

How Modern Markets Differ from Those Past

Structural change is not new to trading in financial instruments. In fact, it is the constancy of innovation that has helped drive the leadership of modern financial institutions. High-frequency trading (HFT) has stepped into the limelight over the past few years and delivered considerable operational improvements to the markets, most of which have resulted in lower volatility, higher market stability, better market transparency, and lower execution costs for traders and investors. This chapter of the book provides the overview of dramatic changes that precipitated in the securities markets over the past 50 years, and defines HFT and core strategies falling under the HFT umbrella.

Over the past two decades, the demand for computer technology in consumer markets has led to significant drops in hardware prices across the board, as discussed in detail in Chapter 2 of this book. As a result, technology-enabled trading has become cost effective, and the ensuing investment into software has made trading platforms more accessible and more powerful. Additionally, the savings from lower errors in message transmission and data input, higher reliability of order execution, and continuity of business through computer code, deliver a business case for deepening financial firms' reliance on their technology systems. The escalating complexity of regulations also requires more advanced reporting capabilities that are becoming prohibitively expensive without substantial platforms. The lower cost base squeezes margins further, and this puts pressure on the traditional full-service model. Figures 1.1 and 1.2 illustrate the financial services landscape circa 1970s and today.

In the 1970s and earlier, the market participants were organizations and individuals now considered "traditional" players. As Figure 1.1 shows, on the portfolio management or "buy" side, the markets engaged

- Discretionary asset managers, including pension funds, mutual funds, and hedge funds.
- Retail flow, including individual “mom-and-pop” investors, and others with comparatively smaller capitalization.
- Manual speculators, individuals involved in intraday proprietary trading for their own account or for the account of their bank.

On the transaction facilitation, middle-men, or “sell” side, the markets supported

- Manual market makers (representatives of broker-dealers), taking short-term inventory risk, providing quotations to the buy side, and generally facilitating the buy-side trading for a fee.




FIGURE 1.1 Financial Markets in the 1970s, before Electronization




FIGURE 1.2 Financial Markets Today

- A single not-for-profit exchange in each asset class was established to curtail wild speculation by the exchange members and to lower transaction costs paid by the investors.¹

The highly manual and therefore labor-intensive financial landscape of the 1970s was characterized by high transaction costs, leading to low turnover of securities; a high degree of error associated with manual processing of orders, and relatively high risk of trading, as traders predominantly relied on their experience and intuition as opposed to science in making their bets on the markets. Yet, the 1970s were also high margin businesses, with brokers receiving a large share of the spoils in the form of large commissions and, ultimately, the proverbial “fat-cat” bonuses to the tune of tens of millions of dollars.

Fast-forward to today’s markets, illustrated in Figure 1.2: new entrants successfully compete using lean technology and science to hash out precise investing models, reshaping the markets in the process:

- Quantitative money managers, such as mutual funds and hedge funds, are using the precise science of economics, finance, and the latest mathematical tools to chisel increasingly close forecasts of securities prices, improving profitability of their investments.
- Automated market makers, for example, broker-dealers and hedge funds, harness the latest technology, studies of market microstructure, and HFT to deliver low transaction costs, taking over market share from traditional broker-dealers.
- Automated arbitrageurs, such as statistical arbitrage hedge funds and proprietary traders, use quantitative algorithms, including high-frequency trading techniques, to deliver short-term trading profits.
- Multiple alternative trading venues, like new exchanges and dark pools, have sprung up to address market demand for affordable quality financial matching services.

These innovations have changed the key characteristics of the markets, and largely for the better:

- The markets now enjoy vastly democratic access: due to proliferation of low-cost technology, anyone can trade in the markets and set quotes, a right formerly reserved to members of the exclusive connections-driven club of broker-dealers.
- Plummeting transaction costs keep money in investors’ pockets; more on this later.
- Automated trading, order routing, and settlement deliver a new low degree of error.

¹ Most exchanges became not-for-profit only in the 1970s. From the time of their formation to the 1970s, however, the exchanges were very much for profit. In fact, the Buttonwood agreement of 1792 that laid foundation to the New York Stock Exchange, specified explicit profitability rules: no broker was allowed to charge less than 0.25 percent of the transaction volume, a staggering commission by today’s standards.

The extreme competition among the new entrants and old incumbent market participants, however, has also resulted in reduced margins for broker-dealers, squeezing out technology-inefficient players.

The way trading is done has changed over time and these newer approaches affected the relative power of consumers and institutions. In the 1970s' marketplace, the trading process would often proceed as follows:

1. Brokers would deliver one-off trading ideas to their buy-side clients. The ideas were often disseminated via countless phone calls, were based on brokers' then-unique ability to observe markets in real time, and were generally required compensation in "soft-dollar" arrangements—if the customer decided to trade on the idea, he was expected to do so through the broker who produced the idea, and the customer would pay for the idea in the form of potentially higher broker commissions.
2. If and when the customer decided to trade on the idea, the customer would phone in the order to the broker or the broker's assistant. Such verbal orders frequently resulted in errors: the noise on the brokers' trading floors often impeded correct understanding of customer instructions.
3. After receiving a customer's order, the broker's next steps would depend on the size of the placed order: while large orders would be taken to the market right away (potentially in smaller parcels), smaller orders would sit on the broker's desk, waiting for other similar orders to fill up a "round lot"—the minimum order size executable on an exchange. Smaller customers were thus often at a disadvantage, waiting for execution of their orders while the favorable market price slipped away.
4. Once the order or several sequential orders comprised the order size acceptable to the broker, the broker would route the order to the appropriate exchange.
5. Next, human representatives of the exchange, known as "specialists," would match the order and send the trade acknowledgments back to the broker. It is well understood that the specialists often created preferential terms for some of their connections, at the expense of orders of others. Such behavior rewarded investment in connections and chummy networks, and resulted in exclusive Wall Street cliques capable of significant price discrimination for in-group versus out-of-group customers. Even though exchanges operated as not-for-profit organizations, influence peddling was common, and the markets were a long way away from anything resembling an equal playing field for all participants.
6. The broker notified the client of execution and collected his commissions and oversized bonuses. The brokers presided over the power of the markets and were compensated as kings.

Figure 1.3 illustrates the traditional investing process prevalent circa 1970s.

Fast-forward 40-something years ahead, and the balance of power has shifted. Customers have increased their expertise in quantitative analysis and are often better equipped for research than brokers. Brokers' area of expertise has decreased in scope from the all-encompassing sell-side research into securities behavior to a more narrow, albeit still important area of algorithmic execution designed to help clients




FIGURE 1.3 Broker-centric Investing Process Prevalent before Electronization

navigate the choppy intraday trading waters. With such buy-side investors, the market flow evolves according to the process in Figure 1.4:

1. Customers, not brokers, generate research based on forecasts of securities movements and their existing allocations, all within the quantitative portfolio management framework.
2. The customer places an order via electronic networks, greatly reducing errors and misunderstandings. The order instantaneously arrives on the broker's desktop.
3. The customer or the broker selects the broker's optimal execution algorithm designed to minimize the customer's execution costs and risk, speed up execution whenever possible, and minimize observability of the customer's trading actions.
4. Selected algorithm electronically parcels out the customer's order and routes the order slices to relevant exchanges and other trading venues.
5. Trading venues match the customer's order slices and acknowledge execution.
6. The broker sends the order acknowledgment back to the customer, and receives his considerably lower commission. (In 1997, the lowest broker commission on retail trades was offered by Merrill Lynch, and the commission was \$70 per trade. Today, Interactive Brokers charges about \$0.70 per trade, a 100-fold reduction in transaction costs available to clients.)

Some customers go even further and prefer to do away with broker service altogether, building their own execution algorithms, keeping a higher share of the profits. Plummeting costs of technology have enabled fast distribution of tick data to all interested parties, and now customers, not just brokers, can watch and time markets and generate short-term forecasts of market behavior. Customers taking the largely




FIGURE 1.4 Modern Investing Process, Scenario 1: Brokers Provide Best Execution for Clients' Orders

broker-independent route are said to engage in “direct access” to the markets, and their execution process consists of the following steps:

1. The broker grants the customer a privilege to access the exchange directly for a negotiated per-trade or per-volume fee. To grant access, the broker may allow the customer to use the broker's own identification with a specific exchange. The customer's order routing systems then use the broker's identification in order messaging with the exchange.
2. Customer computer systems or human analysts generate a high- or low-frequency portfolio allocation decision that involves one or more trading orders.
3. Customer uses his own order splitting and routing algorithms to optimally place his orders directly with exchanges and other trading venues.
4. One or several exchanges and trading venues match the orders, acknowledge execution directly to client.
5. The broker receives settlement information and charges the client for the privilege of using the broker's direct access identification.

Figure 1.5 summarizes these steps.

■ Media, Modern Markets, and HFT

While the market-wide changes have disturbed the status quo on the broker-dealer side and squeezed many a broker out of business, the changes to the society at large have been mostly positive, depositing the saved dollars directly into investor pockets. Gone are the multimillion-dollar bonuses of many brokers taking phone orders




FIGURE 1.5 Modern Investing Process, Scenario 2: Clients Decide on Best Execution, Access Markets Directly

and watching markets on their computer screens. The money has been redirected to bank shareholders and end investors.

Clearly, not everyone is happy about such shifts in the industry, and the least happy bunch happens to be brokers losing their income to automation. Stripped of the ability to extract easy money out of investors' pockets, brokers have been the most vocal opponents of high-frequency trading. Brokers like Arnuk and Saluzzi (2012), for example, denounce automation, yet wax eloquent about those manual error-prone days when investors were not allowed on exchanges and brokers were the fat cats of the world.

Some brokers, whose lifestyle has been significantly reduced by technology, attempt to demonize HFT with an even more sinister goal in mind: they are seeking to lure in investors to their outfits still charging exorbitant transaction costs under the guise of protecting the poor investor lambs from the HFT predators. Investors should take time to compare costs of trading through a broker versus other available options. Chapters 5, 12, and 15 of this book provide specific information to enable low-frequency investors to estimate the risk of potentially adverse HFT, and take educated steps to managing said risk, without relying on self-serving hype of selected brokers who refuse to catch up on technical innovation, resorting to scare tactics at the expense of their clients instead. The remainder of this chapter is devoted to explaining the evolutionary nature of HFT and the definitions and overview of strategies that fall under the HFT umbrella.

■ HFT as Evolution of Trading Methodology

Brokers who speak loudly against the HFT tend to rely on technical analysis in making their decisions of when to enter or exit a position. Technical analysis was one of

the earliest techniques that became popular with many traders and is, in many ways, a direct precursor to today's sophisticated econometrics and other HFT techniques.

Technical analysts came in vogue in the early 1910s and sought to identify recurring patterns in security prices. Many techniques used in technical analysis measure current price levels relative to the rolling moving average of the price, or a combination of the moving average and standard deviation of the price. For example, a technical analysis technique known as moving average convergence divergence (MACD) uses three exponential moving averages to generate trading signals. Advanced technical analysts would look at security prices in conjunction with current market events or general market conditions to obtain a fuller idea of where the prices may be moving next.

Technical analysis prospered through the first half of the twentieth century, when trading technology was in its telegraph and pneumatic-tube stages and the trading complexity of major securities was considerably lower than it is today. The inability to transmit information quickly limited the number of shares that changed hands, curtailed the pace at which information was incorporated into prices, and allowed charts to display latent supply and demand of securities. The previous day's trades appeared in the next morning's newspaper and were often sufficient for technical analysts to successfully infer future movement of the prices based on published information. In post-WWII decades, when trading technology began to develop considerably, technical analysis developed into a self-fulfilling prophecy.

If, for example, enough people believed that a "head-and-shoulders" pattern would be followed by a steep sell-off in a particular instrument, all the believers would place sell orders following a head-and-shoulders pattern, thus indeed realizing the prediction. Subsequently, institutional investors have moved to high-frequency econometric modeling using powerful computer technology, trading away technical patterns. By now, technical analysis at low frequencies, such as daily or weekly intervals, is marginalized to work only for the smallest, least liquid securities, which are traded at very low frequencies—once or twice per day or even per week.

Some technical analysis techniques, such as momentum or Bollinger bands, have been successfully adopted and extended by modern-day quants in all investing frequencies. It has long been shown that human investors tend to pour money into strategies that worked in recent months. As a result, strategies working in the past month are also likely to work the following month, forming a tradable momentum that can be detected using simple technical moving-average-based indicators, as well as more complex quantitative tools. Similarly, Bollinger bands detect deviation of prices the prespecified number of standard deviations away from the mean. The concept of statistical arbitrage extended Bollinger band principle to detect, for example, deviation of price differences from their long-running means. In this trading exercise, commonly known as pairs trading, traders identify the overpriced and underpriced financial instruments when the price of one instrument exceeds the price of another by the prespecified number of standard deviations of price difference changes. More generally, quants use Bollinger band ideas to pinpoint mean-reverting processes and trade financial instruments with the expectation that the measured average quantity will stay stable, or "stationary" in the language of statistics.

Another important investing and trading technique, known as fundamental analysis, originated in equities in the 1930s when traders noticed that future cash flows,

such as dividends, affected market price levels. The cash flows were then discounted back to the present to obtain the fair present market value of the security. Graham and Dodd (1934) were the earliest purveyors of the methodology and their approach is still popular. Over the years, the term *fundamental analysis* expanded to include pricing of securities with no obvious cash flows based on expected economic variables. For example, fundamental determination of exchange rates today implies equilibrium valuation of the rates based on macroeconomic theories.

Fundamental analysis developed through much of the twentieth century. Today, fundamental analysis refers to trading on the expectation that the prices will move to the level predicted by supply-and-demand relationships, the fundamentals of economic theory. In equities, microeconomic models apply; equity prices are still most often determined as present values of future cash flows. In foreign exchange, macroeconomic models are most prevalent; the models specify expected price levels using information about inflation, trade balances of different countries, and other macroeconomic variables. Derivatives are traded fundamentally through advanced econometric models that incorporate statistical properties of price movements of underlying instruments. Fundamental commodities trading analyzes and matches available supply and demand.

Various facets of fundamental analysis are inputs into many high-frequency trading models, alongside market microstructure. For example, event arbitrage consists of trading the momentum response accompanying the price adjustment of the security in response to new fundamental information. The date and time of the occurrence of the news event is typically known in advance, and the content of the news is usually revealed at the time of the news announcement. In high-frequency event arbitrage, fundamental analysis can be used to forecast the fundamental value of the economic variable to be announced, in order to further refine the high-frequency process.

Like selected technical models, some fundamental models were adopted by quants who extended the precision of their models, and often dramatically sped up calculation of the relevant values. Fair values of equities following an earnings announcement were recomputed on the fly, enabling quants to reap the profits, at the expense of fundamental traders practicing longhand analysis in Excel spreadsheets.

Speed, in fact, became the most obvious aspect of quant competition. Whoever was able to run a quant model the fastest was the first to identify and trade on a market inefficiency and was the one to capture the biggest gain. To increase trading speed, traders began to rely on fast computers to make and execute trading decisions. Technological progress enabled exchanges to adapt to the new technology-driven culture and offer docking convenient for trading. Computerized trading became known as *systematic trading* after the computer systems that processed run-time data and made and executed buy-and-sell decisions.

High-frequency trading developed in the 1990s in response to advances in computer technology and the adoption of the new technology by the exchanges. From the original rudimentary order processing to the current state-of-the-art all-inclusive trading systems, HFT has evolved into a billion-dollar industry.

To ensure optimal execution of systematic trading, algorithms were designed to mimic established execution strategies of traditional traders. To this day, the term

algorithmic trading usually refers to the automated “best execution” process—that is, the optimization of buy-and-sell decisions once these buy-and-sell decisions were made by another part of the systematic trading process or by a human portfolio manager. Algorithmic trading may determine how to process an order given current market conditions: whether to execute the order aggressively (on a price close to the market price) or passively (on a limit price far removed from the current market price), in one trade or split into several smaller “packets.” As mentioned previously, algorithmic trading does not usually make portfolio allocation decisions; the decisions about when to buy or sell which securities are assumed to be exogenous.

The advances in computer technology over the past decades have enabled fully automated HFT, fueling the profitability of trading desks and generating interest in pushing the technology even further. Trading desks seized upon cost savings realized from replacing expensive trader head count with less expensive trading algorithms along with other advanced computer technology. Immediacy and accuracy of execution and lack of hesitation offered by machines as compared with human traders has also played a significant role in banks’ decisions to switch away from traditional trading to systematic operations. Lack of overnight positions has translated into immediate savings due to reduction in overnight position carry costs, a particular issue in crisis-driven tight lending conditions or high-interest environments.

Banks also developed and adopted high-frequency functionality in response to demand from buy-side investors. Institutional investors, in turn, have been encouraged to practice high-frequency trading by the influx of capital following shorter lock-ups and daily disclosure to investors. Both institutional and retail investors found that investment products based on quantitative intraday trading have little correlation with traditional buy-and-hold strategies, adding pure return, or alpha, to their portfolios.

Under the Dodd-Frank Act, banks were forced to close many of the proprietary trading operations, but not HFT. In certain banks, the formerly prop-trading HFT is alive and well in the market-making function, where it is now run with client rather than bank capital and is often referred to as *prehedging*.

As computer technology develops further and drops in price, high-frequency systems are bound to take on an even more active role. Special care should be taken, however, to distinguish HFT from electronic trading, algorithmic trading, and systematic trading. Figure 1.6 illustrates a schematic difference between high-frequency, systematic, and traditional long-term investing styles.

Systematic trading refers to computer-driven trading decisions that may be held a month or a day or a minute and therefore may or may not be high frequency. An example of systematic trading is a computer program that runs daily, weekly, or even monthly; accepts daily closing prices; outputs portfolio allocation matrices; and places buy-and-sell orders. Such a system is not a high-frequency system.

Another term often mentioned in conjunction but not synonymous with HFT is *electronic trading*. Electronic trading refers to the ability to transmit the orders electronically as opposed to telephone, mail, or in person. Since most orders in today’s financial markets are transmitted via computer networks, the term *electronic trading* is rapidly becoming obsolete.




FIGURE 1.6 HFT versus Algorithmic (Systematic) Trading and Traditional Long-Term Investing

Algorithmic trading is more complex than electronic trading and can refer to a variety of algorithms spanning order-execution processes as well as high-frequency portfolio allocation decisions. The execution algorithms are designed to optimize trading execution once the buy-and-sell decisions have been made elsewhere. Algorithmic execution makes decisions about the best way to route the order to the exchange, the best point in time to execute a submitted order if the order is not required to be executed immediately, and the best sequence of sizes in which the order should be optimally processed. Algorithms generating HFT signals make portfolio allocation decisions and decisions to enter or close a position in a particular security. For example, algorithmic execution may determine that a received order to buy 1 million shares of IBM is best handled using increments of 100-share lots to prevent a sudden run-up in the price. The decision fed to the execution algorithm, however, may or may not be high frequency. An algorithm deployed to generate HFT signals, however, would generate the decision to buy the 1 million shares of IBM. The high-frequency signals would then be passed on to the execution algorithm that would determine the optimal timing and routing of the order.

Successful implementation of HFT requires both types of algorithms: those generating HFT signals and those optimizing execution of trading decisions. This book covers both groups of algorithms: those designed for generation of trading signals (Chapters 8 through 11) and those for order execution designed to conceal information within (Chapter 15). Chapter 14 of the book also includes latest algorithms for managing risk of HFT operations.

The intent of algorithmic execution is illustrated by the results of a survey conducted by Automated Trader in 2012. Figure 1.7 shows the full spectrum of responses from the survey. In addition to the previously mentioned factors related to adoption of algorithmic trading, such as performance management and reporting, both buy-side and sell-side managers also reported their use of the algorithms to be driven by trading decision and portfolio management needs.




FIGURE 1.7 Reasons for Using Algorithms in Trading

Source: Automated Trader Survey, 2012

True HFT systems make a full range of decisions, from identification of under-priced or overpriced securities through optimal portfolio allocation to best execution. The distinguishing characteristic of HFT is the short position holding times, one day or shorter in duration, usually with no positions held overnight. Because of their rapid execution nature, most HFT systems are fully systematic and are also examples of systematic and algorithmic trading. All systematic and algorithmic trading platforms, however, are not high frequency.

Ability to execute an order algorithmically is a prerequisite for HFT in a given financial instrument. As discussed in Chapter 3, some markets are not yet suitable for HFT, inasmuch as most trading in those markets is performed over the counter (OTC). According to research conducted by Aite Group, equities are the most algorithmically executed asset class, with over 50 percent of the total volume of equities expected to be handled by algorithms by 2010. As Figure 1.8 shows, equities are closely followed by futures. Advances in algorithmic execution of foreign exchange, options, and fixed income, however, have




FIGURE 1.8 Adoption of Algorithmic Execution by Asset Class

Source: Aite Group

been less visible. As illustrated in Figure 1.8, the lag of fixed-income instruments can be explained by the relative tardiness of electronic trading development for them, given that many of them are traded OTC and are difficult to synchronize as a result.

While research dedicated to the performance of HFT is scarce relative to data on long-term buy-and-hold strategies, anecdotal evidence suggests that most computer-driven strategies are high-frequency strategies. Systematic and algorithmic trading naturally lends itself to trading applications demanding high speed and precision of execution, as well as high-frequency analysis of volumes of tick data. Systematic trading, in turn, has been shown to outperform human-led trading along several key metrics. Aldridge (2009b), for example, shows that systematic funds consistently outperform traditional trading operations when performance is measured by Jensen's alpha (Jensen, 1968), a metric of returns designed to measure the unique skill of trading by abstracting performance from broad market influences. Aldridge (2009b) also shows that the systematic funds outperform nonsystematic funds in raw returns in times of crisis. That finding can be attributed to the lack of emotion inherent in systematic trading strategies as compared with emotion-driven human traders.

Furthermore, computers are superior to humans in such basic tasks as information gathering and lightning-fast analysis of a multitude of quotes and news. Physiologically, the human eye cannot capture more than 50 data points per second, as evidenced by an entirely different industry—cinematography. In modern movies, the human eye is exposed to only 24 frames per second, which appear seamless to most moviegoers. Even then, the majority of images displayed on sequential frames involve continuously moving objects. In comparison, modern financial information incorporates drastically bouncing quotes, the number of which can easily exceed 1,000 per second for just one financial instrument. Detecting inter-instrument information spillovers involves processing data for multiple assets and asset classes, as discussed in Chapter 15. Where efficient processing of high volumes of information is key to profitable trading, technology-averse humans have little chance of succeeding. HFT takes over.

■ What Is High-Frequency Trading?

High-frequency trading is an umbrella term comprising several groups of strategies. Given the breadth of HFT, various market participants have somewhat divergent opinions of what HFT actually stands for. This section discusses common definitions of HFT:

- *A definition of HFT that includes all activity utilizing fast algorithmic execution.* For example, the Technology Subcommittee of the U.S. Commodity Futures Trading Commission (CFTC), tasked with compiling a working definition of HFT, came back with the following draft definition in June 2012:

High-frequency trading is a form of automated trading that employs:

- Algorithms for decision making, order initiation, generation, routing, or execution, for each individual transaction without human direction;

- low-latency technology that is designed to minimize response times, including proximity and co-location services;
- high speed connections to markets for order entry; and
- high message rates (orders, quotes or cancellations).

Such definition captures many high-frequency traders, yet also includes 95 percent of investors using algorithmic technology to execute their orders. Even a “mom-and-pop” retail investor entrusting his broker to execute his order in the most efficient algorithmic manner becomes a high-frequency trader under the definition proposed by the CFTC’s subcommittee on HFT. Not surprisingly, this definition generated strong dissent from many members of the subcommittee itself.

- *The definition of HFT as a latency-sensitive subset of algorithmic trading.* Gomber, Arndt, Lutat, and Uhle (2011) proposed to define HFT as shown in Figure 1.9. Under such definition, HFTs are also traders “on steroids,” utilizing super-fast technology to speed up algorithmic processes and drive models in supersonic time. Interestingly, also under this definition, the HFTs do not engage in portfolio construction or management, but generate trading signals, validating models, and execute trades in any one security.




FIGURE 1.9 HFT vs. Algorithmic Trading and Quant Portfolio Management

Source: Gomber, Arndt, Lutat and Uhle (2011)

- *The definition of HFT based on holding period of capital throughput.* A survey of hedge-fund managers, conducted by FINalternatives in 2009, generated the following definition of HFT:

High-frequency trading comprises

- Systematic,
- Quant-based models

- With holding periods from a fraction of a second to one day (no positions held overnight).

The survey was based on nearly 300 responses from hedge fund managers who subscribe to FINalternatives (close to 10,000 questionnaires were sent out). It is also worth noting that at the time, a prominent multibillion-dollar Greenwich-based hedge fund launched a high-frequency fund with an average position holding period of three days, a far departure from submicrosecond frequencies often mentioned in connection with HFT. The fund was later retracted.

- *The definition of HFT based on their observed market activity.* Kirilenko, Kyle, Samadi, and Tuzun (2011) identify high-frequency traders as market participants that generate high market volume, all the while holding low inventory. The researchers use the definition to distinguish HFT from other market participants:
 - Intermediaries, characterized by low inventory, but not high trading volume.
 - Fundamental buyers, who are consistent net buyers intraday.
 - Fundamental sellers, who are consistent net sellers within a given day.
 - Small traders, generating low volume.
 - Opportunistic traders, loosely defined as all other traders, not fitting the definition of HFT or other categories above.

Such definition may rely on somewhat arbitrary cutoffs of low inventory and high volume.

The definition of HFT based on behavior unattainable by human market participants. A common definition used by brokers to segment their clients into HFT and non-HFT, this definition calls for attribution of trading activity of each specific account into human feasible and human infeasible. For example, an account generating 200 orders per second would be deemed HFT, as would an account that consistently succeeds at locking in a penny gain day-in and day-out.

This book considers all of the definitions discussed here.

While a concrete definition of HFT has proven to be a challenge, most market participants are comfortable with the range of strategies deployed by HFT, summarized in Figure 1.10.

■ What Do High-Frequency Traders Do?

Despite the disagreements about the precise definition of HFT, most market participants agree that HFT strategies fall into the following four broad classes:

1. Arbitrage
2. Directional event-based trading
3. Automated market making
4. Liquidity detection




FIGURE 1.10 Major Categories of HFT Strategies

Arbitrage strategies trade away price deviations from long-running equilibria or relative asset mispricing, and can include multiple asset classes, as well as multiple exchanges. Many HF arbitrage strategies detect price discrepancies in multiple securities, as discussed in Chapter 8. Several strategies arbitrage prices of the same asset trading on different exchanges, are known as *latency arbitrage strategies*, and are discussed in Chapter 12. Most arbitrage strategies are based on assumptions of mean-reversion of asset prices.

Statistical arbitrage models comprise a range of models, including cross-asset models, where financial securities have strong statistical relationships. All the models included in the book are deep-rooted in economic theories, ruling out spurious statistical relationships often developed using plain data mining and also known as the Spaghetti Principle of Modeling (if one throws a plate of spaghetti filled with data against the wall of statistics, something may stick; what sticks, however, may not have any sound reason for sticking and is likely to fall apart in production). For example, bonds and interest rate futures have been shown to possess considerable dependencies, and their values, as a result, tend to move in tandem. When prices of bonds or interest rate futures deviate from their long-running price equilibrium for no obvious reason, statistical arbitrage trading may be feasible via buying the instrument with a lower-than-expected price relative to the other instrument(s), and selling the instrument with a higher-than-expected price relative to the other instrument(s). Chapter 8 details many economic models, as well as the model estimation techniques and known results.

Directional strategies identify short-term trend or momentum. This class of high-frequency strategies includes event-driven strategies, discussed in Chapter 9, other strategies based on predictable short-term price movements, discussed in Chapter 11, as well as the controversial ignition strategies, discussed in Chapter 12. Event arbitrage models show the methodology as well as performance of trading on predictable and recurrent effects of news. Various types of news used in event arbitrage are showcased in Chapter 9, which also includes references to the latest relevant studies as well as specific practical examples.

Automated market-making strategies comprise perhaps the most traditional trading strategies, encompassing automated market making, a cost-effective and accurate alternative to human broker-dealers, discussed in detail in Chapter 10. The category of automated market making and liquidity provision includes both inventory-driven and information-driven approaches. Inventory-driven methods tend to focus on joint minimization of the inventory risk and market risk, ensuring that the positions are within a trader's risk tolerance limits given market conditions, and hedged where appropriate. Information-driven market-making models are built with the goal of minimizing the risk of adverse selection, the risk of taking an opposite position to a better-informed party. To minimize the number of such losing positions, high-frequency traders can deploy a wide range of models that help forecast short-term directionality of markets, track the number of well-informed market players in the market waters, and even help forecast impending lumps and shortages of liquidity, covered in Chapter 11. These techniques allow traders to choose the quantities and levels of aggressiveness of their orders based on expectations of surplus or dearth of liquidity.

Perhaps least palatable to the low-frequency investors are liquidity detection strategies, like ping-pong (also known as sniffing and sniping), quote stuffing, and spoofing, addressed in Chapter 12. While this book focuses on explaining sound HFT strategies, the book attempts to draw a balanced perspective and include the methodology behind controversial HFT as well. "Ping-pong" has been shown to exist on selected venues (ping-pong was detected in dark pools). The nature of other strategies like "ignition strategies" have been mostly speculative, and no credible evidence of strategy existence has been produced to date. Still, the hypothetical strategies like ignition strategies have been included for completeness, accompanied by a brief analysis of their feasibility, properties, and impact on the broader markets.

■ How Many High-Frequency Traders Are There?

The number of high-frequency traders largely depends on the definition of HFT used. As mentioned earlier, under the CFTC draft definition proposed in June 2012, 19 out of every 20, or 95 percent of all investors and traders would qualify as HFT. Kirilenko, Kyle, Samadi, and Tuzun (2011), define HFTs as traders who produce large trading volume while holding little inventory, and find that HFTs account for about 30 percent of volume in the Standard & Poor's (S&P) 500 E-Mini markets. Aldridge (2012a) estimates that HFTs comprise just 25 to 30 percent in EUR/USD foreign exchange futures and that in the most liquid exchange-traded fund, the S&P 500 SPDR (NYSE:SPY), high-frequency traders on average represent fewer than 20 percent of market participants.

■ Major Players in the HFT Space

Many HFT participants prefer to stay out of the limelight, all the while generating considerable profits. The most well-known HFT outfits include Getco, Renaissance Capital, and DE Shaw. Lesser-known but still very profitable players dedicated to HFT include specialist firms like IV Capital, DKR Fusion, and WorldQuant.

The line between HFT and other forms of trading, however, can be blurred. As mentioned earlier, HFT, and specifically automated market making, are becoming staples on most trading desks in all the major banks. And the advantages of such developments are easy to see: the new automated market-making “robots” are considerably more accurate, inexpensive, and reliable than their human counterparts. Likewise, HFT can seamlessly blend in with activities of statistical arbitrage. In Canada, for example, banks often list most of their HFT in the statistical arbitrage category in the banks’ annual reports.

■ Organization of This Book

This book is written with the explicit goal of providing the latest, yet applied and ready-to-implement information to management and employees that are interested in starting or enhancing their high-frequency trading operations, individuals and institutions seeking to protect their and their clients’ trading activity against high-frequency traders, as well as casual observers, seeking to better understand modern financial markets.

Chapters 2 through 5 of the book explain the present-day frontiers in financial markets. Chapter 2 describes technological evolution that has enabled algorithmic and high-frequency trading. Chapters 3 through 5 lay the foundation of analysis via description of modern microstructure, high-frequency data, and trading costs.

Chapters 6 and 7 delve into the economics of high-frequency trading. Chapter 6 describes methodologies for evaluating performance and capacity of HFT strategies, and Chapter 7 outlines the business case of HFT.

Chapters 8 through 12 and 14 through 16 are devoted to actual implementation of HFT. Chapters 8 through 12 dissect core models of today’s high-frequency strategies. Chapter 14 focuses on risk measurement and management of high-frequency trading as well as portfolio construction. Chapters 15 and 16 discuss the nuts-and-bolts in implementation of HFT systems, as well as best practices in running and monitoring HFT systems.

Chapters 13 and 15 focus on regulation of HFT and mitigation of HFT externalities. Chapter 13 presents a summary of current regulatory thought on HFT, discusses models for detection of HFT market manipulation, as well as mathematics of foreseeing market-wide events like flash crashes. Chapter 15 of the book offers solutions for low-frequency traders concerned about the impact of HFT on modern markets. Chapter 15 discusses the latest order slicing techniques and their respective ability to avoid information-prying HFTs, and may also prove useful to high-frequency traders seeking to further expand capacity of their trading systems.

■ Summary

- High-frequency trading is an organic evolution of trading technology.
- The technological evolution of financial markets created ability to replace human-driven intermediation function with cost-effective technology, returning broker compensation to end investors and bank shareholders.

- High-frequency trading strategies are well defined, and most of them are beneficial to the markets.

■ End-of-Chapter Questions

1. Describe the major groups of today's market participants. What role do they play? How do they interact?
2. What are the core groups of strategies deployed by high-frequency traders?
3. How do high-frequency trading strategies relate to other trading strategies, such as technical analysis, fundamental analysis, and quant strategies?
4. What are the major changes that have occurred in the financial markets over the past 40 years?
5. What is algorithmic trading?
6. How do end investors benefit from high-frequency trading?

