

CHAPTER 1**Introduction to
Systematic Trading**

Systematic trading is a particular discipline of trading, which is one of the oldest human activities. Trading and the associated arena set by the marketplace coevolved in time to become one of the dominant industries on the planet. At each stage of their development, new efficiencies were introduced.

Starting as barter where goods were exchanged “on sight,” the first major evolutionary step was the introduction of a numeraire (be it gold or fiat money) that literally allowed comparison between apples and oranges. It also allowed the storage of value in a compact way. Then the first organized exchanges in Flanders and Holland introduced several key concepts: first and foremost the concept of the exchange as a risk disintermediator, then the concept of standardization so important in comparing bulk commodities, and finally the technique of open outcry—the famous Dutch Auction at the basis of the modern exchange mechanism. Despite the fact that the concept of interest (via grain loans) was introduced by the Egyptians, the effective leverage in the marketplace only came with the growth of the stock markets and commodity futures markets in the United States in the early twentieth century. Also at that point the nascent global banking system spurred the creation of the money market where short-term loans are traded in a standardized fashion and help to transfer leverage between counterparties. An important factor in the stabilization of the market process was the introduction of floor specialists or market-makers who ensured orderly matching of buyers and sellers. With the advent of increasing computing power, the co-evolution of the marketplace and the trading associated with it has accelerated further. Not only has the banking system evolved into a global network of compensating agents where money can be transferred at the speed of light, but the whole flow of information has become available to a much

larger group. The marketplace and trading have become truly global and gradually more electronic. This evolution has taken its toll on the open outcry system and on specialists, with some of them being gradually crowded out by robotic market-making computer programs and the increasing importance of semi-private matching engines like dark pools and electronic commerce networks (ECNs).

And this is where we are right now, a world some would say of information overflow, of competition for microseconds, of over-leverage and over-speculation. Each evolutionary stage comes with its share of positives and negatives. A new organism has to keep searching for its boundaries independently of its forebears and try to learn from its rewards and mistakes so as to set the stage for its own progress.

This book focuses on a subset of trading techniques that applies to a subset of the marketplace. It explores the systematic automated trading of liquid instruments such as foreign exchange, futures, and equities. It is an activity on the edge of the evolutionary path that also tries to find its current boundaries, technologically and conceptually.

This introductory chapter sets the philosophical context of trading and puts on equal footing the seemingly contradictory approaches of systematic and discretionary trading. They are compared as business activities by presenting a cost-benefit analysis of each, concluding with the viability and similarity of both. The psychological implications of choosing one path over the other is analyzed and it is argued that it is the defining criterion from a rational trader's perspective. The chapter concludes by putting the theoretical Parts One to Three and the practical Part Four of the book into the historic context and showing how the evolution of systematic trading is intimately related to the progress in technology and science.

1.1 DEFINITION OF SYSTEMATIC TRADING

The majority of successful traders design their trading strategy and trading discipline in the most objective way possible but cannot be qualified as systematic, because many of their decisions are based on their perceived state of the world, the state of their mind, and other factors that cannot be computationally quantified. The type of trading that is relying on noncomputable processes will be qualified as discretionary in this book.

As opposed to the discretionary, the qualifier *systematic* encompasses the following two concepts:

1. The existence of a rules-driven trading strategy that is based on objectively reproducible (computable) inputs.
2. The application of that strategy with discipline and outside of the human emotional context.

Systematic trading implies the construction of a mathematical model of a certain behavior of the market. This model is then encompassed in a decision-making algorithm that outputs continuously the allocation of exposure to such a model in the context of the trader's other models' behavior, total risk allocation, and other objective and reproducible inputs. The continuous running of such an algorithm is oftentimes best left to a robot.

Before making further comparisons let us now explore the two trading approaches in a broader philosophical context of the perceived behavior of the market and its participants.

1.2 PHILOSOPHY OF TRADING

The philosophy of trading derives from a set of beliefs about the workings of the human mind, the behavior of crowds of reward-seeking individuals, and the resulting greed-fear dynamics in the market. Trading is a process, a strategy, a state of mind. It is the mechanism by which a market participant survives and thrives in the marketplace that itself is composed of such participants and constrained by political and regulatory fads and fashions.

Choosing a trading style is as much about knowing and understanding the workings of the market as it is knowing and understanding oneself. The nonemotional self-analysis of behavior under stresses of risk, reward, and discipline are part of the personal effort any trader has to evolve through, most often by trial and error. I will defer comments on this self-analysis to later and will now focus on the more objective and observable part related to the market.

1.2.1 Lessons from the Market

Let us first see what conclusions we can derive from observing the market as a whole and the behavior of its participants. The most relevant observations can be summarized as follows:

- *Macroeconomic information unfolds gradually, therefore prices do not discount future events immediately.* Why is it the case that at the peak

of the business cycle asset prices do not discount its next through and vice versa? Because no one knows when the next through is coming despite the seeming regularity of business cycles. Things always look so optimistic on the top and so pessimistic at the bottom. This is why we observe long-term trends in all asset prices and yields.

- *The leverage in the market yields a locally unstable system because individuals have finite capital and are playing the game so as to survive the next round.* This instability is increased by the asymmetry between game-theoretic behaviors of accumulation and divestment of risky positions. When you accumulate a position you have all the incentive in the world to tell all your friends, and it is a self-fulfilling virtuous circle as people push prices in “your” direction, thus increasing your profit. This is the epitome of a cooperative game. On the other hand, when you divest, you have no incentive to tell anyone as they may exit before you, pushing prices away from you. This is a classic Prisoner’s Dilemma game where it is rational to defect, as it is not seen as a repeated game. This is why we observe a great deal of asymmetry between up and down moves in prices of most assets, as well as price breakouts and violent trend reversals.
- *There is a segmentation of market participants by their risk-taking ability, their objectives, and their time frames.* Real-money investors have a different attitude to drawdowns than highly leveraged hedge funds. Pension fund managers rotate investments quarterly whereas automated market-makers can switch the sign of their inventory in a quarter of a second. In general, though, each segment reacts in a similar way to price movements on their particular scale of sampling. This explains the self-similarity of several patterns at different price and time scales.
- *The market as a whole has a consensus-building tendency, which implies learning at certain timescales.* This is why some strategy classes or positions have diminishing returns. When people hear of a good money-making idea, they herd into it until it loses its money-making appeal.
- *The market as a whole has a fair amount of participant turnover, which implies un-learning at certain longer timescales.* A new generation of market participants very rarely learns the lessons of the previous generation. If it were not the case why are we going through booms and busts with the suspicious regularity commensurate to a trading career lifespan of 15 to 20 years?
- *There is no short-term relationship between price and value.* To paraphrase Oscar Wilde, a trader is a person who knows the price of everything but the value of nothing.

1.2.2 Mechanism vs. Organism

The above observations do not reflect teachings of the economic orthodoxy based on the concept of general equilibrium, which is a fairly static view of the economic landscape. They become more naturally accepted when one realizes that the market itself is a collection of living beings and that macroeconomics is an emergent property of the society we live in. The society, akin to an organism, evolves and so does the market with it. The complexity of the macroeconomy and of the market is greater than what is implied by overly mechanistic or, even worse, static models.

In thinking about the market from this rather lofty perspective, one is naturally drawn into the debate of mechanism versus organism, the now classic debate between biology and physics. The strict mechanistic view of economics, where the course of events is determined via an equilibrium concept resulting from the interaction of a crowd of rational agents, has clearly not yielded many robust predictions or even *ex post* explanations of realized events in the last 100 years of its existence. Thus despite the elaborate concepts and complicated mathematics, this poor track record causes me to reject the mechanistic view of the world that this prism provides.

The purely organismic view of the market is probably a far fetch from reality as well. First of all, the conceptual definition of an *organism* is not even yet well understood, other than being a pattern in time of organized and linked elements where functional relationships between its constituents are delocalized and therefore cannot be reduced to the concept of a mechanism (that is, a set of independent parts only linked by localized constraints). There are clearly delocalized relationships in the market, and stresses in one dimension (whether geographic location, asset class, regulatory change, etc.) quickly propagate to other areas. This is in fact one of the sources of variability in correlations between different asset classes as well as participants' behaviors. On the other hand, on average these correlation and behavioral relationships are quite stable. Also, unlike in a pure organism, the removal or death of a "market organ" would not necessarily imply the breakdown of the organism (i.e., market) as a whole. For example, the various sovereign debt defaults and write-downs in the past did not yield the death of the global bond market.

1.2.3 The Edge of Complexity

So, intuitively the market is not as simple as Newton equations nor is it as complicated as an elephant or a mouse. Its complexity lies somewhere in between. It has pockets of coherence and of randomness intertwined in time. A bit like a school of silverside fish that in normal circumstances has an

amorphous structure but at the sight of a barracuda spontaneously polarizes into beautiful geometric patterns.

The good thing is that the market is the most observable and open human activity, translated into a series of orders, trades, and price changes—numbers at the end of the day that can be analyzed ad nauseam. The numeric analysis of time series of prices also yields a similar conclusion. The prices or returns do not behave as Gaussian processes or white noise but have distributional properties of mild chaotic systems, or as Mandelbrot puts it, turbulence. They are nonstationary, have fat tails, clustering of volatility that is due to clustering of autocorrelation, and are non-Markovian. A very good overview of the real world properties of price time series is given in *Theorie des Risques Financiers* by Bouchard and Potters.

1.2.4 Is Systematic Trading Reductionistic?

As per the definition above, systematic trading is essentially a computable model of the market. Via its algorithmic nature it can appear to be a more reductionistic approach than discretionary trading. A model reduces the dimensionality of the problem by extracting the “signal” from the “noise” in a mathematical way. A robotic application of the algorithm may appear overly simplistic.

On the other hand, discretionary traders often inhibit their decision making by strong beliefs (“fight a trend”) or do not have the physical ability to focus enough attention on many market situations thus potentially leaving several opportunities on the table. So discretionary trading also involves an important reduction in dimensionality but this reduction is happening differently for different people and times.

1.2.5 Reaction vs. Proaction

A common criticism of systematic trading is that it is based on backward-looking indicators. While it is true that many indicators are filters whose calculation is based on past data, it is not true that they do not have predictive power. It is also true that many systematic model types have explicitly predictive features, like some mean-reversion and market-making models.

At the same time one cannot say that discretionary trading or investing strategies are based solely on the concept or attempts of prediction. Many expectational models of value, for example the arbitrage pricing theory or the capital asset pricing model, are based on backward-looking calculations of covariances and momentum measures. Despite the fact that those models try to “predict” reversion to some normal behavior, the predictive model is normally backward-looking. As Niels Bohr liked to say, it is very difficult to predict, especially the future.

1.2.6 Arbitrage?

Many times I've heard people arguing that the alpha in systematic strategies should not exist because everyone would arbitrage them away, knowing the approximate models people use. The same could be argued for all the discretionary strategies as most of the approaches are well known as well. Thus the market should cease trading and remain stuck in the utopian equilibrium state. Yet none of this happens in reality and the question is why? Probably exactly because of the fact that people do not believe that other people's strategies will work. So as much as it is seemingly simple to arbitrage price discrepancies away, it is less simple to arbitrage strategies away. Having said that, the market system in itself is cyclical and, as mentioned above, strategies get arbitrated away temporarily, until the arbitrageurs blow up all at the same time because of their own sheer concentration of risk and the cycle restarts with new entrants picking up the very valuable mispriced pieces.

1.2.7 Two Viable Paths

Viewing trading and the market from this level yields a positivist view on the different ways to profit from it. The discretionary traders see in it enough complexity to justify their approach of nonmechanizable intuition, insight, and chutzpah. The systematic traders see in it enough regularity to justify their approach of nonemotional pattern matching, discipline, and robotic abidance to model signals.

Which approach is right then becomes a matter of personal taste, as the edge of complexity the market presents us with does not allow for a rational decision between the two. In fact both approaches are right, but not necessarily all the time and not for everyone. Of course the Holy Grail is to be able to combine the two—to become an übertrader who is as disciplined as a robot in its mastery of human intuition.

This book of course does not offer the Holy Grail to trading; intuition and insight are quite slippery concepts and highly personal. There is no one way. But this work is not interested either in focusing on the same old mechanistic techniques that appeared at numerous occasions in books on systematic trading. It aims at moving further afield toward the edge of complexity, by giving enough structure, process, and discipline to manage a set of smarter, adaptive, and complex strategies.

1.3 THE BUSINESS OF TRADING

If, as was derived in the last section, there is no a priori rational way to choose between discretionary and systematic trading paths, one should then

aim at objectively comparing the two approaches as business propositions. Seeing it this way will lead naturally to a choice based on the trader's own psychology; that is, which of the two business propositions is the most compatible with the inner trust of his own ability to sustain and stand behind that business activity over time.

The goal of a business is to produce a dividend to its stakeholder. Any sustainable business is built on four pillars:

1. *Capital*: provides the necessary initial critical mass to launch the business and sustain it through ups and downs
2. *Product*: the edge of the business, the innovation relative the rest of the competition
3. *Factory*: the process by which the products are manufactured, which is an integral part of the edge itself
4. *Marketing*: the means by which information about the product reaches the outside world and helps replenish the capital, thus closing the loop

Both discretionary and systematic trading businesses should be seen in the context of those necessary contexts. Of course trading is not per se manufacturing of anything other than P & L. So the product is the trader's edge or algorithm and the factory is the continuous application of such trading activity in the market. Marketing is the ability to raise more capital or assets under management based on performance, regulatory environment, or good looks. Here the *trader* can mean an individual, a group, or a corporate body.

So let us do a comparison between systematic and discretionary trading, keeping in mind the above concepts.

1.3.1 Profitability and Track Record

Before one even starts looking at the individual pillars of business, can one say anything about the long-term profitability of the two trading styles? This is an important question as it may provide a natural a priori choice: If one type of business is dominantly more profitable than the other then why bother with the laggard?

Interestingly it is a hard question to answer as the only objective data that exists in the public domain is on hedge fund and mutual fund performance. Any of the profitability data of bank proprietary desks is very hard to come by as it is not usually disclosed in annual reports. Also the mutual funds should be excluded on the basis of the fact that their trading style is mostly passive and index-tracking. This leaves us with comparing discretionary to systematic hedge funds.

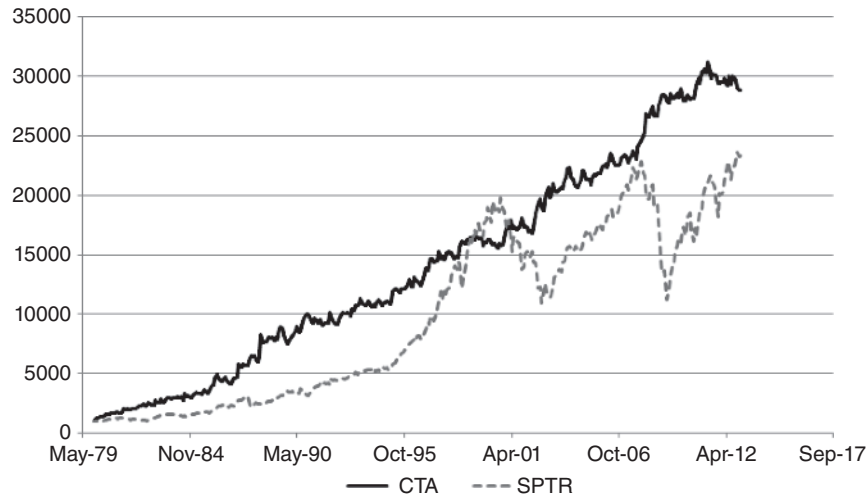


FIGURE 1.1 HFR CTA Index versus SP500 Total Return Index

In both camps there is a wide variety of underlying strategies. In the discretionary camp the strategies are long-short equity, credit, fixed-income relative value, global macro, special situations, and so on. On the systematic side the strategies are commodity trading advisors (CTAs), statistical arbitrage, high-frequency conditional market-makers, and so on. What is the right comparison: absolute return, assets under management (AUM)-weighted return, return on shareholders equity? Because private partnership is the dominant corporate structure for hedge funds, the return on shareholders equity is not a statistically significant comparison as far as publicly available data is concerned. Hence one has no choice but to compare strategy returns. As on average the fee structure is similar in both camps, one may as well compare net returns to investors.

Figure 1.1 shows the comparative total return on the Hedge Fund Research CTA Index and the total return on the SP500 stock index. Table 1.1 shows the comparative statistics of major Hedge Fund Research strategy indices from 1996 to 2013.

Some of the earliest hedge funds were purely systematic and have survived until now despite the well-known attrition in the hedge fund industry as a whole. Many commodity trading advisors and managed account firms have been involved in the systematic trading business for at least 40 years. Their track record represents an interesting testament to the robustness of the systematic approach, from the performance and process perspective. Also systematic strategies have in general low correlation to

discretionary strategies and to other systematic strategies, especially classified by time frame.

In conclusion one sees that the major strategy types tend to be quite cyclical and that there are sizable up-runs and drawdowns in each class, be it in the discretionary or systematic camps. Thus it is difficult to draw any conclusions on the dominance of either style on the basis of profitability alone.

This brings us back to our exploration of how the two styles compare in the context of the four business pillars mentioned above, in the order of product, factory, marketing, and capital.

1.3.2 The Product and Its Design

Research and information processing are the crux of the product's edge for the trader. A trading strategy is first and foremost an educated idea on how to profit from certain situations, be they ad hoc or periodic, and how to mitigate losses from unexpected events. It requires an ability to gather, process, and research a large quantity of information.

Information In the discretionary world, this information is categorized into the following seven areas and the trader forms an intuition based on this set in order to pull the trigger:

1. Macroeconomic
2. Political
3. Asset-class specific
4. Idiosyncratic to a company
5. Specific to a security (share, bond, etc.)
6. Price and transactional
7. Flow and holdings

The majority of the time in the systematic world, the information required is limited to the price and transactional and in rarer occasions on the holdings and flows (such as the Commitment of Traders report in the futures markets). Most of the systematic models base their decision making on the extraction of repeatable patterns from publicly available data on prices and executions. The statistical significance of such patterns is derived from simulation (the action of back- and forward-testing).

Both activities are clearly information-intensive but this intensity manifests itself in quite different dimensions. The discretionary style requires processing of a broad scope of nonnumerical data, and traders read and rely on a range of broker and analyst research along with continuous news and political analysis. A lot of useful information is also seen in the flow and holdings

that are obtained via brokers, that is, who are the transacting participants and how much. This in itself implies that discretionary trading is difficult to do solo and often requires teams of people to digest all the information flow. Interestingly, some firms have started creating numerical sentiment indices based on textual and voice news flows, a technique used initially by intelligence agencies to discover subtle changes in political rhetoric.

For the systematic style, the dimensionality of the information is much lower; the models are in general only interested in price or tick data but they require a continuous feed and automated processing of this data at high speeds, especially in the current context of the ECNs. This means that from a technological perspective, especially for high-frequency business, the required connectivity and throughput needs to be large. This in general has cost implications.

Most systematic models also require prior and continuous recalibration, thus large databases of historical data need to be kept for research purposes.

Research Information is useless if it cannot be interpreted in context, be it intuitive or model based. To be able to form such an educated view, some research needs to be performed on the relevant data.

In the discretionary context, most useful research falls into (1) political and regulatory analysis, (2) macroeconomic analysis, (3) asset-specific research, or (4) quantitative research. Many investment banks and institutions have large departments focused on macroeconomic analysis and asset-specific research. Discretionary traders or teams have access to such research via prime brokerage relationships and those costs are implicitly absorbed into trading and clearing fees. A few smaller private firms run by former bank or government institutions officials provide political and regulatory analysis and macroeconomic analysis for fees and also use their former contacts to introduce clients to current central bankers, finance ministers, and other officials. Such relationships are invaluable for certain strategies such as global macro, where fund managers constantly try to read between the lines for changes of moods or rhetoric in order to form their own expectations on upcoming policy moves. Thus a lot of research that is valuable for discretionary trading is already out there. It needs to be gathered, filtered, read, and distilled to be presented to the portfolio managers. Large discretionary hedge funds hire in-house economists and analysts to do such work but many operate just using publicly available and broker research.

There is a subset of discretionary strategies that is driven by quantitative modeling. Fixed-income relative value, long-short equity, and volatility strategies are such areas, for example. Each require a fair amount of advanced mathematical techniques, pricing tools, and risk management tools. Although there is commercially available software with standard libraries for pricing options, interpolating yield curves, or handling large-scale

covariance analysis, the vast majority of quantitative discretionary operations employ in-house quants to write a series of models and pricing tools as well as to maintain the relevant data and daily process. This has clear cost implications on such businesses.

The systematic approach is entirely research-driven and in a very direct sense research innovation is the backbone of the business. The principal areas of research fall into the following four categories:

1. *Individual Models.* The goal is to produce a set of diversified robust trading agents that exploit various repeatable price and trade patterns. Various techniques of back- and forward-testing are employed for this goal. It is the key area for the success of the whole business. It is the focus of Part One.
2. *Adaptation of Model Portfolios.* The goal is to produce an automated allocation rule for a portfolio of models by studying the persistence of behavioral regimes of individual models. It is an important area for integrated risk management in the high-frequency trading domain. Part Two is dedicated to some of my findings in the matter.
3. *Trading Costs Minimization.* The goal is to minimize market impact from model execution by slicing the trades according to various execution algorithms that derive mostly from liquidity distributional analysis. This is explored in Part Three.
4. *Trading Process Optimization.* The goal is to optimize the trading process from the perspective of computational efficiency as well as to ensure fast recoverability from downtime. It is a vast area to which the practical Part Four is dedicated. It encompasses the design of low-latency order management systems and their coupling with various model engines, domain models for state persistence and recovery, distribution of computational tasks among components, and so on.

These four categories are closely intertwined in automated systematic trading and demonstrating this concretely is an important feature of this book.

1.3.3 The Trading Factory

Process Designing and implementing a disciplined trading process on the basis of either computable or subjective signals is key to the success of the business of trading. The process presupposes an infrastructure and a technology optimized for the production of the trading widget. It is not enough to have a good widget idea; one also has to be able to manufacture it efficiently. Of course, having a great factory producing widgets that no one

wants is a waste of time and money. But as much as great trade ideas or strategies are necessary, they are not sufficient if not implemented correctly. The underlying processes of discretionary and systematic businesses present many similarities but also major differences, as we will show now.

In the discretionary world, choosing the winning set of human traders is key. The traders have to have at least the following four features, with the last three criteria being essentially a strong self-discipline:

1. *Profitability*: Ability to generate revenues in different market conditions
2. *Predator Mentality*: Proactive trade idea generation stemming from continuous information processing accompanied by aggressive sizing into good opportunities
3. *Ego Management*: Proactive risk management and survival skills
4. *Clear Head*: Knowing when not to overstay one's welcome in the market and take time off when the picture is not clear

Several successful traders have published honest and objective self-analyses of their occasional failings in instituting such discipline, courage, or focus and drew lessons for the benefit of the whole trading community. Of course longer-term survival let alone profitability hinges on the discipline of applying the trading process as per the last three criteria.

As mentioned in the previous section, the systematic business is research-driven. The principal goal of that research is to produce a portfolio of profitable models. It implies that the continuous fostering of innovative research is a key element to the process and to the success of the business. Finding a set of robust models in the systematic world is equivalent to hiring a desk of good traders in the discretionary world.

The systematic approach a priori formalizes a lot of the individual trader's discipline as models are run continuously, have embedded stop-losses and profit targets, and can be scheduled to be turned on or off during certain periods. The trading process is thus run as an algorithm. The key four features for success are similar in nature to those mentioned above:

1. *Profitability*: A diversified set of models that are profitable in various market conditions
2. *Continuous Monitoring*: The models are continuously processing data and output either position or order changes, thus opportunities are exploited to the maximum, 24 hours a day and across many markets
3. *Dynamic Risk Management*: Portfolios of models have an embedded dynamic sizing algorithm that controls the exposure as a function of the performance of each model and the portfolio as a whole

4. *Model Fitness*: A higher-level feedback mechanism compares a portfolio of possible models and dynamically chooses a subset to trade on the basis of a fitness measure, thus models are demoted and promoted dynamically from a prior set of potential candidates

The systematic trading process is much more involved than the discretionary one as by its nature it is automated. The increased complexity comes from the fact that many things that are second nature to humans are actually hard to implement in software (for example, automatic recovery mechanisms from data disconnects or loss). It is a technology-driven process as the technology implements the factory element. Thus from the technology perspective, the systematic business requires an investment into software and hardware much larger than for the discretionary business. We focus in Part Four on the analysis of the various key elements one needs to master to put such a process in place.

Cooperation We have come to another important aspect of the nature of communication and cooperation within the two businesses. In a discretionary hedge fund, especially in areas like global macro, there is a tendency to encourage trade diversification by discouraging communication between various traders. This is a noncooperative game scenario and some funds push it even further by encouraging traders to compete for the biggest risk allocation from the same pot thus creating potential friction, jealousy, and mutual dislike between people.

Interestingly, on the systematic side such selection is done implicitly by the higher-order model feedback mechanism. So the noncooperative game is left to the machine and one does not hear models screaming or squealing when they get demoted. The research process, though, has to be a cooperative game where cooperation between team members serves the exact purpose of creating a diversified portfolio of models. Efficient systematic research has to be run on the exemplar of academic institutions where people are given enough leeway to innovate and learn from communication with each other, and are driven by the common good of cooperative success.

There is another important cooperative game going on in the systematic trading business. It is the natural synergy between the research, development, technology, infrastructure, and monitoring teams. Research needs an optimized implementation that in turn needs efficient technology run on a robust infrastructure that is being monitored continuously. All areas need research to come up with money-making models to produce cash and sustain the whole food chain. The success of some large systematic funds is corroborated by my own knowledge of the way such cooperation had been instituted within them.

1.3.4 Marketing and Distribution

The differences in the products and processes discussed above imply differences in the approach to marketing and branding of discretionary and systematic strategies. One could say that the brand of a discretionary trading business falls more into the craft category, whereas the brand of a systematic trading business falls more into an industry category. The last remark could be justified from our analysis of the process, not the product. Both product design processes are crafts, coming from accumulated intuition of traders on one hand and researchers on the other. From a marketing perspective, the element of skill is crucial in both worlds.

One could argue that it is somewhat easier for a newcomer to launch a systematic fund rather than a discretionary fund. The crucial point that comes in all capital raising discussions is the ability to produce a credible track record. It is difficult for discretionary traders to have a track record unless they have traded before, which is of course possible only if they traded their own money or could take their track record from a previous firm (a very tricky exercise in itself). Thus the majority of discretionary traders start in market-making and other sell-side careers then graduate to a proprietary trader status. Only then can they start to build their independent track record.

The situation is quite different in systematic trading as there is a reasonable degree of acceptance among allocators of back-tested and paper-traded track records. This of course supposes that the simulated net asset value (NAV) contains a realistic (or, even better, pessimistic) assessment of transaction costs, scalability, and sustainability of the market access and the trading process in general. The discussion then focuses on how this track record was generated and whether there was a risk of over-fitting and using future information in the process of building the models.

Once the fund has been launched, let us compare the hypothetical clues to answer the four main types of questions clients would usually ask while doing their due-diligence assessment:

1. *Can your profitability be sustained?* In the discretionary world, everything hinges on the ability of the head trader to keep performing, whereas in the systematic world, it is all about the quality and innovation of the head of research.
2. *Is your market risk management robust?* In the systematic world, risk management is embedded in the model and portfolio processes and can be explained very clearly. In the discretionary world, it is usually harder to formalize and a fair amount of due-diligence time goes into drilling the head trader on the reaction to various past and hypothetical

situations, as well as on the discipline with which the trading team abides by the constraints imposed on them by the risk management team.

3. *What is your operational risk?* In the discretionary world, one source of operational risk is the key man risk. Once the head trader gets run over by the proverbial bus so goes the fund generally. In the systematic world, once the automated process has been put in place the focus is on its sustainability and resilience. One has to show that the process has live disaster-recovery sites and can be rerouted or delocalized if need be to protect from data loss and market disconnects. The aspects of accumulated data recovery as well as people relocation in case of premises incapacitation are the same in both worlds.
4. *What is your capacity?* This of course depends on the time frame of trading given the liquidity of products traded. In the systematic world, one could argue that it is easier to estimate the impact of increased trading volume on transaction costs because most products are exchange traded or have excellent price and volume transparency.

In conclusion, I believe that it is somewhat easier to start a systematic fund but it requires a similar marketing effort as for a discretionary business.

1.3.5 Capital, Costs, and Critical Mass

Enlightened by the comparison of the three functional parts of the business, we now come to the crucial questions of necessary initial capital and of running capital for operations. Of course we need to compare the two businesses *pari passu* as far as size and revenue goals are concerned. We use the example of hedge funds because they are stand-alone entities where all costs and revenues can be objectively estimated.

How much is needed to start the business? In 2010, the realistic critical mass of initial capital needed to start a hedge fund business is north of \$50 million and better at \$100 million. The main reason is the structure of allocators—funds of hedge funds, asset managers, and family offices. Most of them will rarely look at a target with AUM below \$50 million because they do not want to participate more than 10 percent in any fund. This helps them to reduce the risk of concentration of other clients in the fund if, of course, the other clients are also invested less than 10 percent each. As they get lesser fees than hedge funds themselves, an investment of less than \$5 million is not worth the costs and time of due-diligence process.

It is actually not a bad thing for the fund itself as it forces it to be diversified in its client base, so that losing one client will not put the fund in

jeopardy. But then the question comes down to the classic chicken-and-egg: How would one start a fund in this difficult environment? One needs to find a set of seed investors, hopefully all at the same time, a lot of performance luck, and a lot of marketing effort! This is the same across various strategies and the systematic business is no different from the discretionary in this respect. Thus the barriers of entry are quite high for either type of stand-alone trading business.

How much is needed to maintain the operations? As the seeders invariably take a cut of the economics, the resulting revenue is probably not the usual 2 percent management:20 percent performance fees structure but closer to 1 percent:15 percent. Assuming raising \$50 million of AUM the first year, the realistic management fee revenue is around \$500,000.

Certain types of trading styles can be perfectly run on minimal infrastructure consisting of the head trader (Chief Investment Officer), a middle-office person (Chief Operating Officer and Chief Risk Officer), and a marketing and client relationship person who can also hold the title of Chief Compliance Officer. Those four functions combined into three people tick the minimally accepted boxes as far as institutional allocators are concerned in their goal toward reducing operational and key-man risks. Other functions can be outsourced, in particular many back-office functions of control, fund administration, IT support, and legal support. The costs of renting a furnished office space plus utilities of course varies but be it offshore or on-shore, it comes roughly to at least \$50,000 per year. The IT and legal support costs, communications (phone and Bloomberg feed), and directors' and officers' insurance also come to at least \$50,000 a year but may be larger. Adding travel and entertainment costs puts the total pre-salaries expenses at around \$150,000 conservatively. The salary expenses then pretty much take up the rest of the fees, with usually \$150,000 to the COO, \$150,000 to the marketing person, and the rest to the head trader, who probably is the sole partner working for the upside call option. The business can survive one or two years on this without making extra trading revenue but if it does not, clients will usually pull the capital anyway. So the \$50 million is indeed the low end, the necessary but not always sufficient critical initial mass.

There are several styles on the discretionary and systematic sides that are doable under the above setup. Those styles usually do not require either a large amount of assets to trade nor a high frequency of trading. Styles that would be difficult to fit in this minimal mold are, for example, equity statistical arbitrage, high-frequency systematic trading, global emerging markets strategies, and global credit strategies. These require more people trading more markets or a more complicated technology that needs to be overseen and run by more people.

For example, the high-frequency systematic business requires at least a team of two full-time researchers and two full-time technologists in addition to the minimal model above. From my personal experience in building such a business, this is required in order to ensure operational stability in a 24-hour operation. That automatically increases the costs by roughly \$500,000 the first year and means that a realistic stand-alone critical mass for such a business becomes \$100 million AUM. The statistical arbitrage style has very similar features. Of course, one could buy an off-the-shelf integrated solution that provides in the same package a financial information exchange (FIX) engine, connectivity setups, a complex event processor, databasing, and an ability to program in your models. Is this really cheaper than developing all the infrastructure in-house? Those packages are actually priced not far from the salary of two technologists. One still needs a technologist to maintain the system and a researcher for innovation, so the off-the-shelf solutions are not dissimilar in costs to building everything in-house.

Cost of Capital and Leverage Another important consideration is the cost of running trading positions. Here the answer tends to be more clearly in favor of systematic strategies simply because they tend to use very liquid low-margin and exchange-traded products. The leverage in the most liquid products can be up to 50-to-1 even taking into account the extra margining imposed by prime brokers. On the contrary, many discretionary strategies exploit risk concessions that arise from liquidity premiums and those strategies essentially benefit from providing liquidity to the market. Some products, such as emerging market credit instruments and insurance-linked securities, have no leverage at all and one has to pay the full price to participate in them. So in general CTAs and other systematic trading strategies are quite long cash on which they can earn a premium whereas credit and arbitrage funds mostly borrow cash which sometimes can be costly.

How much is locally too much? This means how much capital could be allocated to a single strategy style above, which the efficiency starts decreasing because of trading costs? In general scalability is directly proportional to the liquidity of the market traded and inversely proportional to the frequency of trading. Many discretionary strategies actually extract alpha by buying risk concession in illiquid instruments and hence have limited scalability. Systematic high-frequency price-taking strategies have limited scalability because they aggress the market and move it in the process. It is difficult to draw general conclusions and we will discuss this specifically for strategies on stocks, futures, and foreign exchange (FX) markets in Part Three.

How much is globally too much? This means how much global commitment to the same class of strategies yields a decrease of efficiency for

everyone involved? There is no hard answer to this but the study of damage from various strategy bubbles—the 1997 crash of carry trades, the subsequent crash of LTCM, the 2006 crash of Amaranth, the 2007 crash of statistical arbitrage, the 2008 crash of long carry in credit—all point out that this number is growing. The monetary inflation and quantitative easing only add to the fire. The next bubble is probably going to be bigger than the sum of the previous ones.

* * *

In conclusion, the systematic business has an advantage on the discretionary from the lower capital usage, but can have a disadvantage of higher up-front costs if one wants to compete at the cutting edge of technology and research.

1.4 PSYCHOLOGY AND EMOTIONS

The previous section aimed at showing that from an objective business perspective the systematic and discretionary trading activities have a lot in common. They all hinge on finding and maintaining efficient trading agents and instituting the adequate discipline for the trading process. Also the costs to start and maintain the businesses are comparable at comparable scales of revenue. Another section argued that longer-term profitabilities are on average the same. From a more conceptual and philosophic perspective based on our analysis of the market complexity, the two approaches are equally viable.

Thus we come to a subjective and personal point: What is the style of trading that is compatible with one's psychological makeup? In this section we try to suggest what trader psychologies best fit the two trading styles. My personal choice was made a long time ago in favor of systematic trading.

1.4.1 Ups and Downs

As simple as it sounds, the crucial psychological skill is the ability to deal with losses and gains. The volatility of the profits on the trading book is a natural feature of the trading process. The psychological ability to unemotionally deal with the upside and the downside volatility is the crucial aspect of the maintenance of the trading process.

The systematic approach embeds a sophisticated money management strategy. Not only does each model have its own sizing, stop-loss, and profit-taking rules, but the portfolio of models itself is managed on the basis of global sizing rules that allow it to deal with variability in correlations between the individual model returns. Also, automatic selection rules based

on the fitness of each model can be introduced. Model-specific rules are explored in Part One and portfolio and fitness rules in Part Two. Part Three explores the important issue of slippage that must be taken into account to produce realistic downside expectations during periods of stress.

The discretionary approach does not formalize explicitly any such money management rules. Traders take a view on how much they can allow themselves to lose on a particular position and have expectations of how much they can gain. The risk manager is then responsible for making sure that the exposures do not breach certain levels of value at risk (VAR) or some other measure based on historical covariances.

The unformalized approach presents positives and negatives. On the negative side, there is no automated stop-loss. Often traders hang on to their positions because they “know” they are right. It is then only a matter of time before either the management or the traders themselves throw in the towel, and the damage is a multiple of what could have been had a hard stop-loss been respected. Also, because all the VAR calculations are backward-looking and have a large lag, sudden correlational shocks cannot be dealt with in a timely fashion. On the positive side, the trader’s human judgment and intuition can sometimes save the position from a stop-loss forced in by silly market behavior stemming from overreaction to some irrelevant news or rumors. Just as important is the ability to recognize an outsized opportunity relative to historic data and stick with the position for much longer than by respecting an a priori computed profit-taking level.

1.4.2 Peer Pressure and the Blame Game

Emotions always are in overdrive in situations of stress. Choosing the systematic approach exposes the participant to the criticism from his discretionary peers that the approach is formalized in a finite set of rules. When you lose money, your models must be wrong or too simplistic. On the other hand, choosing the discretionary approach, which is opaque to a formal analysis, yields equally strong criticisms from the systematic peers. When you lose money, you must have a lack of discipline or focus or are a macho, fighting obvious patterns and thinking you are smarter than the market.

The ensuing soul-searching comes down to a question: What aspect of your psyche do you trust more, the computational or the intuitive, the right or the left side of the brain?

1.4.3 Trust: Continuity of Quality

On what basis does one trust a trading process? Intuitively, we trust based on two traits:

1. *Quality*: the ability to deliver and survive in different stress scenarios
2. *Continuity*: the ability to “wake up every morning and bite the ass off a bear,” as John Gutfreund allegedly said on the Salomon Brothers trading floor

On the systematic front, the quality aspect comes down to being able to organize and deliver solid, innovative research in profitable strategies. The continuity comes down to being able to automate the application of those strategies in a dynamic and adaptive portfolio context. So the trust comes down to the ability to deliver an efficient research and development (R&D) process. This book’s goal is to provide a very solid base for such a delivery.

On the discretionary front, quality is about the trader’s instinct and discipline. Continuity is about knowing how to choose quality traders and organizing a reward/punishment structure that retains the best over time. Quality nevertheless is a dominant feature as the complexity of the process is lower relative to the systematic world.

Thus from a psychological perspective the discretionary process is more individualistic. As we noted above, it does not have to be a cooperative game at the level of a group of traders as it arguably helps to diversify ideas and risk. On the systematic side, the process has to be cooperative, first between researchers themselves, then between different groups—research, technology, infrastructure, and monitoring.

1.4.4 Learning from Each Other

Given that we are all competent business people and can organize an efficient money-making factory, the question then boils down to which we enjoy more, the thrill of the unexpected just before the non-farm payrolls come out or the quiet humming of our servers crunching terabytes of tick data? It is basically an affinity to mathematical abstraction versus to human language.

With different psychological makeups compatible with systematic and discretionary trading styles, this section suggests what actual elements one style of trading can learn from the other.

The primary intuition about patterns that can be systematically exploited comes from the discretionary side. It is ultimately the analysis of participants’ emotions toward making or losing money that gives clues as to what patterns are exploitable at different price and time scales. Another most important point of adaptation comes from observing human and animal behavior toward problem solving and Part Two of this book explores several avenues to systematize it in a broader context. Thus the

behavior of the discretionary traders is really useful to understand for the systematic researchers.

At the same time the inherent discipline toward money management and the robotic trading process naturally present in the systematic trading are useful role models for the discipline of discretionary traders. Also, the knowledge of systematic models is useful for discretionary traders for predicting stress, trend reversal, and breakout levels.

Thus the two trading disciplines can be seen as a coevolving set. In fact the whole thing did not start as such because the marketplace was initially dominated by discretionary traders. But especially since the coming of the technological mega-trend and computerization of major markets, the landscape is changing by the day with more and more automated systematic trading strategies coming online and in some markets even starting to dominate the traded volume at the expense of discretionary participants.

The question is, who is going to learn faster? The humans with their zillions of neurons and synapses or the cloud-based parallel supercomputers? We are definitely living in very exciting times where the arms race of cold war weaponry has been crowded out by the arms race of trading bots!

1.5 FROM CANDLESTICKS IN KYOTO TO FPGAs IN CHICAGO

Understanding systematic trading in a historical context is interesting and important as it sheds light on the natural progress from the very beginnings of data and pattern representation to the modern highly parallel adaptive processes connected directly to exchanges and crunching data in microseconds. An overview of this history is given here that shows its constant coupling with the developments in the relevant scientific and technological spheres.

Systematic trading as a style has been in existence since the advent of organized financial markets and the associated record keeping of prices and transactions data, long before the introduction of computers and even of the ticker tape. Its origins can be traced to the sixteenth century rice traders in Japan who introduced tools to represent price activity in a visual way that lead to the discovery of certain patterns, often bestowed with poetic names.

Data representation is a very important part of an effective analysis of a situation. It quite often is achieved through the concept of compression, or in other words the removal of irrelevant details. Of course the whole concept of relevance is tightly linked to the goal of the analysis, in other words to the extraction of the signal from the noise. Part One starts by covering various

ways data can be usefully represented for systematic trading and explaining which features are retained and which features are compressed away.

The field was then taken to the next level by Charles Dow in the early 1900s in the study of the U.S. stock markets. Driven by the fast expansion of industry and transportation in the United States in the late nineteenth century, the New York Stock Exchange (NYSE) had acquired a prime position as the center for organized exchange of risk and price discovery, a position it still holds. Charles Dow introduced various indicators based on industry and transportation sub-indices, moving averages, and various other filters, and was the first to formalize certain trading rules coming from the relationships between those indicators. Several people refined those trading rules, resulting in flamboyant trading careers for the likes of Jesse Livermore and W. D. Gann.

We continue Part One by discussing the concept of an indicator and various examples of them. Indicators are filters that presuppose a choice of the data representation methodology and are the building blocks of the underlying signals to the systematic trading models.

Risk management rules were formalized from the observations of widths of trends and extensions of common price patterns. Interestingly, the origins of the money management rules most commonly used in modern systematic trading come from a different crowd than the community of buy-side speculators and asset managers. Namely, the increasing importance in the mid-twentieth century of the highly leveraged futures markets in agricultural products in Chicago attracted a large crowd of pit traders who were mostly scalpers and market-makers, trading hundreds of times per day for a couple of ticks here and there. By the sheer frequency of their trading they had to adopt very strict money management rules to survive till the end of each trading day. Those tricks of the floor trader have been formalized into numerical rules on stop-loss and profit-taking based on the volatility and the liquidity of the market and applied with almost robotic discipline of their implementation.

The design of the indicators is the first step toward building a mechanized strategy. They provide the signals on which the triggers to enter or exit positions are based. Most of the indicators are filters that have an inherent lag and may expose the strategy to the risk of being too slow to react to low probability market moves. Thus a money management overlay is warranted in the majority of cases to build a better trading strategy but at the expense of increasing its complexity and potentially its brittleness.

The advent of increasingly cheap computing has created an avenue to test such more complex strategies. A whole cottage industry of trading systems, indicators, and methods resulted from this in the 1980s and 1990s. After discussing several strategy types, Part One focuses on some key

implementation aspects of the design and testing of trading systems using the full capabilities of modern computing techniques. In particular, a unified representation of strategies as finite-state machines is introduced and scalable back-testing and forward-testing engines are built on that basis.

Starting in the 1960s, new conceptual developments in areas of control and adaptation gradually coevolved with the uptrend in technology and culminated in what was labeled in 1990 as artificial life. Initially, the real defining driving force was the more ambitious endeavor of artificial intelligence (AI) started in the 1950s by John McCarthy, continuing the intellectual lineage of Alan Turing and John von Neumann. But after an initial jump in the progress toward building intelligent machines able to emulate and surpass humans, a plateau was reached in the late 1980s. The main approach at that point was mainly top-down, trying to automatically create a semantic or visual analysis of the surrounding world. It of course brought a lot of benefits to the progress in pattern recognition, computer vision, and graphics as well as attempts at ontological analysis that ultimately links us now to the semantic web. But the ultimate goal of programmed intelligence was still quite far off, so the whole field found itself in a stasis.

A defining moment in progress came from Rodney Brooks at MIT in the early 1990s. He essentially turned the top-down approach upside-down and designed very efficient robots based on a new concept of control. His bottom-up concept of control is based on a subsumption architecture that receives signals from a set of concurrent sensors, ranks the signals' importance, makes the decision, and sends it to the actuators. The control architecture is itself evolved through trial and error via a genetic algorithm or a reinforcement learning scheme that embeds a concept of fitness of the robot's behavior. This approach is distributed and reactive in nature rather than monolithic and proactive. Brooks demonstrated much better results than the original top-down approach on several important examples. It has yielded progress in many other fields, the most important ones in my mind being the distributed agents systems, swarm computing, and of course the optimization of software design patterns and operating systems to tackle parallel processing, multithreading, and the associated concurrency problems.

Despite the fact that the new approach has not solved the old problem of artificial intelligence, it refocused the research community on tackling other no-less-interesting problems and the field was coined artificial life (AL) by Chris Langton in 1990. With a much better understanding of evolutionary computing techniques such as genetic algorithms, genetic programming, and reinforcement learning, a whole new door was opened to breed and play with lifelike creatures that evolve through adaptation and learning and provide a test bed for both Darwinian and Lamarckian ideas. One of the main observations from that exercise is the natural emergence of complexity incarnated

into a higher organizational order, an effect already observed through the study of nonlinear dynamic systems a couple of decades before.

At the same time that these great theoretical advances in adaptation were happening, the global financial markets were going through their own technological revolution. Many exchanges, starting with Eurex, were moving gradually into electronic market access and automated matching engines. The trend accelerated when the competition to major exchanges came from new electronic commerce networks (ECNs) in the late 1990s with competing liquidity at faster access times and lower prices. Exchanges at the end of the day are money-making institutions that thrive on high volume of transactions. So the initial trickle of business away from traditional open outcry in the pits to the screens accelerated faster than some exchanges could predict and sometimes handle. Added to this, technology enabled several cost-saving exercises for large institutions in the form of dark pools that are explored in Part Four.

This innovation has increased the share of systematic trading in three ways, all taking advantage of this technological trend:

1. A new breed of systematic trading strategies appeared mostly in the higher-frequency domain, driven as much by the then existing players as by the cohort of locals leaving the pits for the screens.
2. Major sell-side institutions have implemented automated market-making engines and introduced several new algorithmic execution techniques, replacing many locals.
3. Many hedge funds and bank proprietary desks have increased their share of electronically executed systematic risk taking as the barriers to entry have been decreasing thanks to the advent of electronic connectivity providers and price aggregators.

It is difficult to estimate what exact proportion of global electronic volume is originated by systematic strategies but some anecdotal evidence suggests that certain specific markets have already passed the 50 percent mark thanks to the dominance of automated market-makers.

Of course, like any other fad, this technological trend is feeding on itself. According to Ray Kurzweil, we have not seen anything yet as it is feeding in a super-exponential fashion! While the singularity is a few years away, we still need to adapt, but at an increasingly faster pace.

Part Four focuses on the design of an infrastructure that supports efficient low-latency systematic trading with modern electronic exchanges and ECNs. That infrastructure parallels the architecture of such exchanges and also contains an internal matching engine for competing model orders, that is, a mini dark pool. The increasing dominance of fast electronic

transactions came in hand with many new technological advances in hardware and software. Namely, we discuss how the following six innovations, among others, naturally fit into the design of our integrated low-latency trading infrastructure:

1. Multithreading and concurrency design patterns that allow the emulation of parallel processing
2. Distributed in-memory caching that solves several state persistence and fast recovery issues
3. Message-passing design patterns that are the basis for a distributed concurrent components architecture and help reduce latency
4. Web server technology that allows remote control of components
5. Universal communication protocols that ensure smooth data passing between counterparties, such as the FIX protocol
6. FPGAs (Field Programmable Gate Arrays, which are programmable chips) and GPUs (Graphics Processing Units, which are highly parallel graphics chips) that help optimize certain algorithms

While the fields of electronic finance and artificial life are both experiencing strong independent growth, one cannot exactly yet call them co-evolving entities. The fundamental goal of this book is to suggest avenues to bridge that gap. Philosophically, given that the discretionary world is driven by humans, why not endow our systematic machines with better learning and adaptive skills? The artificial life paradigm is giving us a first genuine step in that direction. It does not yet give us automated foresight, the holy grail of artificial intelligence. We will need to wait for that one a bit longer.

Part Two provides several concrete examples where artificial life techniques can be profitably applied to finance by building robust adaptive systematic trading strategies. Adaptation is studied both from the viewpoint of endowing an individual with more complexity (akin to the subsumption architecture discussed), as well as from an automated choice of individuals from a population (akin to a genetic algorithm). Appropriate concepts of strategy fitness are introduced. Higher-frequency systematic strategies present the best test-bed for such concepts because they quickly generate large and statistically significant trade samples.

Of course implementing concretely such concepts requires nontrivial machinery, and the design of an integrated low-latency trading architecture in Part Four is tuned to the task of processing parallel swarms of adaptive strategies.

Coming back to Earth from musings into the ever-bright future, let us comment on the less positive features that the electronic trading dominance

has left the world with. As mentioned above, the traditional pits with open outcry had almost disappeared by the early 2000s. This fundamental evolution of access and reaction to information has changed some of the long-established features of how the market operates. While increasing efficiency locally it has also introduced a share of global instability. The resulting dynamic is becoming more akin to an arms race where firms compete for speed of data access and delivery for exchanges and clients.

The main change in some markets is that specialists, that is, appointed market-makers, have lost a lot of ground to a set of human and automated agents doing noncommittal conditional market-making and who are not held to provide a two-way orderly market. Automated conditional market-makers provide two-way or one-way liquidity as they see fit and can switch off automatically in situations of stress, leaving the bulk of the flow to more traditional specialists who cannot normally deal with it.

Such structural change is most likely the main cause behind shocking self-fulfilling events like the Flash Crash of May 6, 2010, that spurred a lot of soul-searching as much from the Securities and Exchange Commission (SEC) as from the algorithmic firms community. After initial talk of fat finger or other malfunction, none of the subsequent analysis of transactions and operation logs have managed to pinpoint such kind of cause. It is most probably the automated conditional market-making engines that pulled bids when they saw an increasing selling flow from automated portfolio hedgers and other stop-loss algorithms. It is therefore very difficult to argue that the nature and cause of the Flash Crash was that different from the 1987 meltdown. Yet the recoil from the bottom was indeed different because it happened much faster. That fast recovery was probably helped by a set of automated high-frequency momentum strategies that went long, a feature that was much less present in 1987.

Of course such events will not go without consequences on regulation and self-regulation of the markets. It is well known that some participants do flood the markets with masses of orders outside of the immediate trading range in order to tilt the market-making engines of others and sometimes to slow down the whole system. Also the affair of flash orders that create a false sense of liquidity but cannot be reached by most participants due to their access latency is still being investigated and debated.

The modern Goulds, Drews, and Fisks shall also be found out either by people or by algorithms; it is just a matter of time and evolution. Every new organism has to test its boundaries to adapt and survive. In the current fashion of detox by regulation that started during the hangover from the credit bubble, one will probably see a formal response to high-frequency market abuse soon. The main difficulty will be, as with any other regulation, to ensure fairness without hampering efficiency.

The self-fulfilling nature of the technological uptrend of the marketplace with its associated growth of the automated systematic strategies lets us wonder whether it is a bubble ready to burst. As we commented at the beginning, any sector of the market (as well as of human activities more generally) is quite naturally prone to a boom-bust cycle, and this one should be no exception!

Currently this market arena presents high hurdles for entry and in consequence there is a fragmented technological landscape giving privileged access to only a few. There is still a large scope for competition that will invariably drive further increases in efficiency. This points to the fact that we are not yet at crunchpoint; the trend is still up, thus it is still rational to participate!

It is now time to move on to the heart of the matter and start exploring in Part One data representation, indicators, basic model types, and techniques to test them.