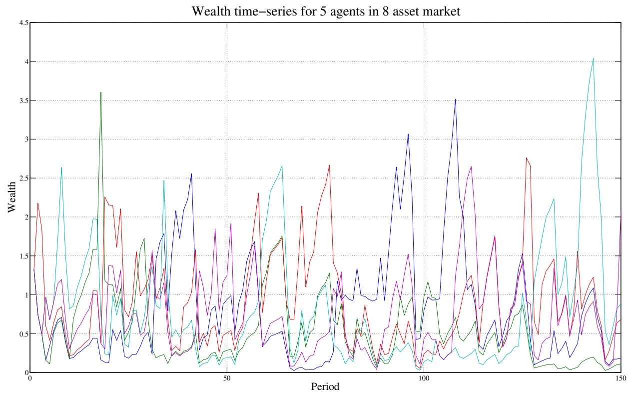

Figure 6.12: Agent wealth time-series.

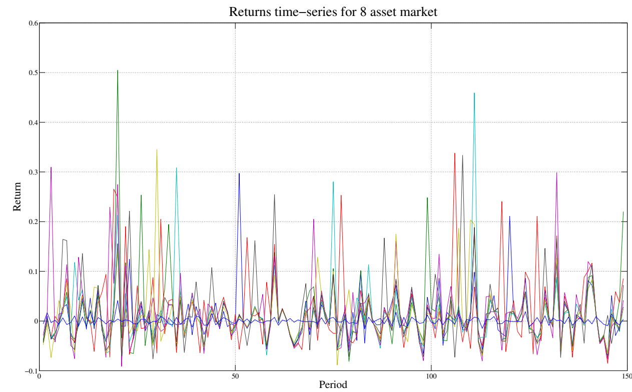

Figure 6.13: Return time-series for 8 asset market.

6.3 Market Efficiency

We now consider if our market is *efficient* as defined by the efficient market hypothesis described in Section 2.3.1. As explained, market efficiency is a very divided subject in economics, there has been evidence to suggest markets are efficient such as low autocovariance's between asset prices, and evidence to suggest markets are not efficient, such as market crashes & bubbles. We have already seen our market, like so many other artificial stock markets, experiences frequent bubbles & crashes. In general these events are explained by referring to investors expectations given their learning methods, however we feel it will be interesting to perform more abstract analysis and evaluate if our market is efficient, this would go some way to explaining the market dynamics we have witnessed.

To try and fully evaluate our market under a variety of conditions we compare two different market setups. The first is a standard 3 asset market with loosely calibrated dividends (which we will call "stable") and a 52 period simulation time. The second is an 8 asset market with highly correlated dividends that closely resembles a Gaussian sample (which we will call "unstable"), this market is run for 150 periods. The stable and unstable dividends can be seen in Figures 6.14 and 6.15 respectively. We can see the large difference between the dividend series, we would like to determine if these differences lead to changes in our market, and in particular if market efficiency can arise.

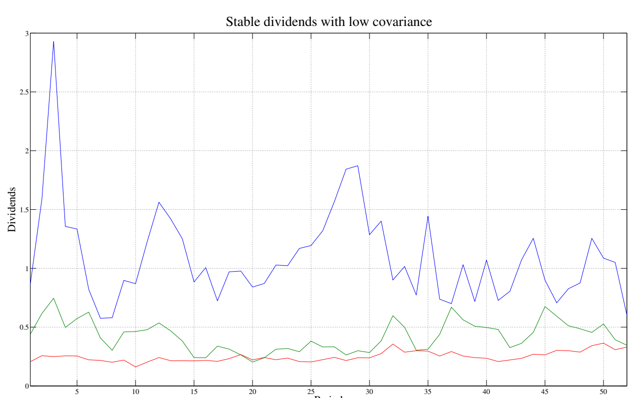

Figure 6.14: Stable dividends.

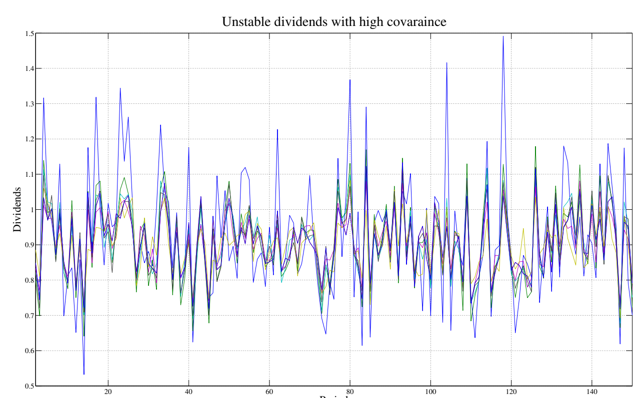

Figure 6.15: Unstable dividends.

When analysing time series for lagged influence it is useful to use a cross-covariance plot which shows the lagged covariance's between all variables in a system, a variables autocovariance is thus a special case. For two variables x and y the cross-covariance of lag m us defined as:

$$\phi_{xy}(m) = \mathbb{E}[(x_{n+m} - \mu_x)(y_n - \mu_y)^*]$$

The vector case follows. We can see the cross-covariance plots for the stable and unstable dividends streams in Figures 6.16 and 6.17 respectably. While both have high covariances, we can immediately see the the differences; the stable dividends has large cross-covariances compared to the unstable dividends suggesting past dividends in the stable case can be used to predict future values, while this is not the case for the unstable dividends.

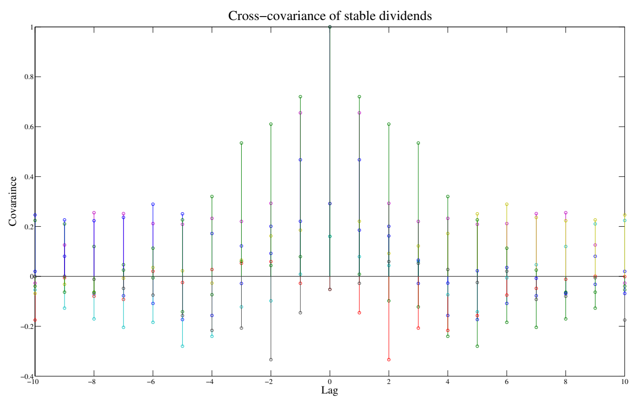

Figure 6.16: Cross-covariance of stable dividends.

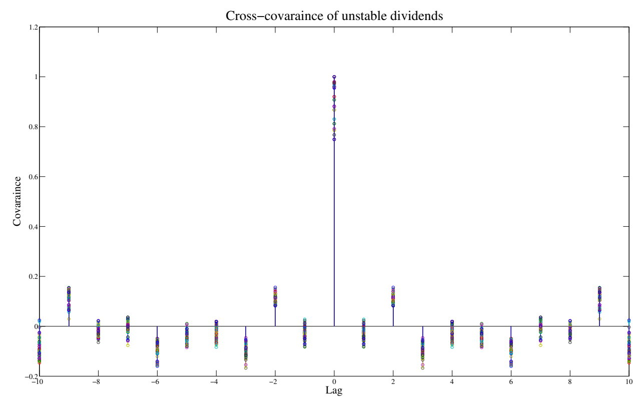

Figure 6.17: Cross-covariance of unstable dividends.

The resulting numeraire normalised price time-series for the simulated markets with stable and unstable dividends are shown in Figures 6.18 and 6.19 respectively. On first glance they seem quite similar, although we can see that the unstable market seems to have a higher covariance than its unstable counterpart. This would suggest our analysis at the beginning of this chapter was indeed correct; the dividends calibration needs refining. Moreover, this further shows the important influence dividends has on our markets dynamics and strongly suggests our market is sensitive to initial conditions, although as we have previously suggested, this is something that needs much further research.

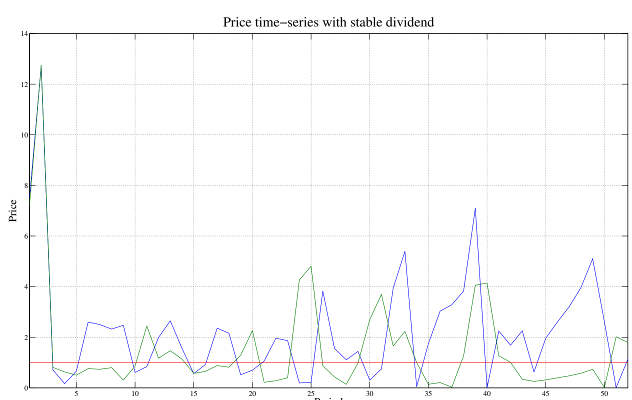

Figure 6.18: Price time-series with first set-up.

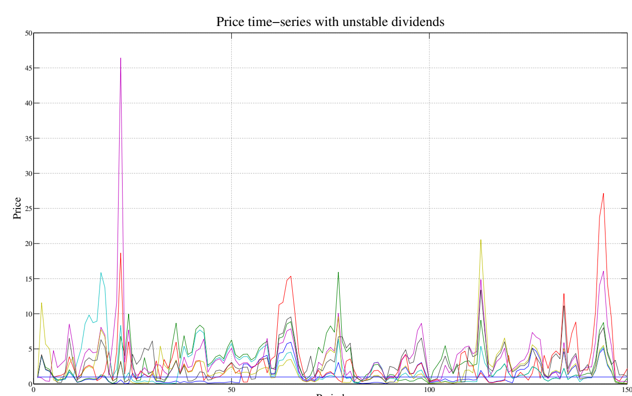

Figure 6.19: Price time-series with second set-up.

To examine the resulting time-series more thoroughly we can look at the each series cross-covariances. The cross-covariance's for the market with stable and unstable dividends are shown in Figures 6.20 and 6.21 respectively. The results are highly surprising, we can see that the market with stable dividends has significantly less cross-covariance than the market with unstable dividends. The stark contrast can be seen by comparing the two figures with Figures 6.16 and 6.17 which shows that the series seem to have inverted; from a dividends stream with high cross-covariance we have a time series with very little information between prices (although there is still some degree of autocovariance); and from a dividends series with virtually no cross-covariance there is seemingly much information between series.

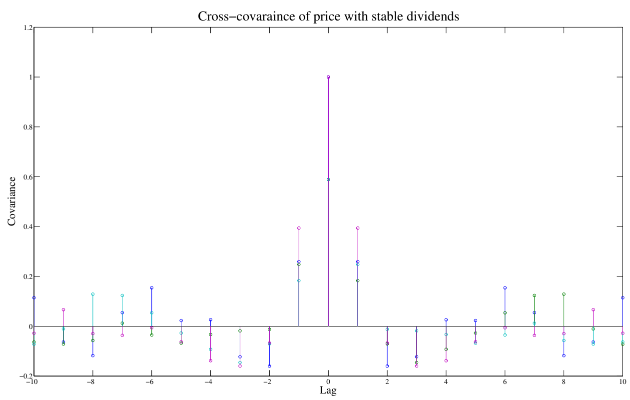

Figure 6.20: Cross-covariance of price with first set-up.

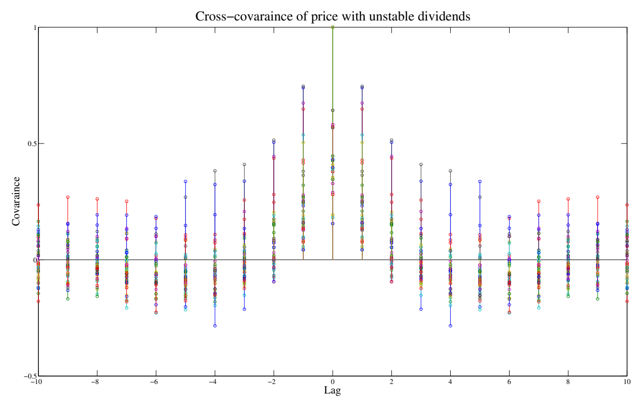

Figure 6.21: Cross-covariance of price with second set-up.

We can go one step further here and apply Granger causality analysis to both markets. Granger causality is a form of hypothesis testing to determine the influence of one variable in a system, over and above others.[8] A high Granger causality indicates one variable can be used to predict another in the system. We applied Granger causality analysis to both systems,[9] the results for the markets with stable and unstable dividends are shown in Figures 6.22 and 6.23 respectively. We can see that Granger causality testing showed there was no influencing variables in the first market, this supports the results of the cross-covariance statistics. On the other hand we can see that Granger causality revealed that there is a strong causal relation between asset 4 and 5 suggesting asset 4's price can be used to predict asset 5's. We found that the influence was actually stronger for longer lags than shorter lags, indicating that long term prices can be predicted. This is supported by what we can see in Figure 6.21 where for lags near 10 there is a rise in covariance between two assets.

These results are perhaps the most surprising and exciting of all our findings. It appear our market is able to produce efficient *and* non-efficient behaviour depending on initial conditions. To the best of our knowledge, no other artificial stock markets have displayed this kind of flexibility, although this is perhaps not surprising given that no other artificial markets are truly multi-asset.[10] While it would seem our market both agrees and disagrees with the efficient market hypothesis, there is much further research that could be

[8]The mathematics is quite involved and we will not discuss it here, instead we refer the reader to Granger original paper Granger (1969).

[9]Results obtained using Matlab toolkit developed by Seth (2010).

[10]We have already noted that Westerhoff (2004) developed a multi-asset market, and although his market did display signs of real market phenomenon, the market is highly simplified; prices are adjusted in response to excess demand and agents and there is no exogenous dividend.

done to reach a more established conclusion (if there is one), for example we have not been able to study the influence of equilibrium approximation, short selling or agent population on our markets efficiency.

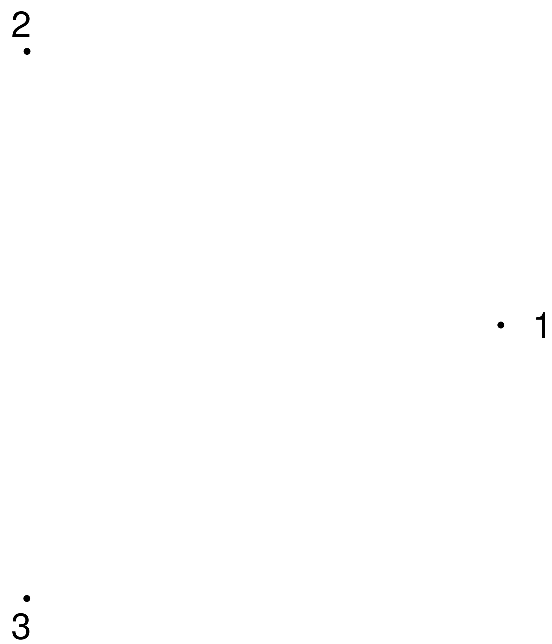

Figure 6.22: Granger causality relations from first set-up.

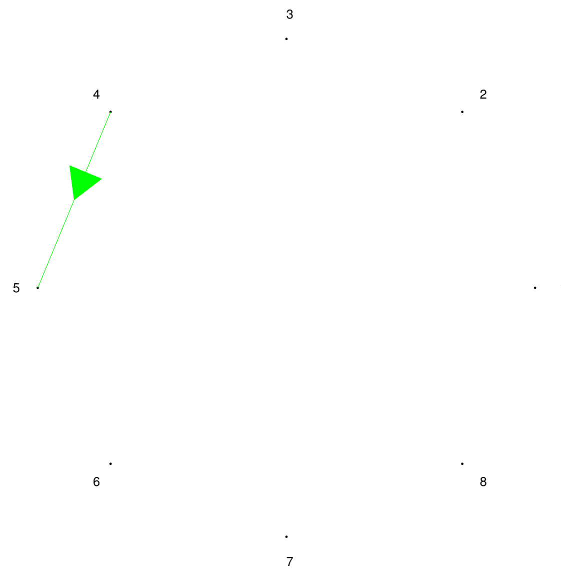

Figure 6.23: Granger causality relations from second set-up.

6.4 Agent Learning Dynamics

We finish our evaluation with a very brief discussion on agent learning dynamics in our market. We have seen how our market is susceptible to bubbles & crashes, and also that our market can display signs of inefficiency. Both would suggest agents are tending towards speculative strategies and away from rational expectations equilibrium. Interestingly, LeBaron (2010) who first developed the forecasting system we used, found similar types of price bubbles & crashes in his market. However, where as he found in his market the long-range ‘buy & hold’ strategies control much of the wealth in the market (although these agents have little influence on price setting), we found that agents almost completely ignore this strategy. Indeed, we found that our agents almost always tended towards the lagged return regression forecast, as shown in Figure 6.24. This is somewhat surprising and goes some way to explaining presence of both ‘fundamental’ and ‘technical’ phenomenon in our market (recall that the lagged return regression forecast lies somewhere in between both extremes).

We postulate that agents like this strategy as it looks to take advantage of any autocovariance’s in the price series, which as we have seen, can be present in our market. This would also suggest that unlike many artificial stock markets, convergence to rational expectations equilibrium is unlikely. This is yet another component of our market that has enough depth to be it’s own line of research. We would be interested to see if agents display similar types of learning when using “many-type” learning strategies such as classifier with a genetic algorithm.

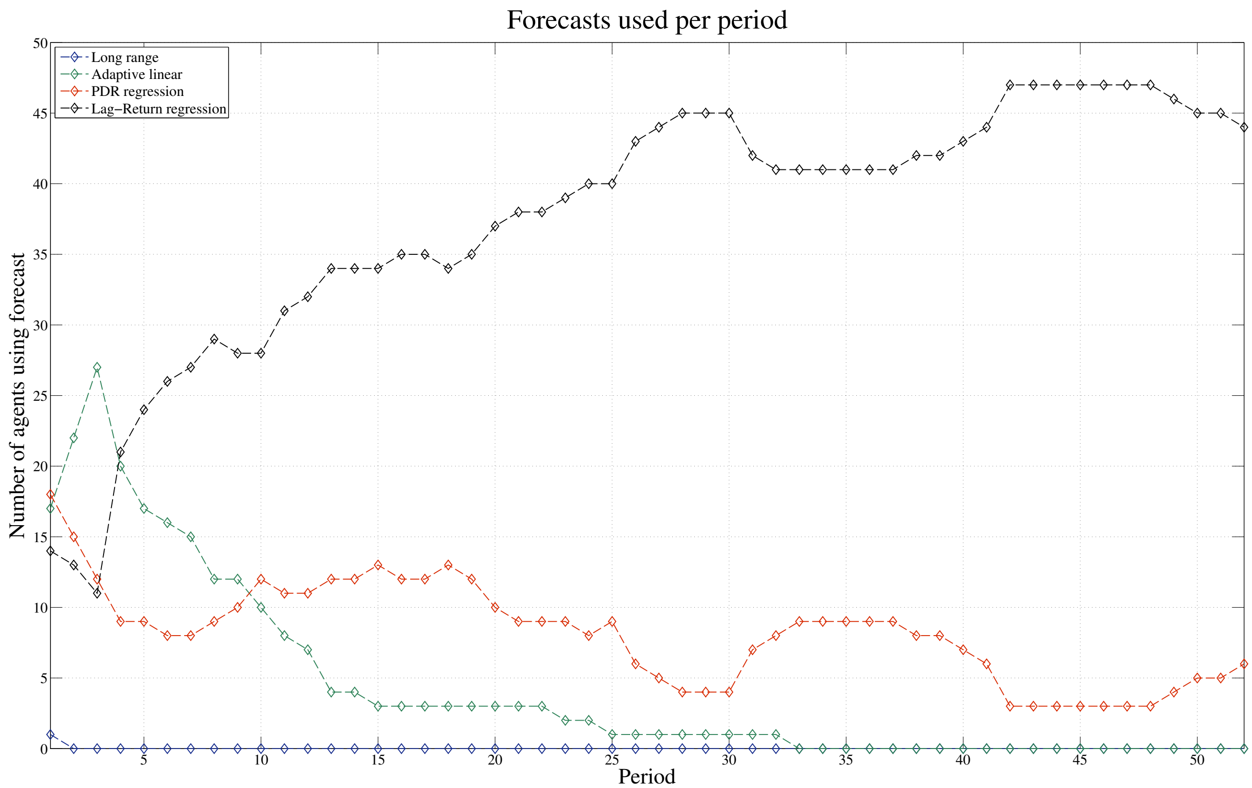

Figure 6.24: Time-series showing the number of agents using forecasts in each period.

6.5 Summary

In this chapter we have concluded the second phase of our evaluation process; our first evaluation phase considered the mechanics involved in our market, namely the equilibrium process that finds clearing prices. In this phase of our analysis we compared our market dynamics against empirical data and found that in many cases, our market was able to reproduce some important market phenomenon. in particular we found that under the correct time-horizon calibration, the price time-series produced by our market displayed **high volatility**. We also found that our market displayed signs of both fundamental *and* speculative events such as bubbles & crashes. Further, we found that our market return has high **kurtosis**. We followed by examining the trading volumes in our market and found that our market had high **cross-correlation**, and for the most part **high trading volume**. Our evaluation on market dynamics concluded with a brief discussion on the wealth dynamics in our market.

Finally, we discussed the **efficiency** of our market and observed signs of both efficiency, and non-efficiency. We concluded our market does not converge to rational expectations equilibrium, although noted that much further research is required to validate this claim.

7

Conclusions & Future Work

In this thesis we have laid the foundations for a new class of artificial stock markets that employs a general equilibrium pricing mechanism. Our model can be populated by a potentially unbounded number of exogenously determined asset, and is therefore to the best of our knowledge the first artificial stock market that does not require the concept of a risk-free asset. By removing the requirement of a risk free asset and introducing the sound principles of general equilibrium theory, our model has the potential to become a valuable tool for agent-based computational research.

7.1 Conclusions

During the process of developing our model, we found evidence to suggest that dividends returns are log-normally distributed, and developed a method for exogenously determined dividends using geometric vector autoregressive models that can be calibrated to closely resemble empirical dividends. Furthermore we showed a simple method for calibrating the geometric vector autoregression using an extension of a least-squared optimisation.

We showed that while our model does not satisfy all the conditions sufficient for general equilibrium existence, it is highly likely equilibrium do existence in our market when short selling is permitted, indeed, one of our main contributions is showing the importance of developing a robust borrowing mechanism for future artificial stock markets that using equilibrium pricing. Furthermore, we developed a non-linear non-convex optimisation problem for finding approximate equilibrium prices in our market and evaluated the robustness of the Nelder-Mead, COBYLA and L-BFGS-B algorithms for finding solutions.

Our artificial shock market is able to faithfully reproduce many well known real market dynamics, and in particular we found when properly calibrated is able to generate simulated market data that is volatile and has return distributions with high kurtosis. Further we found that our model encourages high trading volume and realistic wealth dynamics. Additionally, we found our market supports both fundamental and speculative phenomenon such as bubbles & crashes.

Finally, we showed our market is highly flexible and is able to generate price time-series with varying degrees of efficiency. Indeed, one of the main advantages of our model over existing artificial stock markets is the vast degree of dynamics that can be captured with relatively simple calibration.

7.2 Future Work

The work we have presented in this thesis is extremely broad in scope, and we have tried to highlight areas of potential further research throughout the course of our paper. We summarise these here.

- **Exogenous dividends calibration**: While we have presented an initial study on dividend time-series, and indeed a suitable model for generating dividends in the form of a geometric vector autoregressive model we feel their is much that could still be developed here. In particular, for more accurate dividends time-series our model requires further calibration.

- **Short selling mechanism**: One of the key discoveries of our project has been the importance of short selling for market equilibrium. Modelling borrowing in artificial stock markets is a challenging task. A robust protocol for handling borrowing in our model would have significant benefits for equilibrium approximations.

- **Social interactions**: As we have discussed, modelling agent interaction in a centralised market is a very challenging and interesting research area. By extending our model do allow agent interactions, either over fixed networks or through random intervention, another layer of dynamics would be introduced into our market.

- **Proof of existence**: An interesting mathematical challenge would be to prove existence in our market. This would represent a significant advancement in the advocation of our model.

- **Convex optimisation reformation**: One of the main implementation problems with our market is the equilibrium pricing mechanism requires iterated solutions to a non-linear, con-convex constrained optimisation problem. However, we have showed that the search domain to only show small degrees of non-convexity. It may be possible to take advantage of this and reform our optimisation problem as a convex optimisation problem which would hugely improve market performance and potentially unlock a new layer of market dynamics.

- **Sensitivity investigation**: An interesting line of research would be to investigate our markets sensitivity to initial conditions. It would be very interesting to see if our market is chaotic.

- **Learning mechanism**: In our model we used a relatively simple learning mechanism. This is a component of our market that can be easily modified, and represents an interesting area for research. One could implement more sophisticated learning mechanisms in our model and asses the consequences for market dynamics.

- **Market dynamics**: We have presented initial analysis on the market dynamics of our market. However there is a wealth of possible research area that could be explored here, in particular our model could be used as a tool to study market efficiency and other market phenomenon.

- **Policy evaluation**: Unfortunately we have not been able to study policy evaluation in our model. This is a line of research that should only be completed once our model is well understood.

- **C reimplementation**: One of the main points of criticism of out project is we were unable to populate our artificial market with large numbers of agents. A reimplementation of our market in C, with an eye towards better memory management could see a significant increase in agent populations.

- **Trading costs**: An interesting design challenge would be to include trading costs into our market. This would add a new layer of realism, and is especially important for smaller time-periods.
- **Wealth dynamics**: We found in our market to display very interesting wealth dynamics, however we were not able to dedicate much time to this area and it is certainly something that would be of great potential interest.

8

Appendix

8.1 Implementation Details

We briefly note some of our main implementation details.

- Our market was written in the Python programming language. We chose Python as it has a wide range of libraries available for matrix manipulation and optimisation. Although Matlab arguably has better support for these two features we decided it's poor object support would make it hard to implement an agent-based model. It it also reasonably efficient and has less overhead than a language such as Java (but is not as efficient as a well written C equivalent).
- We used the SciPy optimisation package for the Nelder Mead and L-BFGS algorithms. These implementations are reasonable and easy to use. In addition, we could not find any better open-source equivalents.
- We used the MOSEK quadratic program solver to calculate agent portfolios. The MOSEK solver is arguably one of the best on the market and its interface support with CVXOPT made this a relatively easy choice.

8.2 Proofs

Proof of Lemma 1. (From Balasko (1988)) In order for individual demand to be satisfied, it suffices that demand is satisfied for every commodity

$$\sum_i^m f_i(p, p \cdot \psi_i) \leq \sum_i^m \psi_i$$

Multiplying both sides by the price of commodity j, and assuming at least one equality is strict yields

$$\sum_j^L \left(\sum_i^m f_i(p, p \cdot \psi_i) \right) < \sum_j^L \left(\sum_i^m \psi_i \right)$$

Rearranging both sides gives

$$\sum_i^m p \cdot f_i(p, p \cdot \psi_i) < \sum_i^m p \cdot \psi_i$$

Applying Walras' law (2.2.1) to the left hand side yields

$$\sum_{i}^{m} p \cdot \psi_i < \sum_{i}^{m} p \cdot \psi_i$$

which is a contradiction, thus we can conclude strict inequality does not hold. □

8.3 Mathematical Background

First we define some important properties of functions that the reader should be aware of:

Definition. A function $f : \mathbb{R}^n \to \mathbb{R}$ is *convex* if $\mathbf{dom}\ f$ is a convex set and for all $x, y \in \mathbf{dom} f$, and for λ with $0 \leq \lambda \leq 1$, we have

$$f(\lambda x + (1-\lambda)y) \leq \lambda f(x) + (1-\lambda)f(y) \tag{8.3.1}$$

Definition. A function $f : \mathbb{R}^n \to \mathbb{R}$ is *concave* if $\mathbf{dom}\ f$ is a convex set and for all $x, y \in \mathbf{dom} f$, and for λ with $0 \leq \lambda \leq 1$, we have

$$f(\lambda x + (1-\lambda)y) \geq \lambda f(x) + (1-\lambda)f(y) \tag{8.3.2}$$

Definition. A function $f : \mathbb{R}^n \to \mathbb{R}$ is *quasi-concave* if $\mathbf{dom}\ f$ is a convex set and for all $x, y \in \mathbf{dom} f$, and for λ with $0 \leq \lambda \leq 1$, we have

$$f(\lambda x + (1-\lambda)y) \geq \max(f(x), f(y)) \tag{8.3.3}$$

Definition. A function $f : \mathbb{R}^n \to \mathbb{R}$ is *monotonic* if for any $x, y \in \mathbf{dom} f$ such that $x < y$, we have $f(x) < f(y)$.

The reader may also find the following properties from probability theory useful:

Lemma 4. $\mathbb{E}$ *is monotonic*

$$\text{if } X \leq Y \text{ then, almost surley } \mathbb{E}[X] \leq \mathbb{E}[Y] \tag{8.3.4}$$

Lemma 5. $\mathbb{E}$ *is linear*

$$\begin{aligned} \mathbb{E}[X + c] &= \mathbb{E}[X] + c \\ \mathbb{E}[X + Y] &= \mathbb{E}[X] + \mathbb{E}[Y] \\ \mathbb{E}[\alpha X] &= \alpha \mathbb{E}[X] \end{aligned} \tag{8.3.5}$$

Lemma 6. *The variance of a constant scaled random varaible is*

$$Var[\alpha X] = \alpha^2 Var[X] \tag{8.3.6}$$

Lemma 7. *The variance of the sum of two random variables is given by*

$$Var[X + Y] = Var[X] + Var[Y] + 2Cov[X, Y] \tag{8.3.7}$$

Lemma 8. *If* $X \sim ln\, N(\mu, \sigma^2)$ *then*

$$log \mathbb{E}[X] = \mathbb{E}[log X] + \frac{1}{2} Var[log X] \tag{8.3.8}$$

Lemma 9. *If* $X \sim ln\, N(\mu, \sigma^2)$ *then*

$$Var[X] = \left(e^{Var[log X]} - 1\right)\left(e^{2\mathbb{E}[log X] + Var[log X]}\right) \tag{8.3.9}$$

8.4 Vector Autoregressive Models

Time-series modelling is a critical part of econometrics, in general, we wish to analyse real time series to extract meaningful observations[1], or to forecast future results given some model. A time-series is written as

$$Y_T = \{y_1, y_2, y_3, \ldots, y_T\}$$

where T is the number of observations, note that we are assuming discrete time.

Definition. A process is *covariance-stationary* when the mean μ and the autocovariance's γ_{jt} are independent of t, that is

$$\mathbb{E}(Y_t) = \mu \tag{8.4.1a}$$
$$\mathbb{E}(Y_t - \mu)(Y_{t-j} - \mu) = \gamma_j \tag{8.4.1b}$$

for all t and any j.

Definition. We call a series of i.i.d variables ϵ_t

$$\{\epsilon_1, \epsilon_2, \epsilon_3, \ldots, \epsilon_T\} \tag{8.4.2}$$

where

$$\epsilon_t \sim N(0, \sigma^2)$$

a *Gaussian white noise* process.

Inherently

$$\mathbb{E}(\epsilon_t) = 0 \tag{8.4.3a}$$
$$\mathbb{E}(\epsilon_t^2) = \sigma^2 \tag{8.4.3b}$$

White noise processes are important building blocks for many other time-series. We are most often interested in stochastic dynamic processes that have some autocorrelation and as such cannot be reliably be modelled by a white noise process. The models employed for stochastic time series can be seen as realisations of *first-order difference equations*

$$y_t = \phi y_{t-1} + w_t \tag{8.4.4}$$

Notice we may recursively substitute to give

$$y_t = \phi^{t+1} y_{-1} + \sum_{k=0}^{t} \phi^{t-k} w_t \tag{8.4.5}$$

Thus we may fully define the series given an initial point y_{-1} and all w_t. If we substitute the white noise process (8.4.2) for w (i.e. $w_t = c + \epsilon_t$) we get a *first-order autoregression*.

Definition. A *first-order autoregression*, denoted AR(1) is given by

$$Y_t = c + \phi Y_{t-1} + \epsilon_t \tag{8.4.6}$$

[1] When applied to stock market time-series this is called "technical analyse".

As detailed in Hamilton (1994), when $|\phi| \geq 1$ the system the system is 'explosive', and no covariance-stationary process exists. When $|\phi| < 1$, (8.4.6) is a covariance stationary process and we are able to analyse the mean and covariance's.

Lemma 10. *If $\epsilon_t \sim N(0, \sigma^2)$ and y is covariance-stationary then $y_t \sim N(\mu, \gamma_0)$.*

Proof. y_t is a linear combination of normally distributed i.i.d's. □

Lemma 11. *If y is covariance-stationary the mean μ is*

$$\mu = \frac{c}{1-\phi} \tag{8.4.7}$$

Proof.

$$\mathbb{E}(Y_t) = c + \phi\mathbb{E}(Y_{t-1}) + \mathbb{E}(\epsilon_t)$$

As the process is covariance-stationary

$$\mathbb{E}(Y_t) = \mathbb{E}(Y_{t-1}) = \mu \tag{8.4.8}$$

Thus,

$$\mu = c + \phi\mu + 0$$

or

$$\mu = \frac{c}{1-\phi} \tag{8.4.9}$$

□

Lemma 12. *If y is covariance-stationary the variance γ_0 is*

$$\gamma_0 = \frac{\sigma^2}{1-\phi^2} \tag{8.4.10}$$

Proof. Substituting (8.4.7) into (8.4.6)

$$Y_t = \mu(1-\phi) + \phi Y_{t-1} + \epsilon_t$$

or,

$$(Y_t - \mu) = \phi(Y_{t-1} - \mu) + \epsilon_t$$

Squaring both sides and taking expectation gives

$$\mathbb{E}(Y_t - \mu)^2 = \phi^2\mathbb{E}(Y_{t-1} - \mu)^2 + 2\phi\mathbb{E}[(Y_{t-1} - \mu)\epsilon_t] + \mathbb{E}(\epsilon_t^2) \tag{8.4.11}$$

As $(Y_{t-1} - \mu)$ is a linear function of ϵ and the auto-covariance of ϵ is 0

$$\mathbb{E}[(Y_{t-1} - \mu)\epsilon_t] = 0 \tag{8.4.12}$$

and as the process is covariance-stationary

$$\mathbb{E}(Y_t - \mu)^2 = \mathbb{E}(Y_{t-1} - \mu)^2 = \gamma_0 \tag{8.4.13}$$

substituting (8.4.12) and (8.4.13) into (8.4.11) gives,

$$\gamma_0 = \phi^2\gamma_0 + 0 + \sigma^2$$

or

$$\gamma_0 = \frac{\sigma^2}{1-\phi^2} \tag{8.4.14}$$

□

Lemma 13. *If y is covariance-stationary the j^{th} auto-covariance is given by*

$$\gamma_j = \frac{\sigma^2 \phi^j}{1-\phi^2} \tag{8.4.15}$$

Proof. See proof of **Lemma 16**. □

In many situations we may wish to model multiple processes that have some degree of correlation, naturally this leads to *vector autoregressive models*, denoted VAR. A p^{th} order VAR is written VAR(p) and has the form

$$\boldsymbol{y}_t = \boldsymbol{c} + \Phi_1 \boldsymbol{y}_{t-1} + \Phi_2 \boldsymbol{y}_{t-2} + \cdots + \Phi_p \boldsymbol{y}_{t-p} + \boldsymbol{\epsilon}_t$$

where $\boldsymbol{c}$ is a $(n \times 1)$ vector constant, Φ_i are $(n \times n)$ coefficient matrices and $\boldsymbol{\epsilon}_i$ is a $(n \times 1)$ white noise vector satisfying

$$\mathbb{E}(\boldsymbol{\epsilon}_t) = 0 \tag{8.4.16}$$

$$\mathbb{E}(\boldsymbol{\epsilon}_t \boldsymbol{\epsilon}'_\tau) = \begin{cases} \Sigma & t = \tau \quad (8.4.17a) \\ \mathbf{0} & \text{otherwise} \quad (8.4.17b) \end{cases}$$

Definition. A VAR(1) has the form

$$\boldsymbol{y}_t = \boldsymbol{c} + \Phi \boldsymbol{y}_{t-1} + \boldsymbol{\epsilon}_t \tag{8.4.18}$$

For example, a bivariate VAR(1) model is given by

$$\begin{pmatrix} y_{1,t} \\ y_{2,t} \end{pmatrix} = \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} + \begin{pmatrix} \phi_{1,1} & \phi_{1,2} \\ \phi_{2,1} & \phi_{2,2} \end{pmatrix} \begin{pmatrix} y_{1,t-1} \\ y_{2,t-1} \end{pmatrix} + \begin{pmatrix} \epsilon_{1,t} \\ \epsilon_{2,t} \end{pmatrix}$$

A VAR(1) is covariance-stationary if the eigenvalues of Φ all lie in the unit circle, more formally $|\lambda| < 1$ in

$$|\boldsymbol{\iota}\lambda - \Phi| = 0 \tag{8.4.19}$$

thus for covariance-stationary VARs we can obtain the following

Lemma 14. *If y is covariance-stationary the mean $\boldsymbol{\mu}$ is*

$$\boldsymbol{\mu} = (\boldsymbol{\iota} - \Phi)^{-1} \boldsymbol{c} \tag{8.4.20}$$

Proof. We have

$$\mathbb{E}(\boldsymbol{y}_t) = \boldsymbol{c} + \Phi \mathbb{E}(\boldsymbol{y}_{t-1}) + \mathbb{E}(\boldsymbol{\epsilon}_t)$$

Assuming $\boldsymbol{y}$ is covariance stationary then

$$\boldsymbol{\mu} = \boldsymbol{c} + \Phi \boldsymbol{\mu}$$

Rearranging gives

$$\boldsymbol{\mu} = (\boldsymbol{\iota} - \Phi)^{-1} \boldsymbol{c}$$

□

Lemma 15. *If y is covariance-stationary the covariance's $\boldsymbol{\Gamma}_0$ are*

$$\boldsymbol{\Gamma}_0 = \sum_{k=0}^{\infty} \Phi^k \Sigma \Phi^{T^k} \tag{8.4.21}$$

Proof. Writing $\boldsymbol{y}_t$ in recursive form we have

$$\boldsymbol{y}_t = \sum_{k=1}^{t} \Phi^k c + \Phi^t \boldsymbol{y}_0 + \sum_{k=0}^{t} \Phi^k \boldsymbol{\epsilon}_k$$

Then

$$\boldsymbol{y}_\infty = (\boldsymbol{\iota} - \Phi)^{-1} \boldsymbol{c} + \sum_{k=0}^{\infty} \Phi^k \boldsymbol{\epsilon}_k$$

Taking the variance of both sides gives

$$\begin{aligned}
Var(\boldsymbol{y}_\infty) &= Var\left(\sum_{k=0}^{\infty} \Phi^k \boldsymbol{\epsilon}_k\right) \\
&= Var\left(\sum_{k=0}^{\infty} \sum_{j=0}^{\infty} (\Phi^k \boldsymbol{\epsilon}_k)(\Phi^j \boldsymbol{\epsilon}_j)^T\right) \\
&= \sum_{k=0}^{\infty} \sum_{j=0}^{\infty} \Phi^k Var(\boldsymbol{\epsilon}_k \boldsymbol{\epsilon}_j) {\Phi^T}^k \\
&= \sum_{k=0}^{\infty} \Phi^k \Sigma {\Phi^T}^k
\end{aligned}$$

□

Lemma 16. *If y is covariance-stationary the j^{th} autocovariance $\boldsymbol{\Gamma}_j$ are*

$$\boldsymbol{\Gamma}_j = \Phi^j \boldsymbol{\Gamma}_0 \tag{8.4.22}$$

Proof.

$$\begin{aligned}
\boldsymbol{\Gamma}_j &= \mathbb{E}[(\boldsymbol{y}_t - \boldsymbol{\mu})(\boldsymbol{y}_{t-j} - \boldsymbol{\mu})'] \\
&= \mathbb{E}[(\boldsymbol{c} + \Phi \boldsymbol{y}_{t-1} + \boldsymbol{\epsilon}_t - \boldsymbol{\mu})(\boldsymbol{y}_{t-j} - \boldsymbol{\mu})'] \\
&= \mathbb{E}(\boldsymbol{c}(\boldsymbol{y}_{t-j} - \boldsymbol{\mu})') + \Phi \mathbb{E}[(\boldsymbol{y}_{t-1} - \boldsymbol{\mu})(\boldsymbol{y}_{t-j} - \boldsymbol{\mu})') + \\
&\quad E(\boldsymbol{\epsilon}_t (\boldsymbol{y}_{t-j} - \boldsymbol{\mu})'] \\
&= \boldsymbol{0} + \Phi \Gamma_{j-1} + \boldsymbol{0} \\
&= \Phi \Gamma_{j-1}
\end{aligned} \tag{8.4.23}$$

It is easy to see repeated application of (8.4.23) will yield

$$\boldsymbol{\Gamma}_j = \Phi^j \boldsymbol{\Gamma}_0 \tag{8.4.24}$$

□

8.5 Mathematical Optimisation

Mathematical optimisation is an area of mathematics and computer science concerned with optimising (find the minimum or maximum) some function, sometimes subject to constraints. More formally, an optimisation problem has the form

$$\begin{aligned} \underset{x}{\text{minimise}} \quad & f_0(x) \\ \text{subject to} \quad & f_i(x) \leq 0, \quad i = 1, \ldots, m \end{aligned} \tag{8.5.1}$$

where the vector $x = (x_1, \ldots, x_m)$ is the *optimisation variable*, $f_0 : \mathbb{R}^n \to \mathbb{R}$ is the *objective* function and the functions $f_i : \mathbb{R}^n \to \mathbb{R}$, $i = 1, \ldots, m$ are the (inequality) constraints.

A solution to 8.5.1, x^*, is either *locally optimum* or *globally optimum*. A globally optimum solution satisfies $f_0(x^*) \leq f_0(x)$ for all x. A locally optimum solution satisfies $f_0(x^*) \leq f_0(x)$ for a neighbourhood N of x^*.

An optimisation problem is *convex* if the objective function and constraints are convex. Convex optimisation problems can almost always be efficiently be solved. To see this consider an alternative definition of a convex function:

$$f(y) \geq f(x) + \nabla f(c)^T (y - x)$$

Which is the second order Taylor expansion of f. Therefore if we set the differential of f w.r.t x to 0 we obtain

$$f(y) \geq f(x)$$

for all y. This is the first order optimality condition for convex optimisation. Convex optimisation problems are discussed in great detail by Boyd and Vandenberghe (2004).

An optimisation problem of the form

$$\begin{aligned} \underset{x}{\text{minimise}} \quad & \frac{1}{2} x^T P x + q^T x + r \\ \text{subject to} \quad & Gx \leq h \\ & Ax = b \end{aligned} \tag{8.5.2}$$

is called a *Quadratic Program*. Solution methods for quadratic programs are discussed in Boyd and Vandenberghe (2004) and also in Nocedal and Wright (1999).

If the objective function f_0 of 8.5.1 is non-linear and non-convex then the problem often becomes very difficult to solve. In general one must choose between a local optimiser in polynomial time or a global optimiser in exponential time. Methods for solving these types of problems are highly involved and are discussed in Nocedal and Wright (1999).

Bibliography

Beth Allen and James S Jordan. The Existence of Rational Expectations Equilibrium: A Retrospective. Technical report, Federal Reserve Bank of Minneapolis, August 1998.

Kenneth J Arrow and Gerard Debreu. Existence of an Equilibrium for a Competitive Economy. *Econometrica*, 22, July 1954.

W. Brian Arthur. On learning and adaptation in the economy. 1992.

W. Brian Arthur. Complexity in economic and financial markets. *Complexity*, 1(1):20–25, 1995.

W. Brian Arthur. Out-of-Equilibrium Economics and Agent-Based Modeling. In *Handbook of Computational Economics*, pages 1551–1564. 2006.

W. Brian Arthur, J H. Holland, Blake LeBaron, Richard G Palmer, and Paul Taylor. Asset pricing under endogenous expectations in an artificial stock market. In David Lane, editor, *The Economy as an Evolving Complex System II*, page 15. 1997.

Robert Axelrod. Agent-Based Modeling as a Bridge Between Disciplines. In Leigh Tesfatsion and Kenneth L. Judd, editors, *Handbook of Computational Economics*, volume 2, pages 1565–1585. North Holland, 2006.

Yves Balasko. *Foundations of the Theory of General Equilibrium*. Economic theory, econometrics, and mathematical economics. Academic Press, 1988. ISBN 9780120769759.

John M. Barron, Bradley T. Ewing, and Gerald J. Lynch. *Understanding Macroeconomic Theory*. Routledge advanced texts in economics and finance. Routledge, 2006. ISBN 9780415701952.

S.P. Boyd and L. Vandenberghe. *Convex Optimization*. Berichte über verteilte messysteme. Cambridge University Press, 2004. ISBN 9780521833783.

Margaret Bray. Learning, estimation, and the stability of rational expectations. *Journal of economic theory*, 1982.

William A. Brock and Cars H. Hommes. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control*, 22(8-9):1235–1274, 1998.

William David Anthony Bryant. *General Equilibrium: Theory and Evidence*. World Scientific, 2010. ISBN 9789812818348.

Marco Bunchman. Nearest positive semi-definite covariance matrix, 2011. URL http://www.mathworks.com/matlabcentral/fileexchange/34182-nearest-positive-semi-definite-cov

Richard H. Byrd, Peihuang Lu, Jorge Nocedal, and Ciyou Zhu. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, 16(5):1190–1208, 1995.

John Y Campbell and Luis M Viceira. Strategic Asset Allocation, January 2002.

Nicholas T Chan, Blake LeBaron, and Andrew W. Lo. Agent-Based Models of Financial Markets: A Comparison with Experimental Markets. Technical report, Massachusetts Institute of Technology, 1999.

Myong-Hun Chang and Joseph E Harrington. Agent-Based Models of Organizations. In Leigh Tesfatsion and Kenneth L. Judd, editors, *Handbook of Computational Economics*, volume 2, pages 1273–1300. North Holland, 2006.

R. Cole and L. Fleischer. Fast-converging tatonnement algorithms for the market problem. Technical report, 2007.

Richard H Day and Weihong Huang. Bulls, Bears and Market Sheep. *Journal of Economic Behavior & Organization*, 1990.

John Duffy. Agent-Based Models and Human Subject Experiments. In Leigh Tesfatsion and Kenneth L. Judd, editors, *Handbook of Computational Economics*, volume 2, pages 949–1011. North Holland, 2006.

Samuel Dupernex. Why Might Share Prices Follow a Random Walk? *Student Economic Review*, 2007: 1–13, 2007.

Ronald Fagin, Yoram Moses, Joseph Y. Halpern, and Moshe Y. Vardi. *Reasoning About Knowledge*. MIT Press, 2003. ISBN 9780262562003.

E.F. Fama. Efficient capital markets: II. *Journal of finance*, pages 1575–1617, 1991.

Eugene F Fama. Efficient Capital Markets: A Review of Theory and Empirical Work. *The Journal of Finance*, 25(2), May 1970.

Lisa Fleischer, Rahul Garg, Sanjiv Kapoor, Rohit Khandekar, and Amin Saberi. A fast and simple algorithm for computing market equilibria. *Internet and Network Economics*, pages 19–30, 2008.

Rahul Garg and Sanjiv Kapoor. Auction algorithms for market equilibrium. *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 511–518, 2004.

Christophre Georges. Staggered updating in an artificial financial market. *Journal of Economic Dynamics and Control*, 32(9):2809–2825, 2008.

Nikolay Gertchev. A Critique of Adaptive and Rational Expectations. *The Quarterly Journal of Austrian Economics*, 10(4):313–329, November 2007.

Gerd Gigerenzer and Reinhard Selten. *Bounded Rationality: The Adaptive Toolbox*. MIT Press, reprint edition, 2002.

Benjamin Graham and Jason Zweig. *The intelligent investor*. Collins Business Essentials. HarperBusiness Essentials, 2003. ISBN 9780060555665.

C.W.J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: Journal of the Econometric Society*, pages 424–438, 1969.

Sanford J Grossman and Joseph E Stiglitz. On the Impossibility of Informationally Efficient Markets. *The American Economic Review*, 1980.

Nachi Gupta, Raphael Hauser, and Neil Johnson. Using Artificial Market Models to Forecast Financial Time-Series. *arXiv.org*, 2005.

James D. Hamilton. *Time Series Analysis*. Princeton University Press, 1994. ISBN 9780691042893.

John Henry Holland. *Adaptation 4in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*. University of Michigan Press, 1975. ISBN 9780472084609.

Sean Holly and Martin Weale. *Econometric Modelling: Techniques and Applications*. National Institute of Economic and Social Research. Cambridge University Press, 2000. ISBN 9780521650694.

Cars Hommes. Financial Markets as Nonlinear Adaptive Evolutionary Systems. *Quantitative Finance*, 1: 149–167, 2001.

Cars Hommes. Heterogeneous Agent Models in Economics and Finance. In Leigh Tesfatsion and Kenneth L. Judd, editors, *Handbook of Computational Economics*, volume 2, pages 1110–1146. 2006.

Nicholas Jennings. Agent-oriented software engineering. In Francisco Garijo and Magnus Boman, editors, *Multi-Agent System Engineering*, volume 1647 of *Lecture Notes in Computer Science*, pages 1–7. Springer Berlin / Heidelberg, 1999. ISBN 978-3-540-66281-5.

Kazuya Kamiya. A Globally Stable Price Adjustment Process. *Econometrica*, 58:1481–1485, 1990.

Donald W. Katzner. *The Walrasian Vision of the Microeconomy: An Elementary Exposition of the Structure of Modern General Equilibrium Theory*. University of Michigan Press, 1989. ISBN 9780472064090.

Alan P Kirman. Whom or What Does the Representative Individual Represent? *Journal of Economic Perspectives*, 6(2):117–136, July 1992.

Blake LeBaron. Evolution and Time Horizons in an Agent Based Stock Market. *Macroeconomic Dynamics*, 5(2):225–254, 2001a.

Blake LeBaron. Empirical Regularities From Interacting Long and Short-Memory Investors in an Agent-Based Stock Market. *IEEE Transactions on Evolutionary Computation*, 5:442–455, 2001b.

Blake LeBaron. A Builder's Guide to Agent Based Financial Markets. *Quantitative Finance*, 1(2):254–261, 2001c.

Blake LeBaron. Building the Santa Fe artificial stock market. *Physica A*, 2002.

Blake LeBaron. Agent-based Computational Finance. In Leigh Tesfatsion and Kenneth L. Judd, editors, *Handbook of Computational Economics*, volume 2, pages 1188–1227. 2006.

Blake LeBaron. Heterogeneous Gain Learning and the Dynamics of Asset Prices. Technical report, International Business School, Brandeis University, Waltham, MA, 2010.

Blake LeBaron. Active and passive learning in agent-based financial markets. *Eastern Economic Journal*, 37(1):35–43, 2011.

Blake LeBaron, W. Brian Arthur, and Richard Palmer. Time series properties of an artificial stock market. *Journal of Economic Dynamics and ...*, 1999.

Martin Lettau. Explaining the facts with adaptive agents: The case of mutual fund flows. *Journal of Economic Dynamics and Control*, 1997.

David G. Luenberger. *Investment Science*. Oxford University Press, 1998. ISBN 9780195108095.

Burton G Malkiel. The Efficient Market Hypothesis and its Critics. *Not Published*, April 2003.

Benoit Mandelbrot. The variation of certain speculative prices. *The journal of business*, 1963.

John F Muth. Rational Expectations and the Theory of Price Movements. *Econometrica*, 29, July 1961.

J.A. Nelder and R. Mead. A simplex method for function minimization. *The computer journal*, 7(4):308–313, 1965.

Eva Tardos Noam Nisan, Tim Roughgarden and Vijay V. Vazirani. Cambridge University Press, 2007a.

Eva Tardos Noam Nisan, Tim Roughgarden and Vijay V. Vazirani. Cambridge University Press, 2007b.

J. Nocedal and S.J. Wright. *Numerical Optimization*. Springer Series in Operations Research. Springer, 1999. ISBN 9780387987934.

Richard G Palmer, W. Brian Arthur, John H Holland, and Paul Taylor. Artificial economic life: a simple model of a stockmarket. *Physica D: Nonlinear . . .* , 1994.

Robert S. Pindyck and Daniel L. Rubinfeld. *Econometric Models and Economic Forecasts*. McGraw-Hill international editions. Irwin/McGraw-Hill, 1998. ISBN 9780071158367.

MJD Powell. A view of algorithms for optimization without derivatives. *Mathematics Today-Bulletin of the Institute of Mathematics and its Applications*, 43(5):170–174, 2007.

Rory Radner. Rational Expectations equilibrium: Generic existence and the information revealed by prices. *Econometrica*, 47, May 1979.

Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall series in artificial intelligence. Prentice Hall, 3rd edition, 2010. ISBN 9780136042594.

Herbert Scarf. On the Computation of Equilibrium Prices. In *Cowles Foundation discussion paper*. Cowles Foundation for Research in Economics at Yale University, 1967.

Herbert Scarf. The Computation of Equilibrium Prices: An Exposition. In *Cowles Foundation discussion paper*. Cowles Foundation for Research in Economics at Yale University, 1977.

Herbert Scarf and Terje Hansen. *The Computation of Economic Equilibria (Cowles Foundation Monograph Series No. 24) (Cowles Foundation Monograph Series No. 24 Illus)*. Yale Univ Pr, first edition edition, June 1973.

Anil K Seth. A MATLAB toolbox for Granger causal connectivity analysis. *Journal of Neuroscience Methods*, 186(2):262–273, February 2010.

Steven M. Sheffrin. *Rational Expectations*. Cambridge surveys of economic literature. Cambridge University Press, 1996. ISBN 9780521479394.

J.B. Shoven and J. Whalley. *Applying General Equilibrium*. Cambridge surveys of economic literature. Cambridge University Press, 1992. ISBN 9780521319867.

Herbert A Simon. Rational Decision Making in Business Organizations. *The American Economic Review*, 69(4):493–513, September 1979.

Steve Smale. A Convergent Process of Price Adjustment and Global Newton methods. *Journal of Mathematical Economics*, 3(2):107–120, 1976.

Ross M. Starr. *General Equilibrium Theory: An Introduction*. Cambridge University Press, 1997. ISBN 9780521564731.

Nassim Nicholas Taleb. *Fooled by Randomness: The Hidden Role of Chance in Life and in the Markets*. Random House Publishing Group, 2008. ISBN 9781400067930.

Nassim Nicholas Taleb. *The Black Swan: The Impact of the Highly Improbable*. Random House Trade Paperbacks, 2010. ISBN 9780812973815.

Leigh Tesfatsion. Agent-Based Computational Economics: A Constructive Approach to Economic Theory. In Leigh Tesfatsion and Kenneth L. Judd, editors, *Handbook of Computational Economics*, volume 2, pages 831–880. 2006.

Nocolaas J. Vriend. ACE Models of Endogenous Interactions. In Leigh Tesfatsion and Kenneth L. Judd, editors, *Handbook of Computational Economics*, volume 2, pages 1047–1079. North Holland, 2006.

Léon Walras. *Elements of Pure Economics: Or, the Theory of Social Wealth.* Routledge library editions: Economics. Taylor and Francis, 2003. ISBN 9780415313407.

Frank Westerhoff. A simple agent-based financial market model: direct interactions and comparisons of trading profits. *Nonlinear Dynamics in Economics*, 2009.

Frank H Westerhoff. Multi-Asset Market Dynamics. *Macroeconomic Dynamics*, (8):596–616, November 2004.

Allen Wilhite. Economic Activity on Fixed Networks. In Leigh Tesfatsion and Kenneth L. Judd, editors, *Handbook of Computational Economics*, volume 2, pages 1013–1045. North Holland, 2006.