

Cambridge University Press
978-1-107-09114-6 - Algorithmic and High-Frequency Trading
Álvaro Cartea, Sebastian Jaimungal And José Penalva
Excerpt
[More information](#)

Part I

Microstructure and Empirical Facts

Cambridge University Press
978-1-107-09114-6 - Algorithmic and High-Frequency Trading
Álvaro Cartea, Sebastian Jaimungal And José Penalva
Excerpt
[More information](#)

Introduction to Part I

In the first part of the book we give an overview of the way basic electronic markets operate. Chapter 1 looks at the main practical issues when trading: what are the main assets traded and the main types of participants, what drives them to trade, and how do they interact. It also looks at the basic functioning of an electronic exchange: limit orders, market orders, and other types of orders, as well as the limit order book, and basic fee structures. It concludes by looking at the way the limit order book is organised and the basic experience of executing a trade.

Chapter 2 provides an overview of the theoretical economics of trading: what are the economic forces driving the competitive advantage of market makers and other traders and how do they interact. It covers the basic market making models that describe how liquidity is affected by inventory risk or the presence of better informed traders. It also looks at the market maker's trade-off between execution frequency and expected profit per trade, and how informed traders optimally exploit their informational advantage by trading gradually to limit the information leakage of their impact on order flow.

Chapters 3 and 4 look at equity market data to provide an overview of some of the basic empirical regularities that can be observed. Chapter 3 focuses on the time series properties of prices and returns, at daily and intraday frequency. It considers such issues as latency and the effects of limitations on price movements, as well as the dynamic structure of price changes, market fragmentation in US markets, and the comovement of asset prices that drives trading in pairs of assets. Chapter 4 focuses on volume and market quality. It looks at the relationship between volume and volatility, as well as known patterns in volume and prices. This is followed by an overview of different measures of liquidity and market quality: spreads, volatility, depth and trade size, and price impact. The chapter concludes by looking at other issues related to trading such as patterns in messages, order cancellations, executions and hidden orders.

1 Electronic Markets and the Limit Order Book

To understand how electronic markets work we must first understand the context in which trading in financial markets occurs. In this chapter, Section 1.1, we provide an overview of how electronic markets function, including short discussions on stocks, preferred stocks, mutual funds and hedge funds. We also discuss types of market participants (noise traders, informed traders/arbitrageurs, market makers) and in Section 1.2 how they interact. Next, in Section 1.3 we describe how electronic exchanges are structured, what limit and market orders are (as well as other order types), how exchanges collect orders in the limit order book (LOB), and the fees charged to market participants. Finally, Section 1.4 provides details of how the LOB is constructed and how market orders interact with it.

1.1 Electronic markets and how they function

Many types of financial contracts are traded in electronic markets today, so let us briefly and very superficially consider the main ones. The most familiar of these are shares or company stocks. Shares are claims of ownership on corporations. These claims are used by corporations to raise money. In the US, for these shares to be traded in an electronic exchange they have to be ‘listed’ by an exchange, and this implies fulfilling certain requirements in terms of the number of shareholders, price, etc. The listing process is usually tied to the first issuance of the public shares (initial public offering, or IPO). The fundamental value of these shares is derived from the nature of the contract it represents. In its simplest form, it is a claim of ownership on the company that gives the owner the right to receive an equal share of the corporation’s profits (hence the name, ‘share’) and to intervene in the corporate decision process via the right to vote in the corporation’s annual general (shareholders’) meetings. Such shares are called **ordinary shares** (or **common stock**) and are the most common type of shares.

The other primary instrument used by large corporations to raise capital is **bonds**. Bonds are contracts by which the corporation commits to paying the holder a regular income (interest) but gives them no decision rights. The differences between stocks and bonds are quite clear: shareholders have no guarantees on the magnitude and frequency of dividends but have voting rights, bondhold-

1.1 Electronic markets and how they function

5

ers have guarantees of regular, pre-determined payments and no voting rights. There are other instruments with characteristics from both these contracts, the most familiar of which is **preferred stock**. Preferred stock represents a hybrid of stocks and bonds: they are like bonds in that holders have no voting rights and receive a pre-arranged income, but the income they receive has fewer guarantees: its legal treatment is that of equity, rather than debt. This difference is especially relevant when the corporation is in financial distress, as debt is senior to all equity, so that in case of liquidation, debt holders' claims have priority over the corporation's assets –they get paid first. Equity holders, if they get paid, are paid only after all debtholders' claims are settled.

The universe of financial contracts is separated into different asset classes or categories according to the characteristics of the underlying assets. Shares and preferred stock belong to Equities. Bonds belong to their own asset class and are usually differentiated from cash (investments characterised by short-term investment horizons and usually with very heavy guarantees and low returns, such as money market accounts, savings deposits, Treasury bills, etc). There are also more exotic asset classes such as Foreign Exchange (FX), Commodities, Real Estate or Property. An investor will find these different types of assets in electronic exchanges, usually in the form of specialised securities such as mutual funds and exchange-traded funds (ETFs), which allow investors to invest in these asset classes in a familiar, equity-like market which simplifies the process of diversification and is associated with greater liquidity.

A **mutual fund** is an investment product that acts as a delegated investment manager. That is, when an investor buys a mutual fund, the investor gives her cash to a financial management company that will use the cash to build a portfolio of assets according to the fund's investment objective. This objective includes the fund's assets and investment strategy, and, of course, its management fees. The fund's assets can belong to a large number of possible asset classes, including all those described above: equities, bonds, cash, FX, real estate, etc. The fund's investment strategy refers to the style of investment, primarily whether the fund is actively managed or passively tracks an index.

An investor who puts money in a fund participates in both the appreciation and depreciation of the assets as allocated by the fund manager. In order to redeem her investment, i.e. to convert her investment into cash, the investor's options depend on the type of fund she purchased. There are two main types of mutual funds: open-end and closed-end funds. Closed-end funds are mutual funds that are not redeemable: the fund issues a fixed number of shares usually only once, at inception, and investors cannot sell the shares back to the fund. The fund sells the shares initially through an IPO and these shares are listed on an exchange where investors buy and sell these shares to each other.

Open-end funds are funds with a varying number of shares. Shares can be created to meet the demand of new investors, or destroyed (bought back by the fund) as investors seek to redeem theirs. This process takes place once a day, as the value of the fund's (net) assets (its Net Asset Value, NAV) is determined

after the market close. Thus, closed-end funds, that do not have to adjust their holdings in response to investor demand, have different liquidity requirements than open-end funds and thus may trade at prices different from their NAV.

A very popular type of fund that, like closed-end funds, are traded in electronic exchanges, are ETFs. Like mutual funds, ETFs act as delegated investment managers, but they differ in two key respects. First, ETFs tend to have very specific investment strategies, usually geared towards generating the same return as a particular market index (e.g., the S&P500). Second, they are not obligated to purchase investors' shares back. Rather, if an investor wants to return their share to the fund, the fund can transfer to the investor a basket of securities that mirrors that of the ETF. This is possible because the ETF sells shares in very large units (Creation Units) which are then broken up and resold as individual shares in the exchange. A Creation Unit can be as large as 50,000 shares. Overall, the general perception one gets is that investors who are looking to reduce their trading costs and find diversified investments prefer ETFs, while investors who are looking for managers with stock-picking or similar unusual skills and who aim to beat the market will prefer mutual funds.

Some investment firms feel that the regulation that is imposed on mutual fund managers to ensure they fulfill their fiduciary duties to investors are too constraining. In response to this they have created **hedge-funds**, funds that pursue more aggressive trading strategies and have fewer regulatory and transparency requirements. Because of the softer regulatory oversight, access to these investment vehicles is largely limited to accredited investors, who are expected to be better informed and able to deal with the fund's managers. Although these funds are not traded on exchanges, their managers are active participants in those markets.

There are also other securities traded in electronic exchanges; in particular, there is a great deal of electronic trading in derivative markets, especially futures, swaps and options, and these contracts are written on a wide variety of assets (bonds, FX, commodities, equities, indices). The concepts and techniques we develop in this book apply to the trading of any of these assets, although we primarily focus our examples and applications on equities. However, when designing algorithms and strategies one must always take into account the specific issues associated with the types of assets one is trading in, as well as the specifics of the particular electronic exchange(s) and the trading objectives of other investors one is likely to meet there.

1.2 Classifying Market Participants

When designing trading strategies and algorithms, it is important to understand the different types of trading behaviour one will probably encounter in these exchanges. For instance, one must consider who trades in these exchanges and why. Everyone's motivation is clear, they want to make money, but it is essential

to consider what drives them to trade – how it is that they may be looking to make money – because in many cases this will interact with our algorithm design choices and affect whether and how different algorithms achieve the desired trading objective.

Let us start from the creation of the objects of trade we have just discussed. The most familiar of these are shares. We have seen that corporations, or rather, their managers, issue stocks or equity in order to raise capital. These stocks are one of the primary objects of trade which are created when a company goes public and goes through the process of having them listed on an exchange, usually via an IPO. A corporation issues shares to raise capital for diverse economic activities, ranging from manufacturing electronic music players to mining ores in remote places. It is important to remember that these shares are claims on a corporation and as such are subject to the decisions of the company. Hence, one type of participant is corporate managers who create some of the assets that are traded in the exchanges, and who will, at times, actively participate in the market by increasing or reducing the supply of their corporation's shares, e.g., through secondary share offerings (SSOs), share buybacks, stock dividends, conversion of bonds into shares (and vice versa), etc.

We have also seen that there are other objects traded in exchanges. In equity markets we find funds (mutual funds, ETFs) created by financial management companies to commercialise their services. These funds manage large numbers of financial contracts, are very active participants in electronic exchanges, and originate a substantial fraction of the trading observed in exchanges. These 'supply-side' traders can have long-term investment goals (e.g., funds which focus on 'value investing', the kind of strategies epitomised by Warren Buffett) or focus on very immediate strategies (e.g., ETFs that replicate the returns of the S&P500). There are also **proprietary traders** who trade on a (sometimes real, sometimes illusory) trading advantage, which range from the large hedge funds we saw earlier, to small individual 'day-traders' moving in and out of asset positions from their home-offices. Proprietary traders trade on their competitive advantage: be it identifying fundamentally mispriced assets, identifying price momentum or sentiment-based price changes, having special technical abilities to process market information and identify patterns (technical traders), being able to time price movements based on news (be it the announcement of government economic figures or processing Twitter feeds), or identifying fleeting unjustified price discrepancies between equivalent assets (arbitrageurs).

Another, very important, group of market participants are 'regular investors' and '**fundamental traders**'. These are investors who have a direct use for the assets being traded. They may be individuals who buy stocks in the hope of being able to share in their growth as the corporation increases its economic value-creation and its shares appreciate in value. Or, they may want to rebalance their investments because of a change in circumstances (in response to a sudden need for cash, a change in their taste for risk or their outlook for the future). They may be corporations that use financial contracts to hedge risks such as changes in

the prices of inputs and outputs from their production activity. Traders in Brent, copper, or electricity futures worry about non-financial issues such as the number of refineries going offline for repairs, the discovery of new methods for safely transmitting electricity, or whether that tropical storm off the coast of Florida is going to turn into a hurricane and make landfall near Miami or Dade. And, one cannot ignore that governments also have a stake in market outcomes. They may want to manage their currency, issue debt to finance public expenditures, or repurchase assets to increase liquidity or maintain market stability.

The effects of the interaction amongst all these traders is one of the key issues studied in the field of market microstructure, which we will familiarise ourselves with in Chapter 2, and which helps us structure the concepts and issues behind our approach to trading. We differentiate three primary classes of traders (or trading strategies) below.

1. **Fundamental** (or **noise** or liquidity) **traders**: those who are driven by economic fundamentals outside the exchange.
2. **Informed traders**: traders who profit from leveraging information not reflected in market prices by trading assets in anticipation of their appreciation or depreciation.
3. **Market makers**: professional traders who profit from facilitating exchange in a particular asset and exploit their skills in executing trades.

Usually, one may consider arbitrageurs as a fourth type of trader, though, for our purposes we subsume arbitrageurs into informed traders moving in anticipation of price changes. Also, although it is not unusual to bundle noise and liquidity traders together, it is unusual to put them together with fundamental traders. The term ‘Noise traders’ is frequently employed to describe trading that is orthogonal to any events driving market prices, and ‘Liquidity traders’ is used for traders driven by the need to liquidate or accumulate a position for liquidity reasons orthogonal to market events.

‘Fundamental traders’, on the other hand, is a term usually reserved for traders that have medium- and long-term investment strategies based on detailed analysis of the actual business activity that underlies the asset being traded. This would naturally classify them as informed traders. However, a large fraction of their trading strategy arises from portfolio management and risk-return trade-offs that have very little short-term price information beyond that contained in the sheer size of their positions. Thus, from the point of view of a high-frequency trading algorithm, it is reasonable to consider them as ‘noise’ trades relative to the specific market events within the algorithm’s horizon. Having said this, as long as a fundamental trader is trading on information with a short-term price impact (such as knowledge of the volume of a substantial change in positions) they may also be included in the Informed trader category.

We can think of market maker types as ‘passive’ or ‘reactive’ trading. This is trading that profits from detailed knowledge of the trading process and adapts to ‘the market’ as circumstances change, while the first two types represent

more ‘active’ or ‘aggressive’ trading that only takes place to exploit specific informational advantages gained outside of the trading environment noise and fundamental traders having only a fleeting effect on short-run movements, while informed traders anticipate short-run price movements. This distinction is useful when setting up a trading strategy, although the boundary between the two is not always clear. Professional traders often leverage informational advantages gained from trading practice into the trading strategies they use for market making.

A common error is to equate market making with liquidity provision and informed trading with the taking of liquidity. Market making activity generally favours the provision of liquidity but a particular market making strategy may at times provide liquidity while at others demand it. Similarly, informed trading does not always occur via aggressive orders, and may at times be better implemented via passive orders that add liquidity. In Chapter 10 we develop algorithms for market makers who always provide liquidity to the market. These algorithms can be extended to show how market making changes when the market maker may take liquidity from the market. Moreover, in Chapter 8 we develop models of optimal execution where the agent’s strategies both take and provide liquidity.

1.3 Trading in Electronic Markets

After the who and what of electronic markets, let us look at the how. There are many ways to implement an electronic market, though essentially they all amount to having a way for people to signal their willingness to trade, and a matching engine to match those wanting to buy with those wanting to sell.

1.3.1 Orders and the Exchange

In the basic setup, an electronic market has two types of orders: **Market Orders** (MOs), and **Limit Orders** (LOs). MOs are usually considered aggressive orders that seek to execute a trade immediately. By sending an MO, a trader indicates that she wants to buy or sell a certain quantity of shares at the best available price, and this will (usually) result in an immediate trade (execution). On the other hand, LOs are considered passive orders, as a trader sending in an LO indicates her desire to buy or sell at a given price up to a certain, maximum, quantity of shares. As the price offered in the LO is usually worse than the current market price (higher than the best buy price for sell LOs, and lower than the best sell price for buy LOs), it will not result in an immediate trade, and will thus have to wait until either it is matched with a new order that wants to trade at the offered price (and executed) or it is withdrawn (cancelled).

Orders are managed by a matching engine and a limit order book (LOB). The LOB keeps track of incoming and outgoing orders. The matching engine uses a well-defined algorithm that establishes when a possible trade can occur, and if so, which criterion is going to be used to select the orders that will be executed. Most

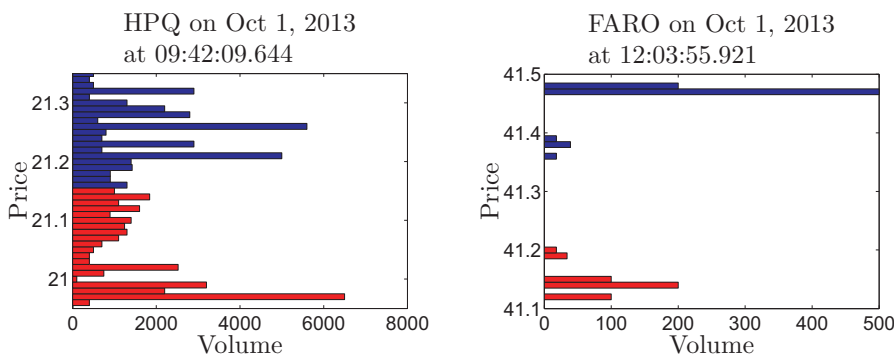


Figure 1.1 Snapshots of the NASDAQ LOB after the 10,000th event of the day. Blue bars represent the available sell LOs, red bars represent the available buy LOs.

markets prioritise MOs over LOs and then use a price-time priority whereby, if an MO to buy comes in, the buy order will be matched with the standing LOs to sell in the following way: first, the incoming order will be matched with the LOs that offer the best price (for buy orders, the sell LOs with the lowest price), then, if the quantity demanded is less than what is on offer at the best price, the matching algorithm selects the oldest LOs, the ones that were posted earliest, and executes them in order until the quantity of the MO is executed completely. If the MO demands more quantity than that offered at the best price, after executing all standing LOs at the best price, the matching algorithm will proceed by executing against the LOs at the second-best price, then the third-best and so on until the whole order is executed. LOs that have increasingly worse prices are referred to as LOs that are deeper in the LOB, and the process whereby an entering market order executes against standing LOs deeper in the LOB is called ‘walking the book’. Section 1.4 provides a more detailed view on how the LOB is built, and how MOs walk the book.

Figure 1.1 shows a snapshot of the limit order book (LOB) on NASDAQ after the 10,000th event of the day for two stocks, FARO and HPQ, on Oct 1, 2013 (see subsection 3.1.1 for a description of how this is constructed from the raw event data). The two are quite different. The one in the left panel corresponds to HPQ, a frequently traded and liquid asset. HPQ’s LOB has LOs posted at every tick out to (at least) 20 ticks away from the midprice. In the right panel, we have FARO’s LOB. FARO is a seldom traded, illiquid asset. This asset has thinly posted bids and offers and irregular gaps in the LOB. We discuss further details of this example in Section 1.4.

1.3.2 Alternate Exchange Structures

The above approach is not the only possible way to organise an exchange. For example, one could use an alternative matching algorithm, such as the **prorata** rules used in some money markets. With a prorata rule, MOs are matched