Chapter 1 — Introduction

This book is about developing trading systems using scientific techniques. It applies to any of the commonly traded financial issues—stocks, bonds, mutual funds, exchange traded funds, futures, currencies, FOREX, commodities. The systems are based on analysis of the price and volume of previous transactions made in open markets.

As the subtitle says, this book describes the importance of using the scientific method, and its derivatives related to engineering design and data science, in the process of developing trading systems and trading management systems.

The purpose of this book is to outline a few basic and relatively simple, but not necessarily simplistic, ideas that will assist readers in their system trading. The focus is developing and managing systems that have a good trade-off between reward and risk, and systems in which the trader will have confidence. And to draw attention to some common misconceptions related to trading system development and trading that can easily lead to systems that are likely to be difficult to trade confidently, riskier than expected, and less profitable than expected.

The process is quantitative, data driven, evidence based, and statistically sound. The resulting systems are rule based, monitored daily, and managed to hold risk to a level that is in keeping with the trader's risk tolerance.

This book is not intended to be encyclopedic. It lays out the foundations and their justifications. There is little computer code in this book. Implementation details, including fully disclosed examples, can be found in my other materials. See the Bibliography.

As every engineer will tell you, in order to design and develop a product, you must know how it will be used; in order to design and develop a process, you must know how it will be measured and managed.

Our product will be a profitable trading system. Our process will be designing and verifying the system, then monitoring its performance and determining the maximum safe position size. Our metrics will be account growth, normalized for risk.

This document is a chapter of the book Foundations of Trading
Published by Blue Owl Press, Inc.
www.FoundationsOfTrading.com

Goal

There is a goal common to all traders:

> To have confidence that the signals generated by the trading system precede trades that provide rewards adequate to compensate for the risk.
>
> The key word is confidence.
>
> The primary limitation is risk.

The remainder of this book explains how to best achieve that goal using rigorous and statistically sound quantitative technical analysis with well defined rules generated from analysis of the price data developed using principles of the scientific method and its related techniques.

Major Changes

Some major changes are taking place in trading system development and trading management. Each of these phrases identifies the "from" and the "to."

Broadly—Galileo to Hubble
Charts to Equations
Subjective to Objective
TSDP* to Machine learning
Indicators to Patterns
Stationary to Dynamic
Position size into Trading mgmt
Single backtest to Monte Carlo
Frequentist to Bayesian
Idea driven to Data driven
Profit oriented to Risk oriented
Deterministic to Probabilistic
Reaction to Prediction
Decision tree to Non-linear
p-value to Confusion matrix
Equity curve to Distribution
Impulse signals to State signals
Trade-by-trade accounting to marked-to-market daily

* TSDP: Traditional Trading System Development Platform

Decisions and Uncertainty

Most of the decisions we make in life are choices that involve weighing opportunity against risk. Most of the calculations are extremely complex and involve estimating costs and values of things not easily quantified—whom to choose as a partner, where to live, what employment to pursue. All are specific applications of making decisions under uncertain conditions. It seems that the more important the decision, the less opportunity we have to practice and the more important it is to be correct early in the process.

How we handle our finances is certainly an important area, and one where we don't get many practice runs. For traders, the goal is maximizing trading profits while minimizing the risk of bankruptcy. In the spectrum of life's activities, this is a problem that is relatively easy to quantify and analyze. The major aspects already have easily measured units of value—dollars. And, given a little understanding of probability and statistics, along with some computer data analysis, we can outline a plan.

Trading and Investing

By trading I mean buying and selling financial instruments with the intention of increasing wealth. While the terms are often used interchangeably, I distinguish between trading and investing.

An investment is a use of funds to acquire some asset, often a tangible asset such as real estate, that will not be sold during the investor's lifetime. The details of the transaction and directions for inheritance are recorded in the investor's will.

A trade is a use of funds to acquire some asset that will be sold at some point in the trader's lifetime. In the context of this book, trading is buying and selling financial instruments—stocks, bonds, mutual funds, exchange traded funds, futures contracts, currency contracts—with fairly short holding periods.

There is a third category—expense. When funds are spent, or the item purchased is consumed or becomes worthless, the transaction is an expense. There is no residual value.

Transactions sometimes change categories.

- Stock or real estate may be purchased with the intention of holding forever and passing it on to heirs, but sold for any number of reasons.
- A stock may be purchased in anticipation of a rise in its price and sale for a profit, only to have its price drop to a very low level, so the trader turns the trade into an investment by holding the stock indefinitely hoping for a recovery.
- An automobile may be purchased expecting it to become a classic and collector's item—later to become damaged and a total loss.

A person who buys a stock with the expectation of selling it at a profit is a trader, regardless of what he thinks at the time of the purchase. I believe that almost everyone who owns stocks, bonds, mutual funds—even real estate—initially intending to hold indefinitely is a trader; just one who has not yet planned the exit from his position.

With the exception of an investment made with the explicit intention of passing it on to heirs, all positions will eventually be closed out. Even investments made during prime working years which are intended to fund expenses in the far future, such as retirement, will eventually be sold.

Technical Analysis and Fundamental Analysis

There are two broad categories of financial data—technical data and fundamental data. Correspondingly, there are two broad techniques for determining which issues to select and when to buy and sell—technical analysis and fundamental analysis.

Technical data consists of the price and volume of transactions freely made, recorded, and published. *Technical analysis* examines price and volume data hoping to detect patterns that precede profitable trading opportunities.

Fundamental data consists of government economic reports, corporate management reports, and analyst estimates of future economic and business activity. *Fundamental analysis* examines economic series and corporate data hoping to identify companies whose share price will increase and economic conditions favorable to holding stocks.

Efficient Market Hypothesis

The efficient market hypothesis holds that asset prices fully and immediately reflect all information about the asset. The degree of efficiency is described as being one of three forms:

- Strong form. All information about a company—public and confidential—is reflected in the share price of the company. There is no data that provides a trader, investor, or even an insider an opportunity to profit from further analysis of the data.
- Semi-strong form. All public information is reflected in the share price. Insiders may have knowledge that gives them an advantage.
- Weak form. The current price of a stock incorporates the historical prices. No analysis of price history provides a trader an opportunity to profit.

There is a large body of discussion devoted to whether the financial markets are efficient or inefficient. Efficiency, in this context, refers to the question of whether variations in prices are just random noise, or whether they represent potentially profitable trading opportunities, and what categories of information are valuable to a trader.

Those favoring the point of view that markets are *strongly* efficient say that even insider information is already reflected in the price of the stock and is not valuable enough to create profitable trades. The capital asset pricing model and mean-variance portfolio construction follow from strongly efficient views.

Those favoring the point of view that markets are *semi-strongly* efficient say that public information, including both fundamental information and historical price information, is already priced into the market and cannot be used to make profitable trades.

Those favoring the point of view that markets are *weakly* efficient say that historical price and volume information is not valuable enough to create profitable trades.

Clearly, insiders do very well trading on information that only they have, so the markets are not strongly efficient. And there is enough controversy about the validity of any form of the efficient market hypothesis that all three forms may be disproved. In order for quantitative trading systems of the type this book discusses to be profitable, the market must have some inefficiency. If there are not persistent patterns and trends that we can identify and trade profitably, then we are all wasting our time.

Warren Buffett of Berkshire Hathaway is reported to have said, "I'd be a bum in the street with a tin cup if the markets were always efficient.[1]"

David Harding of Winton Capital Management asks "Efficient Market Theory—When Will it Die?[2]"

Fundamental Analysis

I am among those who feel that fundamental data provides no information from which a trader can profit. Fundamental analysis is based on the premise that a stock, bond, fund, commodity, or a market as a whole, has an underlying intrinsic value. By analysis of the fundamental characteristics, such as the assets, liabilities, income, supply, or demand, that value can be determined.

Fundamental data for a company includes earnings, sales, inventory turnovers, price to earnings ratios, price to sales ratios, dividend payout ratios, and any other information that might be reported on a balance sheet or income statement. Fundamental data extends to government and private research bureau reports, including gross domestic product, inflation, balance of payments, and any other data reported periodically. In general, fundamental data is only gathered, summarized, and reported—it does not represent trades.

Economists and security analysts who focus on fundamental information have developed mathematical models of the fair price for one share

1 Fortune, April 3, 1995

2 https://www.wintoncapital.com/getmedia/ccba88f6-a914-45d9-8b0d-fcf9db47644c/Efficient-Market-Theory-when-will-it-die.pdf

of the stock—it is the current book value of the stock plus the present value of all future dividends that will be paid to that one share. Any difference between the actual price of the stock and this fundamental value of the stock represents an opinion on the part of shareholders.

The fundamental analyst uses data that is reported by a company or agency to create subjective models. She may use charts to gauge overall price activity, but with few mathematically defined indicators.

While fundamental analysis may have value in its own right, there are several problems associated with incorporating it into trading models. See Appendix A for more detail.

Technical Analysis

Technical analysis, whether from the most subjective chart analysis, through quantitative analysis, to the most sophisticated machine learning and statistical analysis, is based precisely on the assumption that there is information in the publicly available price history. The underlying assumptions of technical analysis are:

- ✰ The markets we model are sufficiently inefficient for us to make a profit trading them.
- ✰ There is information—*patterns*—in the historical price series that can be used to identify profitable trading opportunities.
- ✰ Trading systems can be designed to recognize the patterns and give buy and sell signals.
- ✰ Patterns similar to those found in the historical data will continue to be found in future data.

Technical analysis began as chart analysis, and has developed a large body of subjective interpretation of chart artifacts such as flags, retracements, head-and-shoulders, and trend lines, to name just a few. We traders are all very good at selective vision—we see what we want to see. We can look at a chart and see examples of a big gain following, say, the breakout of a triangle pattern. Daniel Kahneman describes this tendency in his excellent book *Thinking, Fast and Slow*[3]. Thinking we have found a good trade entry technique, we can define those conditions in very precise terms and have the unbiased computer search all the data for instances of that pattern. The results do indeed show a profit for the pattern we saw so clearly. Often it also shows losses from many similar patterns that we either did not see or chose not to acknowledge; and it sometimes shows signals that appear, and then disappear as additional data points are added to the chart.

Quantitative analysis refines technical analysis by:

- Removing the judgment associated with ambiguous chart patterns.

---

3 https://www.amazon.com/Thinking-Fast-Slow-Daniel-Kahneman/dp/0374533555/

- Defining unambiguous, mathematically precise indicators.
- Requiring an indicator value for every data bar.
- Requiring that no indicator or signal may change in response to data that is received after it has been initially computed.
- Making extensive use of mathematical models, numerical methods, and computer simulations.
- Applying statistical validation techniques to the resulting trading models.

Many Trading Techniques

There has been trading for thousands of years. Beginning with the earliest markets—food, wine, spices, gems, cloth—people have bought and sold. Not only for their own use or as raw materials for the products they make, but also with the intent of making a profit through the trades.

Deciding what to trade, when to buy, and when to sell has been, and continues to be, the trader's craft. Even before personal computers, traders monitored the price and size of trades, created charts with plots based on prices, and looked for indications that prices were likely to change.

Over the years, there have been many refinements—standard contracts, regulated exchanges, central clearing, low commissions, powerful and inexpensive personal computers, trading-specific software, and publications of trade data. The barriers to becoming a trader are quite low and the reward for being successful are quite high. It is no wonder that trading is an attractive activity.

Why Traders Stop Trading

Assume a trader has a method—mechanical or discretionary—that she has been using successfully. Also assume that she understands both herself and the business of trading, and wants to continue trading. Why would she stop trading that particular system?

Here are a few possibilities:
1. The results are too good.
2. The results are not worth the effort.
3. The results are not worth the stress.
4. She has enough money.
5. There is a serious drawdown

1. Results are too good

She is afraid that this cannot possibly continue.

Her system—any system—works when the logic of the model and the data it analyzes are synchronized. There are many reasons why systems fail and should be taken offline, but a sequence of winning trades should be seen as a success.

She should continue trading it until one of the other reasons to stop happens.

2. Results are not worth the effort

There is not much gain, but not much loss either. Other things in life are more important. On balance, the time, energy, and resources would be more productively applied doing something else.

3. Results are not worth the stress

Performance is satisfactory, but at a high cost—worry and loss of sleep. Regardless of the position size indicated by the metrics of risk, the positions being taken are too large.

She should either reduce position size or have someone else execute the trades.

4. She has enough money

Not matter how good a system is, there is always a risk of serious loss.

When she has reached her goal, she should retire from trading.

5. There is a serious drawdown

The magnitude of the drawdown needed for it to be classified as serious is subjective. Among my colleagues and clients, those who manage other people's money typically want drawdown limited to single digits. Those trading their own money may be willing to suffer drawdowns of 15 or 20 percent.

But there is a level at which everyone stops trading the system—preferably while the account still has a positive balance.

My view is that experiencing a large drawdown is the primary reason people stop trading a system.

What causes a large drawdown and how should the trader react to it?

- The system is broken.
- There was an unexpected sequence of losing trades.
- The system is out of sync.
- The position size is too high.

As the account balance drops from an equity high into a drawdown, it is not possible to determine which is The reason.

All of the reasons are true to some extent. A system that is broken breaks because the logic and the data become unsynchronized, causing an unexpected sequence of losing trades and at a time when position size was too high for conditions.

The solution is two-fold.
1. Continually monitor system performance and system health.
2. Modify position size to reflect recent performance.

During the trading system development process, a baseline of system performance is established. The out-of-sample trades from the walk forward phase are a good source of this data. Personal risk tolerance and system risk, taken together, determine position size for that system performance. As system performance changes, position size must also change.

Position size varies in response to system health.

Do not continue to trade a system that has entered a serious drawdown expecting that it will recover. It may recover on its own; it may require readjustment; or it may be permanently broken and never work again.

Take it offline and either observe it until recent paper-trade results demonstrate that it is healthy again, or send it back to development.

The correct position size for a system that is broken is zero.

The Scientific Method

The scientific method, in use since the work of Kepler and Galileo in the 17th century, is a process to learn general principles about some system through analysis of data and experimentation.

As Figure 1.1 illustrates, it consists of a sequence of steps, including:
- Observation and measurement of data.
- Hypothesis statement and model specification.
- Prediction and validation.

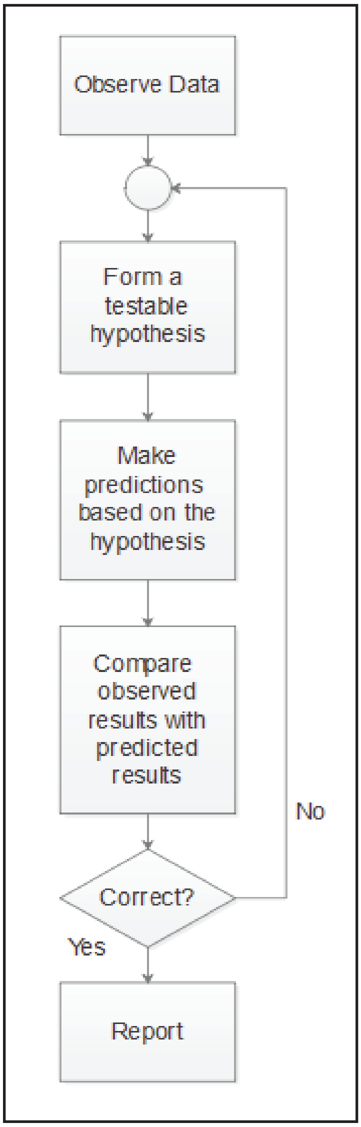

Figure 1.1 — Flowchart of the scientific method

A requirement of the scientific method is that the hypothesis is testable and is tested. A common error is to mistake the hypothesis to be an explanation of a phenomenon, exempt from performing tests—to believe that no test is needed.

An advantage of the scientific method is that it is unprejudiced. Faith, defined as belief that does not rest on logical proof or material evidence, is not required to determine whether to accept or reject the hypothesis. The results obtained using the scientific method are repeatable. It is not necessary to believe a published result; the reader can redo the experiment and replicate the result.

Importantly:

- The scientific method is rule-based and can be replicated. Given the model and the data, the results obtained are always the same without regard for who performs the tests or what his or her biases and subjective judgements might be.
- The model might consist of rules that are understandable, but that is not a requirement. Even if the rules are understandable, there is no requirement that they are in agreement with a previously defined body of cause-and-effect.

Software Engineering

Scientists study how nature works. Engineers create new things. Because their objectives are different, we might expect the processes used to be different. But the steps followed in the process of an engineering design are very similar to those of the scientific method.

- Define the problem.
- Specify requirements.
- Prototype alternative solutions.
- Test.
- Redesign as necessary.

Data Science

The term *data science* is in current favor to describe processes similar to the scientific method. One in particular, *CRISP-DM*, (Cross Industry Standard Process for Data Mining) describes steps of a data science lifecycle[4]:

- Business understanding.
- Data understanding.
- Data preparation.
- Modeling.
- Evaluation.
- Deployment.

The CRISP-DM process had been a standard supported by an industry consortium Special Interest Group, but that has dissolved. In 2015, IBM introduced a new methodology called *ASUM*, Analytics Solutions Unified Method for Data Mining, which refines and extends CRISP-DM. The IBM announcement explains that ASUM continues to sup-

4 Wikipedia, https://en.wikipedia.org/wiki/Cross_Industry_Standard_Process_for_Data_Mining

port the strong analytics of CRISP-DM while introducing additional tasks and activities in the deployment phase.

The diagram[5] in Figure 1.2 representing the CRISP-DM process is adequate for our purposes.

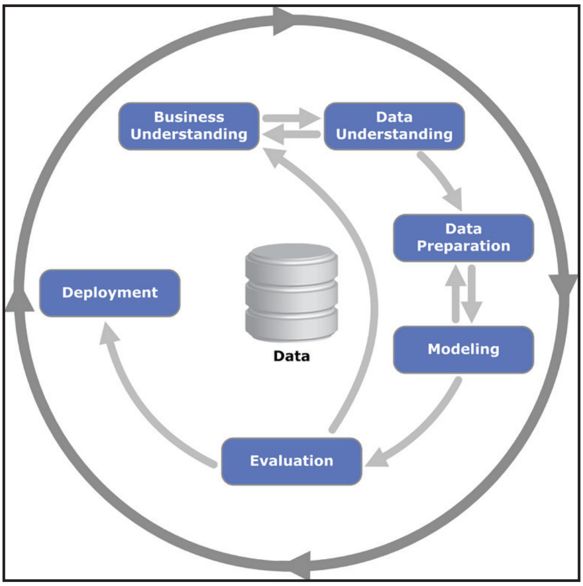

Figure 1.2 — The CRISP-DM process

The steps of data preparation, modeling, and evaluation are iterated until the results, based on evaluation of new and previously unused data, are satisfactory.

Trading System Development

Figure 1.3 shows the flowchart of the process of developing a trading system and its associated trading management system. The diagram clearly follows the scientific method, engineering process, and data science approach. It is the process discussed in this book.

- Observation and measurement of data. Analysis of historical data searching for patterns that precede profitable trades.
- Hypothesis statement and model specification. Define indicators, adjust parameters, propose rules that define trading signals, generate trades, analyze trades.
- Prediction and validation. Predict results for yet unseen data. Analyze accuracy of predictions.

5 Created by Kenneth Jensen, use per creative commons

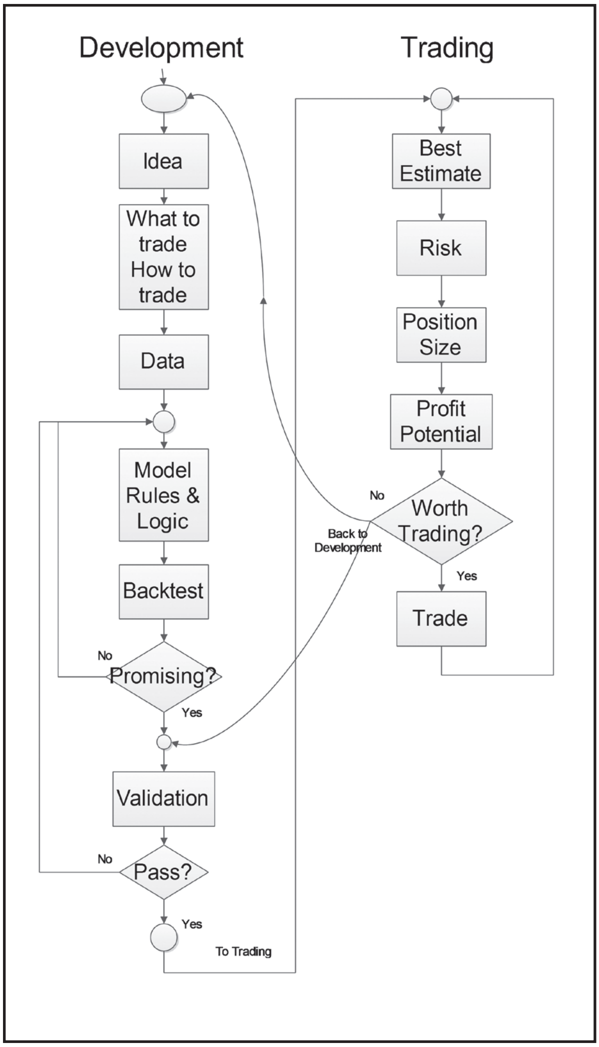

Figure 1.3 — Flowchart of trading system development and trading management

Each Side of the Flowchart is a System

Trading system development has long been thought of as a relatively simple process of applying some chart pattern or indicator to a lengthy series of historical price data, often including a search for the best rules and parameter values, and often including calculation of the position size to be used for each trade.

As both the thinking about that process and the tools available for use with that process have evolved, we can develop better trading sys-