1 Financial markets as complex systems

~~HANDOUT 2 OF 7~~

1.1 Real problems in finance

Suppose we have an interest in a dot-com company *risk-e.com* that has been around on the stock market for a few years. The 'we' could be: private investors who hardly ever buy stocks, on-line traders who trade several times a day for a living, pension-fund managers looking to the longer term, finance officers in a multi-national company looking to speculate with some extra funds, or traders in a commercial bank looking to either balance portfolio risk for customers or carry out proprietory trading. We all have different time-horizons in mind, different capital, different ease-of-access to the market, different levels of transaction costs, different minimum-execution times, and different perceptions and definitions of potential risk. But we all have one thing in common: *We certainly don't want to lose any money.* After all, if we expected to lose money in the market then we should either put our cash in a bank account, or stick it under the bed. Suppose *risk-e.com* had its Initial Public Offering (IPO) in January 1996, hence there is some limited history available of past performance. The history of the stock-price to date is shown below:

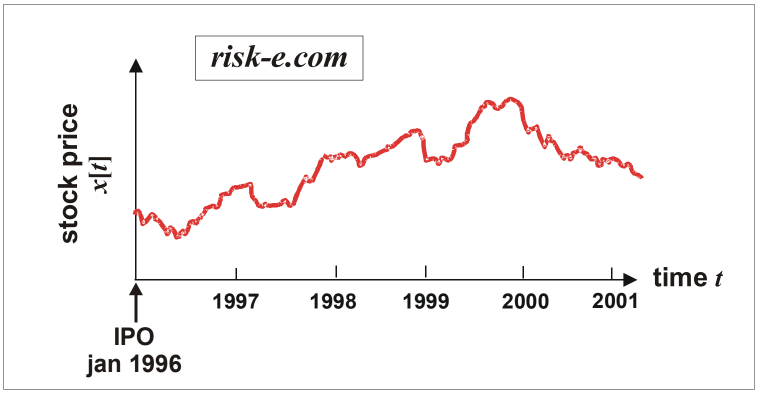

Figure 1-1 Past price-series for the dot-com company *risk-e.com*

We may not just be looking to buy stock in the hope that the price will go up, although it would be nice if it did of course. We might also be wanting to assess future risk, hedge risk, or compare performance to that of other supposedly similar companies. While we could, and should, check out the company's earnings and growth reports, we also know that this information is limited. And anyway, according to standard finance theory, since this information is essentially public surely it will already have been incorporated into the current price? While pondering the situation and staring at the price-chart, we can't help but noticing the roller-coaster ride that the stock seems to have had since the IPO. Questions start popping into our heads:

Does any of that roller-coaster ride in the stock-price $x[t]$ actually tell us anything? If so, what? Should we buy (i.e. go long) the stock? Or sell it (i.e. go short)? Maybe we want to use this stock to hedge our risk in other dot-com companies, or in other technology sectors, or in another market. What should we then do to minimize our risk? What do we each mean by 'risk' – risk of what exactly? Is there any predictability in the stock's behaviour? How, and to what extent, would such predictability manifest itself on the hourly, daily, or monthly scale? Is the roller-coaster ride driven by crowd behaviour? Can we infer what the crowd is thinking? If so, might we then be able to forecast the future based not on the past price series, but on what we believe the crowd will do?

These are a sampling of the types of practical questions that anyone interested in the markets faces regularly, from large investment institutions through to the private investor. These questions are however not the types of questions that are easily answered, or even easy to address, within standard finance theory.Why? Embedded in each of these questions is the issue of *time*, and timing of decisions or actions. Hence in order to really address these questions, we will need to understand something about the time-evolution of the particular asset's price, and probably also the market in general. But surely it is impossible to say exactly what will happen to a given asset or market in the future? Yes, of course -- but even a limited quantitative description might be useful, assuming that we understand what these limitations are. The big problem is that it is very difficult to make even reasonably accurate statements about the future time-evolution of such a system, because that time-evolution is generally so complicated. In fact it would be hard to find anyone who disagreed with the statement that a financial market is indeed a 'complex' system. This underlies why standard theory falls short: it has nothing to do with the ability of finance theorists, but instead has to do with the fact that nobody yet knows how to describe mathematically the time-evolution of such complex systems in a general yet useful way. And herein lie the academic and practical motivations for this book: across a broad range of disciplines, researchers are now realizing that some of the hardest problems that they each face have key common elements. These elements are the key elements of what is now being called a 'complex' system. So in order to understand the claim that a financial market is also such a complex system, we will spend a few moments looking at what these key ingredients are.

1.2 Complex systems and Complexity

When we say 'financial markets are complex systems', we mean more than just 'complicated systems'. Making a pizza or fixing a bike puncture are both 'complicated', but neither is 'complex'. However put these tasks together, and let the next step in one task depend on the present state of the other, and you start to incorporate at least a glimmer of complexity. Although there is no universally accepted definition of 'complexity' or 'complex system', most people would agree that any candidate complex system should have most or all of the following ingredients:

- **Feedback**. The nature of the feedback can change with time -- for example, becoming positive one moment and negative the next -- and may also change in magnitude and importance. It may operate at the macroscopic or microscopic level, or both. The presence of feedback implies that on some level, buried in the details of the dynamics, the system is 'remembering' its past and responding to it, albeit in a highly non-trivial way.
- **Non-stationarity**. We cannot assume that the dynamical or statistical properties observed in the system's past, will remain unchanged in the system's future[1].
- **Many interacting agents**. The system contains many components or participants, known as 'agents', which interact in possibly time-dependent ways. Their individual behaviour will respond to the feedback of information, which is possibly limited, from the system as a whole and/or from other agents. Since these agents may effectively be competing to win, it is unlikely that there is any such thing as a 'typical' agent.
- **Adaptation**. An agent can adapt its behaviour in the hope of improving its performance.
- **Evolution**. The entire multi-agent population evolves, driven by an ecology of agents who interact and adapt under the influence of feedback. The system typically remains far from equilibrium, and hence can exhibit 'extreme behaviour'[2].
- **Single realization**. The system under study is a single realization, implying that standard techniques whereby averages over time are equated to averages over ensembles, may not work.
- **Open system**. The system is coupled to the environment, hence it is hard to distinguish between exogenous (i.e. outside) and endogenous (i.e. internal, self-generated) effects.

Most, if not all, these criteria are applicable to a financial market, implying that a financial market can be thought of as a complex system. The market price-series provide a record of the system's global dynamics, while the market's participants (traders) represent the agents. Given the above criteria, it is often said that such complex systems are *more than the sum of their parts*. In other words, just as we say 'two's company, but three's a crowd' to denote how a collection of N people changes its character and behaviour as N increases from 2 to 3 or more, so a complex system exhibits so-called emergent properties which could not be easily predicted based on the behaviour of the individual constituent parts. A crowd of people is a good example of an emergent property. Certainly there are many, often tragic, examples in history of the power of a crowd. In fact the connection to financial markets is quite

---

[1] In financial calculations for example, historic volatility levels obtained from past price-series may be poor estimates of future volatility. In a complex system, statistical niceties such as stationarity cannot be taken for granted. As suggested by the disclaimer in financial advertisements, past performance provides no guarantee as to future success.
[2] The financial example of such extreme behaviour, is a market crash. See Chapters 3, 4, 5 and 7.

direct, since it is often claimed that crowd behaviour causes crashes. This topic will be discussed in more detail in Chapters 4, 5 and 7. For now, having dealt with the idea of complexity, we spend the rest of this Chapter reviewing the nuts-and-bolts of a financial market (see, for example, [P] for an in-depth account).

1.3 Financial market overview

1.3.1 The role of financial centres

Financial centres are places where the demand for financial services, both within the domestic and international communities, is met. Examples of financial centres worldwide include London, New York and Tokyo. Financial centres increasingly find themselves competing in a global marketplace both to retain their domestic market and to win international business. Governments seek to promote their financial centres, not only because of the influx of substantial amounts of foreign capital but also because they provide employment for vast numbers of people. Effective financial centres can also help the economy by channelling capital into investments with high returns. In general, the most important role of a financial centre is to transfer funds or goods from agents with a surplus to agents with a deficit, in the most efficient way possible. In order for such transfer to be efficient, the financial centre must offer a diverse range of products and services to meet the diverse range of borrowers' and investors' needs. Among the services most in demand are: primary and secondary markets in bonds and equities, foreign exchange, risk management, derivative products, domestic and international bank lending.

1.3.2 Types of financial market

A distinction is made between primary and secondary markets: a primary market deals in issues of new assets whereas in a secondary market, existing assets are traded. In a primary market the issuer of the asset benefits directly from the capital raised by its sale. In a secondary market the issuer does not receive any proceeds from the resale of the asset. The secondary market is nonetheless important to the original issuer. The price of the issuer's asset in the secondary market reflects how willing financial agents are to buy the asset and hence how much capital could be raised by the original issuer through a further new sale on the primary market. A secondary market also provides what is known as liquidity. Liquidity is essentially the freedom to transact assets. Without a liquid secondary market, investors would not be willing to pay as much for new assets on the primary market since they know it would be difficult to sell the assets again. Vital to the provision of liquidity in the secondary market is the presence of so-called 'market-makers'. A market-maker will quote buy and sell prices for assets and be willing to accept large trades in either direction in response to market supply and demand.

Consequently it is usually easy for a buyer or seller of the asset to find a counter-party for the trade, thus allowing high 'volumes' of transactions and increased liquidity. We will henceforth only consider secondary financial markets, since their dynamics are a direct result of these interactions between buyers and sellers. Not only is there competition between buyers of the asset, but also between sellers. The nature of this two-sided competition will be of great importance in casting the financial market into the form of a general complex system. We henceforth refer to secondary financial markets simply as 'financial markets'.

Secondary markets can have diverse forms in terms of both their macroscopic character and microscopic structure. In 'screen-based markets' trading takes place electronically through a (possibly geographically dispersed) IT infrastructure. In a 'call market' orders are batched together at infrequent intervals (just once a day in some markets) and a market price is decided by an auction process, either oral or written. In a 'continuous market' prices are quoted continuously by market-makers throughout the trading day. Some markets have a mixture of systems.

1.3.3 Financial assets

A financial asset is simply a legal claim to a future cash flow. Financial assets are interchangeably referred to as financial securities, financial instruments, financial products or financial claims. The issuer of the financial asset undertakes a legal agreement to make cash payments to the holder (interchangeably referred to also as the investor) of the asset in the future. Financial assets, and hence the financial markets that trade them, are named according to the details of the contract for future cash flow. Broadly speaking, financial assets can be classified according to the following criteria.

1.3.3.1 Debt, equity and foreign exchange

With debt claims, the holder has two pre-determined sources of future cash flow. At the 'maturity time' of the contract, the issuer must deliver back the original value of the claim. Also, at regular fixed times until maturity, the issuer must pay the holder interest on the loan – this interest may be fixed or variable. Typically, holding a debt claim will be lower risk than holding an equity claim since the debt claim will usually be secured against the assets of the issuer. For this reason debt instruments are often referred to as 'fixed-income' since the future cash flow is more assured than for equity instruments. Examples of debt instruments are: government bonds or bills, corporate bonds, mortgages and bank loans. Debt instruments with less than one year maturity are known as 'bills' and are traded on 'money markets' whereas debt instruments with more than a year's maturity are known as 'bonds' and are traded on 'capital markets'. Cash and foreign exchange can be viewed as debt instruments issued by the central bank of the country in question. For this reason they are viewed as the simplest and most secure financial assets available, since repayment is (almost always!) definite and the interest rate is fixed at zero.

With equity claims, the holder also has two sources of future cash flow, but neither is guaranteed in terms of the actual amount. First, the holder has the right to a regular dividend payment which is paid once the holders of all the issuer's debt claims have been paid. Second, the holder has the right to a portion of the value of the issuing institution should it be sold or placed in liquidation. Of course the holder of an equity claim has no guarantee that there will be *any* future cash flow. Equity claims are thus more risky instruments than debt claims since the future cash flow will depend on the issuing institution's financial success. By far the most prevalent form of equity instruments are ordinary company shares (common stock).

1.3.3.2 Time of settlement

If the details of the financial asset contract stipulate that settlement should be made at the time the contract is agreed between the two parties, then the asset is a 'cash (or spot) asset' and as such is traded on a 'cash (or spot) market'. If however the contract is settled at some time in the future from the time it was agreed, the asset is known as a 'forward contract' or 'future contract'. Forward and future contracts can be made on a wide range of financial asset types such as foreign exchange, stock, bonds and even indices (i.e. assets whose value depends on the collective value of a basket of other assets). A forward contract is usually an agreement forged between two distinct parties, hence the contracts cannot easily be sold on to a third party. This is in contrast to futures contracts which are usually for set contract sizes, based on set delivery dates and sold by a 'futures exchange'. Futures contracts can be sold from party to party and do not usually end in the delivery of the 'underlying asset' (i.e. the currency, stock or bond). They are used by investors more for speculation and risk management than for future delivery of the underlying asset. Because the delivery of an underlying asset is secondary to a futures contract, it is classed as a 'derivative instrument' since its value is derived from an underlying asset.

1.3.3.3 Obligation to exchange

For some financial assets, the contract has an associated obligation to deliver another product immediately, or at some time in the future. For example, a bond contract has associated with it the obligation for the issuer to pay to the holder the original bond value at maturity, plus interest at regular intervals. For a currency forward, there is the obligation that the issuer must deliver the holder a quantity of currency at a given future time. However, some financial assets simply give the holder the *right*, but not the *obligation*, to exchange one product for another at a certain time. An example of this financial asset type is the 'option'. The holder of an option contract has the right but not the obligation to buy the 'underlying asset' for a given price (the strike price). As such, an option contract is a derivative instrument just as for the future contract described above.

1.3.4 Financial market agents

The participants who are active in any given financial market, can be loosely split into those whose job it is to provide the service and those who seek to use the service. The main providers of the service will be the brokers, market-makers and regulators. The main users of the service form a wide group and may include: individuals, investment and commercial banks, investment companies, insurance and pension funds, businesses, local and central governments, and international institutions such as the World Bank.

1.3.4.1 Market service providers

'Brokers' act as the legal agents of investors, and as such can be more closely regulated than the investors themselves. A broker is the intermediary between the investor and the financial assets. The brokers not only offer a service for the buying and selling of financial assets, but may also offer relevant market news, research reports and custody of the asset. Brokers make their money by charging a commission for their services. 'Market-makers' act as dealers for particular financial assets. Market-makers will offer a 'bid price' at which they are prepared to buy the financial asset from the investors (or other market-makers) and an 'ask price' at which they are prepared to sell the financial asset. The difference between the ask and bid price is known as the 'spread' and represents the profit margin of the market-maker. The market-maker's job is to aid in the efficient running of the market by offering competitive prices for the financial asset, and also to act as a source of liquidity. 'Market regulators' aim to guarantee the legality and correct execution of all market interactions. Further to this role, regulators may try and ensure a degree of stability in the market. For example, if the movement in a particular financial asset (or the market as a whole) exceeds certain bounds, the regulator may seek to intervene and halt trading in that asset (or the whole market) in the hope that subsequent trading will be calmer.

1.3.4.2 Market service users

Although market users are diverse in size, time horizon and objective, they can usefully be divided into three main groups depending on their motivations and preferences for return on their investment versus risk. The three main groups are:

- Speculators: These are market agents who are aiming for return on their financial investment above all else. Speculators will have a view of the future evolution in price of a financial asset and will consequently buy the asset if they decide it is under-valued with respect to their beliefs, and sell the asset if they decide it is overvalued.
- Hedgers: These are market agents who choose to participate in the market in order to reduce the risk of their uncertain portfolio value. These agents are consequently not as concerned with absolute return as the speculators, but instead choose to buy or sell financial assets whose movement in price they believe will dampen out the movement in value of their existing portfolio.
- Arbitrageurs: These are market agents who are concerned not only with attaining a high return from their investments, but also doing so with minimal (or zero) risk. Arbitrage has a specific meaning within the context of finance: it is the process of exploiting mis-pricing of financial assets in order to gain riskless profit. Suppose an asset has price X_A on market *A* but price $X_B < X_A$ on market *B* -- the arbitrageur would, at the same point in time, buy the asset on market *B* and sell it on market *A*. The arbitrageur thus makes $X_A - X_B$ profit without risk. The very presence of arbitrageurs implies that the opportunities for arbitrage are scarce, since the action of buying on market *B* forces the price X_B upwards whereas selling on market *A* forces the price X_A downwards until $X_A = X_B$. This is known as the principle of no arbitrage[3].

1.3.5 The price of an asset

One of the most important roles of a financial market is to supply, on a continuous basis, a price for a financial asset at which both buyers and sellers are willing to trade. Classically the 'value' of a financial asset is the current value of the total expected future cash flow from that asset. This is the 'rational expectations' price of the asset and there are many models in the economics and finance literature for calculating this price (see for example [CLM]). However, it is an obvious fact that the price of every financial asset moves on a day to day, hour to hour and often second to second basis and that usually there is very little relation between the mean price of the asset and the 'rational expectations' price. The question therefore arises as to what determines the price of a financial asset in practice.

1.3.5.1 Role of the market-maker

We have already discussed the important role of the market-maker in the process of setting a price level for financial assets at which buyers and sellers are willing to trade. The market-makers supply a

[3] M. O'Hara, *Market Microstructure Theory* (Blackwell Publishers, 1995).

willing counterparty for each trade, at a price set by themselves. In the simplest scenario a market-maker, by contrast to the other market participants, will have no interest in generating a financial return from holding a position in the financial asset and then raising its price. Ideally, the market-maker's manipulation of the price should be solely for the purposes of matching the supply of assets from willing sellers, with the demand for assets from willing buyers. In this way, market-makers can maximize the number of assets traded (i.e. the volume) hence generating market liquidity, and at the same time maximizing their own profit. (Recall that market-makers set a spread between bid and ask prices which generates revenue for them from each traded asset). It is a fairly undisputed fact[4] that *on average*, the market-maker will raise the asset price in the presence of an excess of buyers, and will lower the asset price in the presence of an excess of sellers. This follows from the assumption that fewer agents will be willing to buy the asset if the ask price is raised, though more agents will be willing to sell it if the bid price is raised as well. Although this assumption is not necessarily always true, it serves as a good approximation. Hence we have the scenario wherein the demand for assets largely determines the traded price through the action of the market-maker. In order to pursue our goal of understanding what determines the price level of financial assets in the market, we must therefore answer the question: 'what determines the demand for assets?'

1.3.5.2 Demand for assets

We will refer to the number of financial assets sought minus the number offered, as the 'excess demand' for the asset. The excess demand, by construction, will be a very complex function of the investors' beliefs about the asset: in particular, its expected future price levels and its expected future earnings. These two issues are the principal factors that will be in investors' minds when they decide to place an order with their broker to buy or sell the asset. How important each of these issues is to an investor, and how his/her individual perception of these issues arises, will in turn be a function of very many possible parameters. For example, while one investor may view a news bulletin about the performance of company *X* as a deciding influence on the future price movement of company *X*'s assets, another investor may see this news as neutral or may even believe it will have the opposite effect. Consequently any realistic model of the excess demand for a financial asset, and hence a model for the asset-price itself, must capture these possible behavioral issues in a very general way.

1.3.6 Orders and market clearing

Financial market investors who wish to buy or sell assets, do so by contacting their broker who acts as their legal agent in the market. However, there are different types of request an investor can make of

---

[4] See, for example, T. Chordia, R. Roll and A. Subrahmanyam, J. Fin. Econ. **56**, 501 (2001)

his/her broker. Broadly speaking, investors can place an order with their broker to execute a trade irrespective of the price at which the trade is made, or they can place an order which is dependent on the traded price which can be achieved. We refer to the first (unconditional) type of order as a ‘market order’ and the second (conditional) type as a ‘limit order’.The existence of these different types of order arises from the fact that investors cannot immediately execute their desired trade at the current price they observe. It will physically take a certain amount of time for the broker to be contacted and for the trade to be secured. Within this time interval, which is small nowadays but still non-zero, the price may well move away from the value the investor saw at the time the order was placed. In addition, the price observed by the investor at the time of placing the order may not be an accurate representation of the current asset price. The reasons why this might be so are many. For example, the price observed by the investor may be delayed by the data provider, be an average over the last few traded prices, be an average over the bid and ask prices or simply be incorrectly recorded. These factors contribute to create the situation whereby an investor may not have accurate knowledge of the traded price. However, there is a third and more fundamental reason why the actual traded price can be different from the asset price observed at the time of the order. This is due to the mechanism by which the price is set in the market. When the order is placed by the broker, it enters the market-maker’s ‘order book’. The market-maker must then update his/her calculation of the excess demand in the market, by adding up all the buy orders in the order book and subtracting all the sell orders. Due to the revised calculation of the excess demand, the market-maker may then consequently wish to move the price of the asset. All this will be done *prior* to the execution of the orders. *The temporal ordering of this entire process is central to the task of constructing a useful model of the complex system represented by the financial market*. We re-iterate this temporal ordering in the schematic representation shown below:

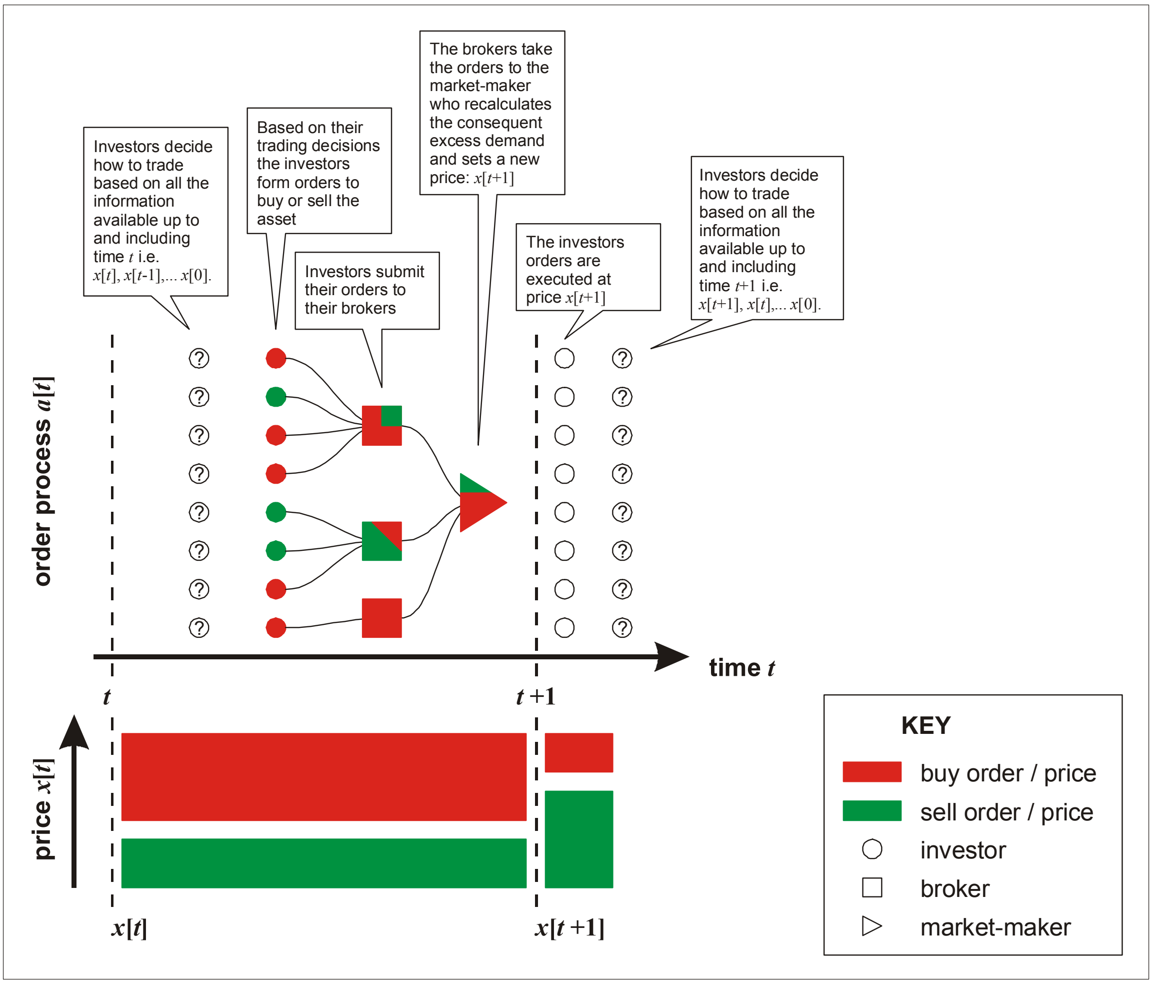

Figure 1-2 A schematic representation of the temporal order in which events occur in the trading of financial assets through the investor--broker--market-maker chain. For simplicity we assume all orders are market orders and thus are all transacted at the new price set by the market-maker. We also have assumed that all orders are placed at the same time; this is purely for clarity as described in Chapter 4.

1.3.6.1 Market impact

The mechanism of the market-maker moving the asset price in response to the excess demand *prior* to executing the orders, leads to a very important feature of the complex system representing the financial market: this feature is termed 'market impact'. Broadly speaking, market impact is the adverse affect of an order on its own traded price. To recap, the market-maker moves the price up in the presence of a positive excess demand, and down in the presence of a negative excess demand. Let us imagine a scenario where the order a_i placed by one investor *i* is statistically independent from the order a_j placed by another agent $j \neq i$. The act of agent *i* placing an order to buy assets ($a_i > 0$) will, when we average over all the possible orders from agents $j \neq i$, increase the excess demand $D = \sum_k a_k$. The opposite is also true, i.e. agent *i* placing an order to sell assets ($a_i < 0$) will on average lower the excess demand. The consequence of this is that an order to buy assets will on average raise the price at which they can be bought, and an order to sell assets will on average lower the price at which they can be sold. This is the market impact effect.

1.3.6.2 Clearing the market

We have discussed the fact that orders can be placed by an investor in two distinct ways, either as 'market orders' or 'limit orders'. Whether the order is a market or limit order determines how it is treated by the market-maker and hence how the order itself affects the market. A general framework for the way orders are placed with the market-maker is as follows. An order a_i (whether it be a market or limit order) represents the willingness of an investor *i* to either buy ($a_i > 0$) or sell ($a_i < 0$) a quantity $|a_i|$ of financial assets. In general we can say[5] that the order at time *t*, which is $a_i[t]$, is a function of the price at which the order will be executed, which is $x[t+1]$. Hence the general order $a_i[t, x[t+1]]$ can now be defined mathematically in terms of the preferences of the investor. For example, a simple market order that says 'transact at any traded price' has the form $a_i[t, x[t+1]] = a_i[t]$, and a simple limit order that says 'transact only if the traded price is equal to or greater than value y_i' has the form $a_i[t, x[t+1]] = a_i[t]\,\mathrm{H}[x[t+1] - y_i]$ where $\mathrm{H}[x]$ is the Heaviside function[6]. In most circumstances the market is dominated by market orders together with simple limit orders of the form 'buy if the traded price is lower than y_i' and 'sell if the traded price is higher than y_i'. The market-maker simultaneously defines a bid price x^- and an ask price x^+, where of course $x^- < x^+$ to prevent arbitrage. After collecting all the orders $a_i[t]$, re-assessing the excess demand $D[(t+1)^-] = \sum_i a_i[t]$ (where $(t+1)^-$ indicates the time immediately before time $t+1$) and hence moving the bid and ask prices ($x^-[t+1]$ and $x^+[t+1]$ respectively), the market-maker will execute:

- all market orders
- all buy orders ($a_i[t, x[t+1]] > 0$) for which $y_i \geq x^+[t+1]$
- all sell orders ($a_i[t, x[t+1]] < 0$) for which $y_i \leq x^-[t+1]$

The market-maker's book will then contain buy orders for which $y_i < x^+[t+1]$ and sell orders for which $y_i > x^-[t+1]$. Some of these may lie within the market-maker's spread, that is $x^-[t+1] < y_i < x^+[t+1]$. Although these orders were unacceptable to the market-maker, they can be filled by other agents with opposite orders within the spread. This trading inside the spread tends to remove all orders in this price region, leaving a collection of buy orders for which $y_i < x^-[t+1]$ and

[5] Units of time are arbitary here, and reflect the discrete nature of the process of placing an order and having it executed, as demonstrated schematically in Figure 1-2.
[6] $\mathrm{H}[x]=1, 0$ for $x\geq 0, x<0$ respectively.

sell orders for which $y_i > x^+[t+1]$. An example of such an order book is shown in Figure 1-3, and graphically in Figure 1-4:

Last Trade 23.300			Time 07:31	Shares 100	
Buy Orders			**Sell Orders**		
Broker	**Shares**	**Price**	**Broker**	**Shares**	**Price**
A	800	23.299	C	2900	23.300
B	100	23.290	B	400	23.340
C	4000	23.280	A	100	23.360
B	1000	23.221	A	800	23.379
B	1100	23.220	B	500	23.400
B	1000	23.200	C	4000	23.410
B	1000	23.200	B	4000	23.410
A	500	23.120	A	2000	23.440
B	100	23.110	B	100	23.590
A	2100	23.040	B	500	23.600
A	4900	23.040	B	100	23.690
A	100	23.040	B	200	23.730
B	400	23.010	A	2000	23.900
A	100	22.920	B	70	23.980
B	100	22.910	B	54	23.990
12 more			26 more		

Figure 1-3 An example of a market-maker's order book showing the (limit) orders remaining after all possible trades have been executed. Data taken from http://www.3Dstockcharts.com shows market activity on the financial index asset QQQ on 13/07/2002.

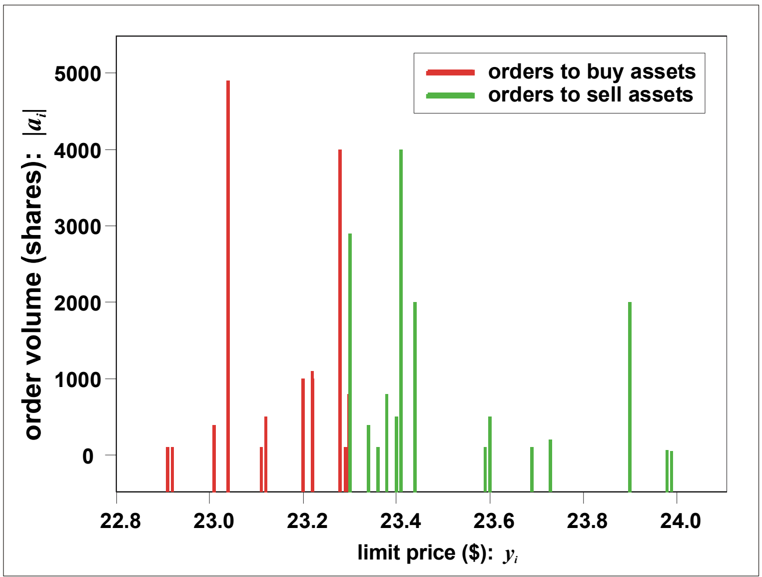

Figure 1-4 A graphical representation of the order book shown in Figure 1-3

The (limit) orders which still remain on the market-maker's order book after all possible trades have been executed, will remain there until the investors wish to withdraw them. As such, the orders remaining on the market-maker’s order book still contribute to the excess demand. One can therefore look upon charts such as Figure 1-4 as showing the ‘pressure’ on an asset price.

So far it has been assumed that the market maker can, and always will, process all possible orders. This in turn implies that the market-maker may have to assume either a long position (holding assets) or a short position (owing assets). The position held by the market-maker represents an unwanted risk; the market-maker, as a pure liquidity provider, should not seek to take a speculatory position in the asset. It is therefore likely that the market-maker will also take his/her own position in the asset into account when deciding a new price level to set. This effect can further complicate the relationship between excess demand and price change[7].

1.3.7 Chartism versus fundamentalism

Section 1.3.5.2 discussed how demand for assets arises from investors' expectations of either the future earnings of the asset, or the future price movements. It could be argued that the future movements in the price of the asset reflect movements in its value which are in turn determined by its future earnings. Hence the two sources of demand might appear to be the same. However, it is an empirical fact that in virtually all cases, the frequency of movements in the price of a financial asset is orders of magnitude greater than the frequency with which news arrives which could relate to the future earnings of that asset. In addition, a large proportion of all asset positions are held for only a short period of time, so short in most cases as to never benefit from the asset's direct earnings (dividends). This seems to indicate instead that the two sources of demand are actually fundamentally different. Furthermore, the changes in investors' expectations of future price movements could indeed generate a demand for assets which fluctuates with the high frequency that is actually observed. It therefore seems reasonable to consider the principal source of demand as the investors' expectations of future price movements. Broadly speaking, expectations of future price movements can be generated through two distinct schools of thought: chartist principles and fundamentalist principles. A chartist bases his/her expectations of future price movements on the past states of a series of 'technical indicators'. These technical indicators[8] are usually functions of the past series of asset prices $x[0],\ldots,x[t]$, and sometimes also the volumes $V[0],\ldots,V[t]$. A fundamentalist on the other hand bases his/her expectations of future price movements on the past states of a series of 'fundamentals'. The fundamentals of an asset are usually parameters relating to the asset's earnings and the economic strength of the institution issuing the asset. Most investors belong to either the chartist or the fundamentalist school of thought with very few employing both principles. It is an ongoing question, and often the cause of much angst, as to whether an investor should trust chartist principles or fundamentalist principles in any given market. A principle of economics known as the Efficient Market Hypothesis (see for example [CLM]) claims that all the information concerning an asset -- technical and fundamental, public and private -- is already incorporated into the current price of the asset. If this principle were true, it would imply that it is theoretically impossible for any investor from

[7] P. Jefferies, M.L. Hart, P.M. Hui and N.F. Johnson, Eur. Phys. J. B **20**, 493 (2001).
[8] See for example, A. Blair, *Investors Chronicle guide to charting* (Pitman, 1996).

any school of thought to make a risk-free profit by following their own trading principles. What one can however say in the debate of chartism versus fundamentalism, is that fundamentalism is considered the 'old' school of thought and chartism the 'new'. Consequently the more modern, liquid, markets seem to be dominated by chartist investors[9]. This is largely due to a belief amongst investors that the more modern, liquid and speculative markets are the most fertile grounds for applying chartist principles. This clearly leads to a vicious circle which draws more and more chartists into the market as the degree of liquidity and speculation grows. There then arises the following interesting problem: if there is a population of traders *looking* for patterns in past prices, *believing* they see such patterns[10], and then *acting* as though those patterns were real, will this process itself induce patterns into the market? Since there is no unique interpretation of such price series in terms of pattern identification, the market of chartists becomes a heterogeneous one. Will some patterns therefore survive, hidden away safely since no chartist happens to be around who could identify them? And what happens in the opposite case where too many chartists are identifying, and acting on, the same apparent pattern? These questions provide the motivation for much of the discussion of market models in Chapter 4.

1.4 Observing the market

Having given an overview of the inner workings of financial markets, we now step back and assume the position of an outsider who is observing the 'output' of a particular market. Despite the large number of variables driving the market itself, as discussed in this Chapter, the outside observer has a very limited number of output variables at his disposition. Hence he may be forced to regard the market as a 'black box'. For the stock market, the output variables are prices and possibly volumes of trades. In recent years the frequency of the available data has increased enormously: whereas previously only daily prices were disclosed, now it is possible to get trade-by-trade price data (so-called *tick-data*) albeit with a finite delay time if one doesn't subscribe to a data provider. Since the price as a function of time $x[t]$ is the primary experimental output variable, or 'observable', it is worth discussing it in some detail[11].

---

[9] Research in collaboration with Dr J. James of Bank One, London.

[10] Even a purely random price-series may appear to have patterns, if one is only observing a finite segment of it. The analogy in a coin-tossing game is as follows: what is the next outcome in the sequence ..HHHHHHHH? The answer is of course H (heads) or T (tails) with equal probability if the coin is fair. But it is very tempting, and arguably part of human nature, to believe that it will definitely be H (i.e. trend-following behaviour), or definitely be T (i.e. contrarian behaviour).

[11] We tend to use $y[t]$ or equivalently y_t to denote a financial variable which is defined at discrete times: e.g. the closing price or total trading volume each day. We tend to use $y(t)$ to denote the special case in which the financial variable is treated as a function of continuous time. Note that $y[t]$ may not be smooth and differentiable, and hence may need to be described by mapping equations (see Chapters 4-7) as opposed to the typical differential equation approach of standard finance for $y(t)$ (see Chapter 2). The uses of $y[t]$, y_t and $y(t)$ will be made clear from their context throughout the book.

Consider the typical situation in which we, as outsiders, are provided with a set of high-frequency price-data $\{x\}$ over a given, finite time-period. Even though the major markets are very liquid, each transaction occurs at a specific moment in time. Hence the time-series will be a set of data-points at discrete times $x[t_1], x[t_2], x[t_3]\ldots$, typically separated by irregular clock-time intervals (e.g. 1 secs, 5 secs, 2 secs ….). To add to the complication, many exchanges are not open 24 hours per day -- plus there are holidays and weekends. Even the ones that do remain open will possibly change their 'character' (and hence characteristic price distributions) when one time-zone shuts down and the other opens up, e.g. when New York closes or London opens. This is not only very awkward to handle in practical terms, it also raises some tricky conceptual questions regarding a priori assumptions of stationarity, and opens up the possibility of picking up spurious dynamical trends. Even within one single geographic zone, it is unclear how we should 'connect' the value from 4pm on Friday and 8am the following Monday. Should they be considered to be truly consecutive in the same way as 2:15pm and 2:16 pm are on a given day? Unfortunately there is no unique solution to this time-related problem, and any empirical investigation must be careful to ensure that unintentional correlations are not introduced by whichever approach is chosen.

Another difficult question relates to which *function* of the price we, as observers, should use in order to perform statistical tests. In particular we wish to extract maximum meaningful information and, possibly more importantly, introduce the minimum amount of spurious correlations or bias by this choice. The following choices are typically the most common when looking at price fluctuations:

- Linear price-change:

  $$\Delta x[t, t-\Delta t] = x[t] - x[t-\Delta t] \qquad (1.1)$$

  or equivalently, using shorthand notation in order to reduce the number of parentheses[11]

  $$\Delta x_{t,t-\Delta t} = x_t - x_{t-\Delta t}$$

- Discounted or de-trended price-change: this attempts to remove the effects of inflation or supposedly deterministic bias etc. by introducing a factor $F[t]$

  $$\Delta x^{(F)}[t, t-\Delta t] = F[t]x[t] - F[t-\Delta t]x[t-\Delta t] \qquad (1.2)$$

The problem then arises as to what form $F[t]$ should take. Fortunately if $F[t]$ is slowly varying as compared to the time-window over which the data is collected, then we can usually assume $\Delta x^{(F)}[t, t-\Delta t] \propto \Delta x[t, t-\Delta t]$.

- Return: this focuses on the fractional change in price in order to distinguish the relative importance of a given absolute price-change, e.g. in order to distinguish between the cases of a \$1 change in price for a \$400 stock as compared to a \$1 change for a \$10 stock.

$$R[t,t-\Delta t] = \frac{\Delta x[t,t-\Delta t]}{x[t-\Delta t]} = \frac{(x[t]-x[t-\Delta t])}{x[t-\Delta t]} \tag{1.3}$$

or, using our shorthand equivalent,

$$R_{t,t-\Delta t} = \frac{\Delta x_{t,t-\Delta t}}{x_{t-\Delta t}} = \frac{x_t - x_{t-\Delta t}}{x_{t-\Delta t}} \quad .$$

In the typical case that the change in price is much smaller than the price itself, i.e. $\Delta x[t,t-\Delta t] \ll x[t-\Delta t]$, we can assume $R[t,t-\Delta t] \propto \Delta x[t,t-\Delta t]$.

- Log-return:

$$z[t,t-\Delta t] = \ln\frac{x[t]}{x[t-\Delta t]} = \ln x[t] - \ln x[t-\Delta t] \tag{1.4}$$

or, using our shorthand equivalent,

$$z_{t,t-\Delta t} = \ln\frac{x_t}{x_{t-\Delta t}} = \ln x_t - \ln x_{t-\Delta t} \quad .$$

In the limit that $\Delta x[t,t-\Delta t] \ll x[t-\Delta t]$, we have

$$z[t,t-\Delta t] = \ln\frac{x[t]}{x[t-\Delta t]} = \ln\left(1+\frac{\Delta x[t,t-\Delta t]}{x[t-\Delta t]}\right) \sim \frac{\Delta x[t,t-\Delta t]}{x[t-\Delta t]} = R[t,t-\Delta t]$$

Hence we can again assume $z[t,t-\Delta t] \propto \Delta x[t,t-\Delta t]$.

Each of these price-related definitions has its advantages and disadvantages. Any particular choice may or may not bias a given data set, hence confusing the usefulness of the conclusions. Again this is a practical problem and there is no easy way out, nor is there any ‘right’ and ‘wrong’ answer. Throughout this book we will switch between these quantities according to the data available, and the problem under consideration. Fortunately these quantities all have similar statistical properties in practice, since the change in price $\Delta x[t,t-\Delta t]$ is typically small compared to the absolute values $x[t-\Delta t], x[t]$. This reason, combined with the fact that we do not focus on specific stock data, means that *we will not tend to worry about the distinction between these various price-change measures in this course.*