I J C T A, 9(34) 2016, pp. 857-868

Quantitative Finance - Black Box Trading models

Rajan Lakshmi A.* and V. N. Sailaja**

ABSTRACT

Research is the heart of quant trading. It is in large part because of well-designed, rigorous, and tireless research programs that the best quants earn their laurels. This paper gives an overview of what research really means for black-box traders. It focuses mostly on research targeted at developing the alpha models of trading strategies. Research is also done with regard to risk models, transaction cost models, portfolio construction models, execution algorithms, and monitoring tools. Relevant research topics in these other areas will be mentioned as necessary, but the general principles from this section hold true throughout the black box.The purpose of research is to scrutinise a well-conceived investment strategy. A strategy is a long-term course of action designed to achieve an objective, usually success or victory. In most applied settings, strategies are chosen from a limitless number of alternatives.

Keywords: Alpha model, risk model, transaction cost model, portfolio construction models.

INTRODUCTION

A computer-based trading system for individual investors that uses a set of fixed, proprietary rules to generate buy and sell signals. Black box systems are named for the secrecy surrounding the methodology employed in the analysis. The "black box" portion of the system contains formulas and calculations that the user does not see nor need to know to use the system. Black box systems are often used to determine optimal trading practices. These systems generate many different types of data including buy and sell signals.The trading system has three modules—an alpha model, a risk model, and a transaction cost model—which feed into a portfolio construction model, which in turn interacts with the execution model. The alpha model is designed to predict the future of the instruments the quant wants to consider trading for the purpose of generating returns. Quantitative trading strategies¯known to many as "black boxes"¯have gained a reputation of being difficult to explain and even harder to understand. While there is a certain level of complexity to this approach, with the right guidance, one can successfully overcome potential obstacles and begin to excel in this arena.

DATA SOURCES

One can get data from many sources. Most direct, but also perhaps most challenging, is to get raw data from the primary sources. In other words, a quant would get price data for stocks traded on the New York Stock Exchange directly from the NYSE. This has the benefit of allowing the quant maximum control over the cleaning and storing of data, and it can also have significant benefits in terms of speed. However, there is also a massive cost to doing things this way. It would require building connectivity to every primary source, and if we are speaking about trading multiple types of instruments (e.g., stocks and futures) across multiple geographical markets and exchanges, the number of data sources can explode. With each, software must be built to translate the primary sources' unique formats into something usable by the quant's trading systems.

* Research Scholar, KL University, Guntur AP, *E-mail: rajyakavi@gmail.com*

** Assistant Professor, K L U B S, K L University, *E-mail: drvedalasailaja@kluniversity.in*

LITERATURE REVIEWED

The interplay between algorithmically driven and traditional trading strategies affects the returns of all investors. Brian Brown2 s new book provides a very clear overview of how these new strategies work and more importantly, how they influence liquidity, volatility, and prices in the global equity market.

— Andrew J. Morton, Co–creator of the Heath–Jarrow–Morton (HJM) Framework

Technology advances over the past decade have dramatically changed the dynamic world of stock market trading. Most analysts have failed to account for this "brave new world" in their "Monday morning quarterback" analysis of the recent worldwide financial systems collapse. Brian Brown has written the first book that clearly and colorfully describes the new technologically–driven way of doing business on the Street, and he does this with great precision and street knowledge. This new world played a central role in the Wall Street collapse, and, paradoxically, will help drive the next ascent.

— Thomas F. Coleman, Dean and Professor, Faculty of Mathematics, Director, Waterloo Research Institute in Insurance, Securities, and Quantitative Finance, University of Waterloo

Much has been made of the activities of "High Frequency Traders" during the Global Financial Crisis. In many cases they have been vilified, but often out of ignorance about the vital function that they perform in today2 s hyper–speed financial markets. Brian sets out to demystify High Frequency Trading and does so in an eminently readable fashion. This book will appeal to anyone, market professional or not, who wants to understand this often secretive group.

— E. John Fildes, Chief Operating Officer, Asia, Instinet Pacific

Schematic of the Black Box

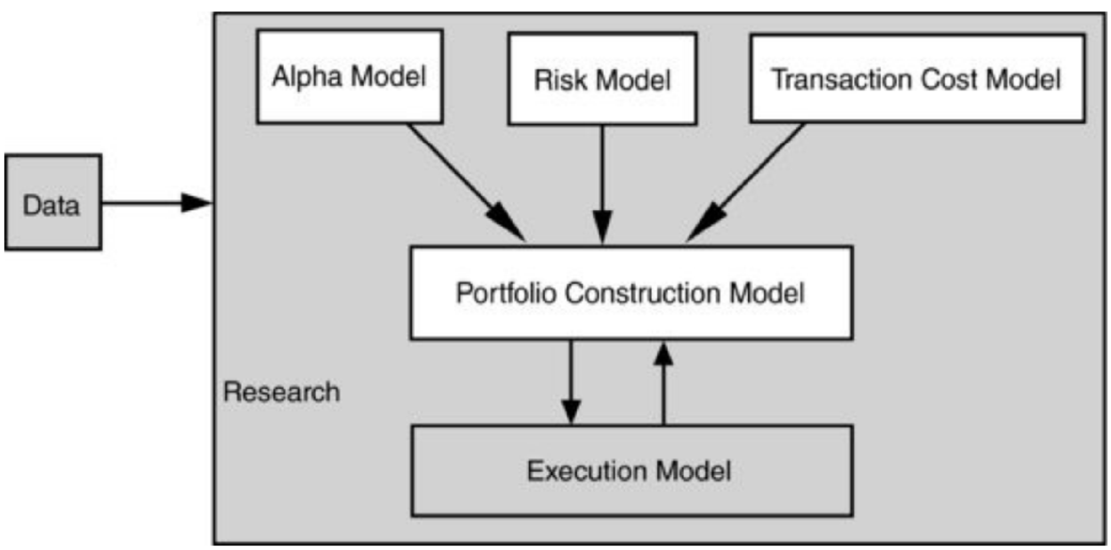

ALPHA MODEL

The black box can be understood by understanding the heart of the actual trading systems that quants use. This first piece of a quant trading system is its alpha model, which is the part of the model that is looking to make money and is where much of the research process is focused. Alpha, the spelled-out version of the Greek letter á, generally is used as a way to quantify the skill of an investor or the return she delivers independently of the moves in the broader market.The software that a quant builds and uses to conduct this timing systematically is known as an alpha model, though there are many synonyms for this term: forecast, factor, alpha, model, strategy, estimator, or predictor.

There are two types of Alpha Models

A. Theory Driven and

B. Data Driven

Theory Driven Alpha Models

Its starts with some economically feasible explanation of why the markets behave in a certain way and test these theories to see whether they can be used to predict the future with any success.

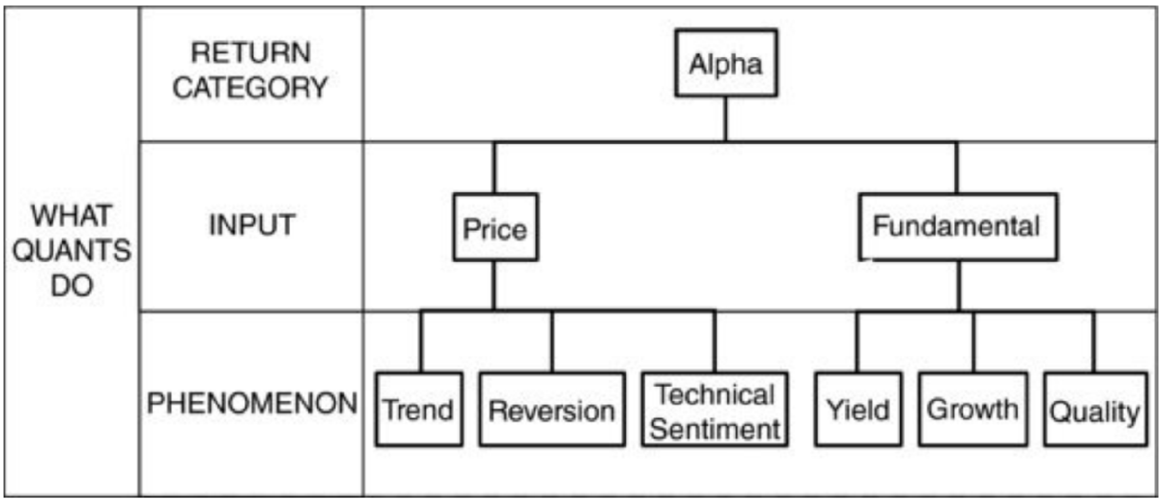

Data Driven Alpha Models

These strategies are far less widely practiced for a variety of reasons, one of which is that they are significantly more difficult to understand and the mathematics are far more complicated. Data mining, when done well, is based on the premise that the data tell you what is likely to happen next, based on some patterns that are recognisable using certain analytical techniques. When used as alpha models, the inputs are usually sourced from exchanges (mostly prices), and these strategies typically seek to identify patterns that have some explanatory power about the future.

Time Horizon

The next key component to understanding implementation of the alpha model is time horizon. Some quant models try to forecast literally microseconds into the future; others attempt to predict behaviour a year or more ahead. Most quant strategies have forecast horizons that fall in the range of a few days to several months.

Risk Models

Risk exposures generally will not produce profits over the long haul, but they can impact the returns of a strategy day to day. More important still, the quant is not attempting to forecast these exposures, usually because he cannot do so successfully. But the fact remains that one of the great strengths of quant trading is to be able to measure various exposures and to be intentional about the selection of such exposures.

There are several kinds of quantitative risk models that limit size, and they vary in three primary ways:

1. The manner in which size is limited.
2. How risk is measured. ?
3. What is having its size limited ?

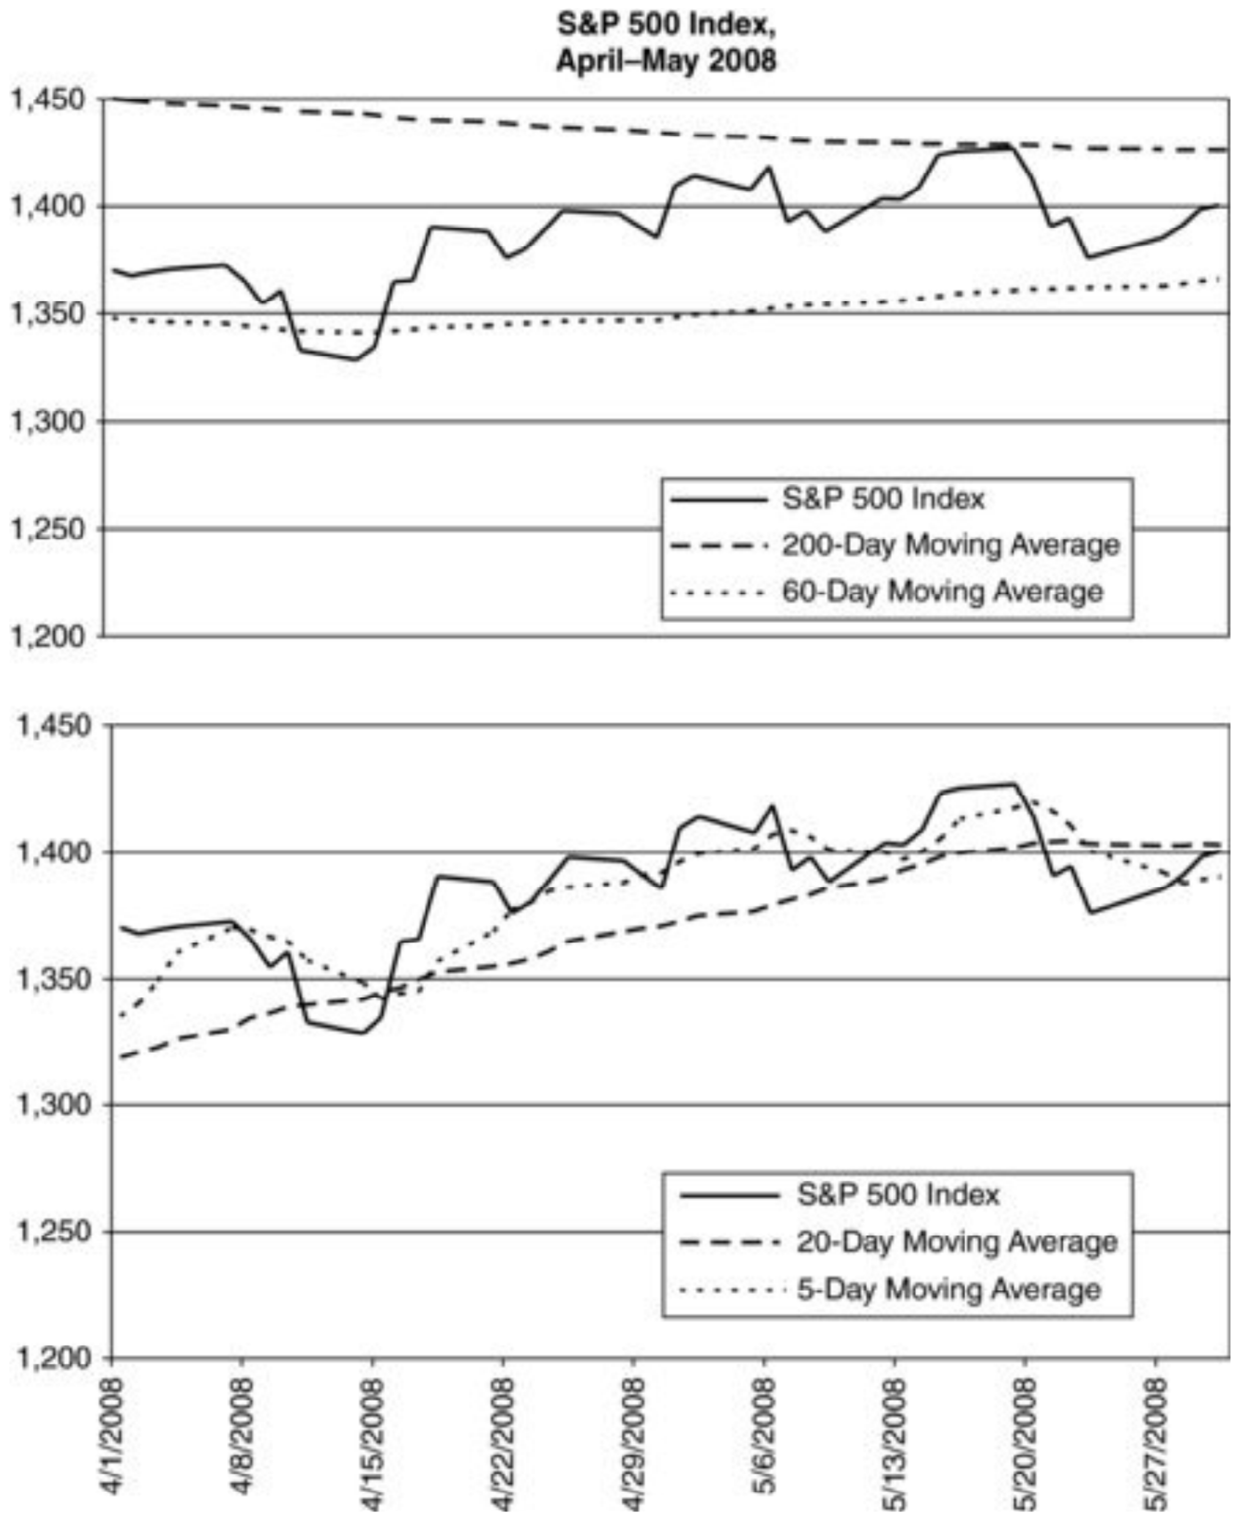

Measuring the Amount of Risk

There are two generally accepted ways of measuring the amount of risk in the marketplace.

1. The first is longitudinal and measures risk by computing the standard deviation of the returns of various instruments over time, which is a way of getting at the concept of uncertainty. In finance circles, this concept is usually referred to as volatility. The more volatility, the more risk is said to be present in the markets.

2. The second way to measure risk is to measure the level of similarity in the behaviour of the various instruments within a given investment universe. This is frequently calculated by taking the cross-sectional standard deviation of all the relevant instruments for a given period. The larger the standard deviation, the more varied the underlying instruments are behaving.

Theory-Driven Risk Models

Theory-driven risk modelling typically focuses on named or *systematic* risk factors. Just as in the case of theory-driven alpha models, systematic risks that are derived from theory are those for which the quant can make a reasonable, economic argument. Theory-driven risk modelling uses a set of predefined systematic risks, which enables the quant to measure and calibrate a given portfolio's exposures.

Empirical Risk Models

Empirical risk models are based on the same premise as theory-driven models, namely that systematic risks should be measured and mitigated. However, the empirical approach uses historical data to determine what these risks are and how exposed a given portfolio is to them. Using statistical techniques such as principal component analysis (PCA), a quant is able to use historical data to discern systematic risks that don't have names but that may well correspond to named risk factors.

Transaction Cost Models

Transaction cost models is that it costs money to trade, which means that one should not trade unless there is a very good reason to do so. This is not an overly draconian view of trading costs. Many highly successful quants estimate that their transaction costs eat away between 20 and 50 percent of their returns.

A transaction cost model is a way of quantifying the cost of making a trade of a given size so that this information can be used in conjunction with the alpha and risk models to determine the best portfolio to hold.

Defining Transaction Costs

It is useful to understand what the costs of trading actually are, since we are describing ways to model them. Transaction costs have three major components:

A. commissions and fees,

B. slippage, and

C. market impact

Commissions and Fees

Commissions and fees, the first kind of transaction costs, are paid to brokerages, exchanges, and regulators for the services they provide, namely, access to other market participants, improved security of transacting, and operational infrastructure. For many quants, brokerage commission costs are rather small on a per-trade basis. Quant traders typically do not utilise many of the services and personnel of the bank but instead use only the bank's infrastructure to go directly to the market. The incremental cost of a trade to a bank is therefore very small, and even very low commissions can be profitable. Given the volume of trading that quants do, they can be extremely profitable clients for the brokerages, despite the diminutive commissions they pay. Some quants utilise significantly less of the bank's infrastructure and therefore pay even lower commission rates than others who use more and pay higher rates.Commissions are not the only costs charged by brokerages and exchanges. Brokers charge fees (which are usually a component of the commissions) for services known as clearing and settlement. Clearing involves regulatory reporting and monitoring, tax handling, and handling failure, all of which are activities that must take place in advance of settlement. Settlement is the delivery of securities in exchange for payment in full, which is the final step in the life of a trading transaction and fulfils the obligations of both parties involved in the transaction. These services take effort and therefore cost money. And, given that many quants are doing tens of thousands of trades each day, there can be a significant amount of work involved.

Slippage

Commissions and fees certainly are not negligible. But neither are they the dominant part of transaction costs for most quants. They are also basically fixed, which makes them easy to model. If the all-in commissions and fees add up to, say, $0.001 per share, the quant must simply know that the trade in question is worth more in terms of alpha generation or risk reduction than this $0.001 per-share hurdle. On the other hand, slippage and market impact are considerably trickier to measure, model, and manage.

Strategies that tend to suffer most from slippage are those that pursue trend-following strategies, because they are seeking to buy and sell instruments that are already moving in the desired direction. Strategies that tend to suffer least from slippage, and for which slippage can sometimes be a *positive*, are those that are mean reverting in orientation, because these strategies are usually trying to buy and sell instruments that are moving against them when the order is placed. A quant trader's latency or speed to market has a large effect on the level of slippage his strategy will experience over time. This is because slippage is a function of the amount of time that passes between the order being decided and the order reaching the market for execution. The more latency in a trader's system or communications with the marketplace, the more time passes before her order gets to the market and the further the price of an instrument is likely to have moved away from the price when the decision was made. Worse still, the more accurate a forecast, particularly in the near term, the more damaging slippage will be.

Market Impact

Market impact, the third and final major component of transaction costs, is perhaps the most important for quants. The basic problem described by market impact is that, when a trader goes to buy an instrument, the price of the instrument tends to go up, partly as a result of the trader's order. If the trader sells, the price goes down as he attempts to complete his trade. At small order sizes, this price movement usually bounces between the current best bid and offer. However, for larger orders, the price move can be substantial, ranging in the extremes, even to several percentage points. Market impact, then, is a measurement of how much a given order moves the market by its demand for liquidity. Market impact is normally defined as the difference between the price at the time a market order enters the exchange and the price at which the trade is actually executed.

Electronic communication networks (ECNs) are examples of platforms for customers to trade directly with one another. The challenge for ECNs is to attract enough customer order flow to show abundant liquidity on their exchanges. ECNs also must provide robust technology so that their exchanges can continue to function without disruption. To attract providers of liquidity, most ECNs in equity markets have established methods to pay traders who provide liquidity and take payment from traders who demand liquidity. It might cost something like three-tenths of a penny per share for a trader who buys shares at the offer or sells shares at the bid, whereas those providing the bids and offers that are getting hit are earning closer to two-tenths of a penny. The ECN keeps the difference, around one-tenth of a penny per share, as its source of revenue. Some kinds of trading strategies (usually mean reversion strategies) actually call for a mostly passive execution approach in which this act of providing liquidity is modelled as a source of profit due to the rebate programs that ECNs put in place to attract liquidity providers.

Dark pools also allow customers to interact with one another. Dark pools are created by brokers or independent firms to allow their customers to trade directly with each other in an anonymous way. They arose in part because of concerns about the market impact associated with large orders. On a dark pool, there is no information provided about the limit order book, which contains all the liquidity being provided by market makers and other participants. Customers are simply posting their orders to the pool and if someone happens to want to do the opposite side of those orders, the orders get filled. As a result of this anonymous process of matching orders, the market is less likely to move as much as it would in a more public venue, where automated market making practitioners require compensation to take the other side of large orders. One fact that makes dark pool transactions somewhat unusual is that they are over-the-counter, off-exchange transactions in instruments that are exchange traded. Dark pools could not exist without the public markets, because the securities traded on dark pools are listed on public exchanges. Furthermore, the public markets provide the only transparent sense of price discovery, without which dark pool participants would have a significantly harder time determining what prices to bid and offer. Partly because of these issues, coupled with the fact that dark pools are available only to selected customers, controversy surrounds dark pools.

There are four basic types of transaction cost models

A. Flat Transaction Cost Model

B. Linear Transaction Cost Model

C. Piecewise-Linear Transaction Cost Model

D. Quadratic Transaction Cost Models

Flat Transaction cost model

The first kind of transaction cost model is a *flat* model, which means that the cost of trading is the same, regardless of the size of the order. This is extremely straightforward computationally, but it is rarely correct and is not widely used.

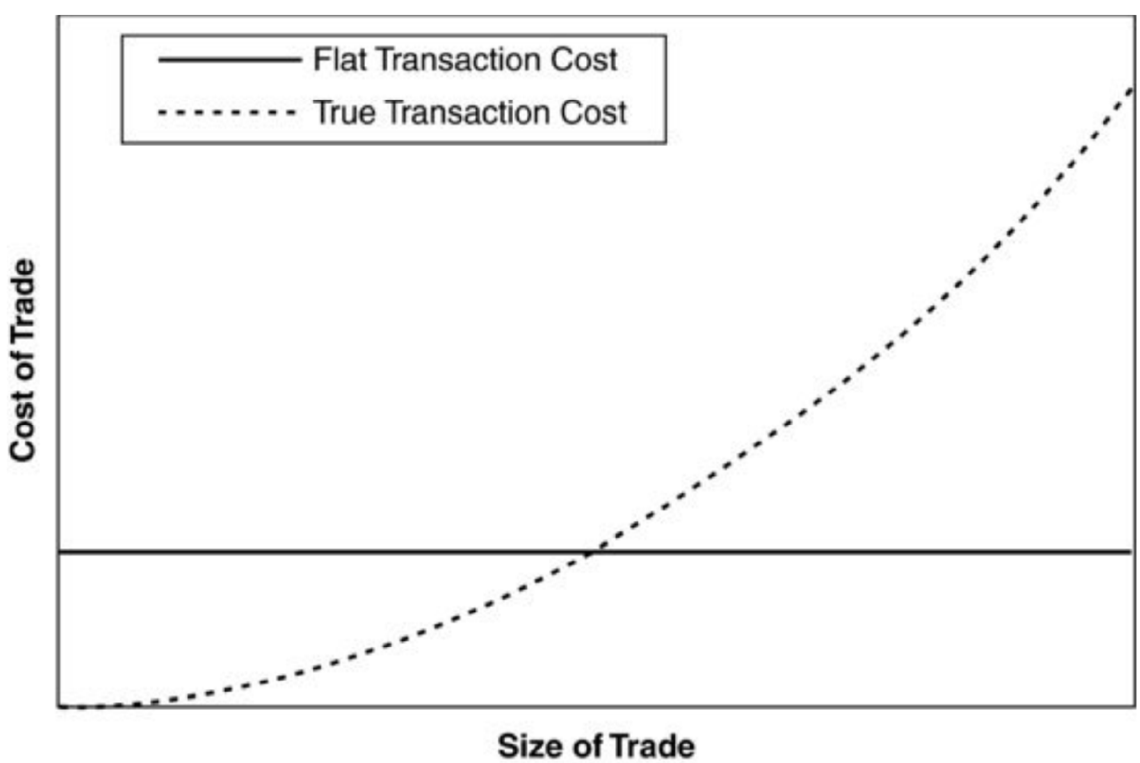

Linear Transaction cost model

The second kind of transaction cost model is linear, which means that the cost of a transaction gets larger with a constant slope as the size of the transaction grows larger, as shown in Exhibit This is a better fit relative to the true transaction cost, but it is still mostly useful as a shortcut to building a proper model.

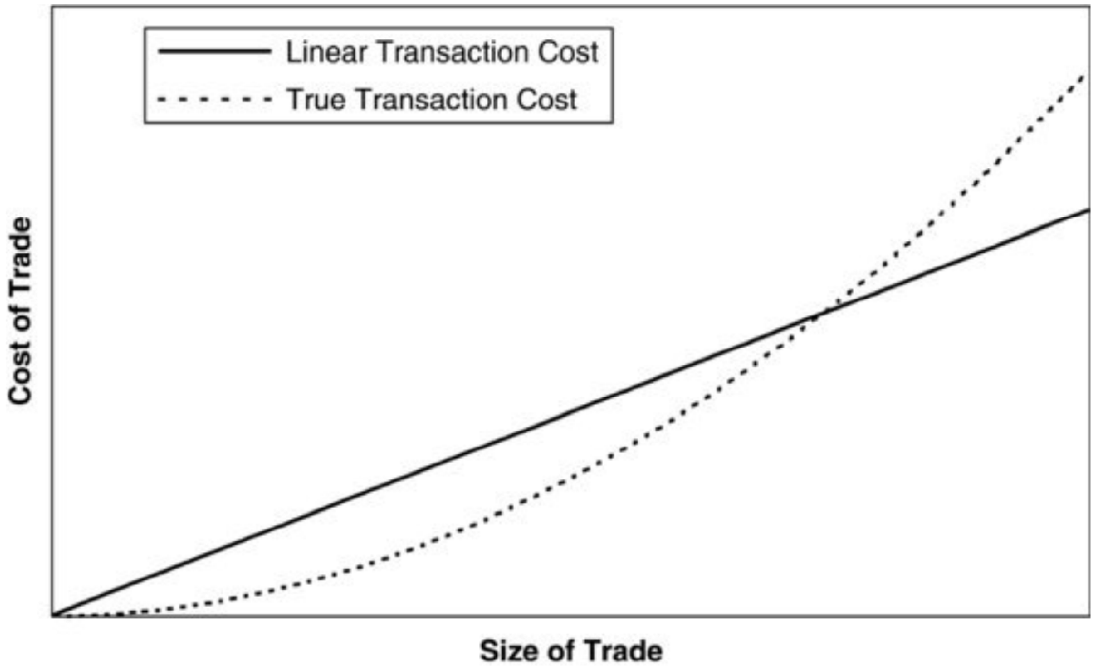

Piecewise-Linear Transaction Cost Models

Piecewise-linear transaction cost models are used to help with precision while using reasonably simple formulas to do so. The idea of a piecewise-linear transaction cost model is that, in certain ranges, a linear estimate is about right, but at some point, the curvature of the quadratic estimator causes a significant enough rise in the slope of the real transaction cost line that it is worthwhile to use a new line from that point on.

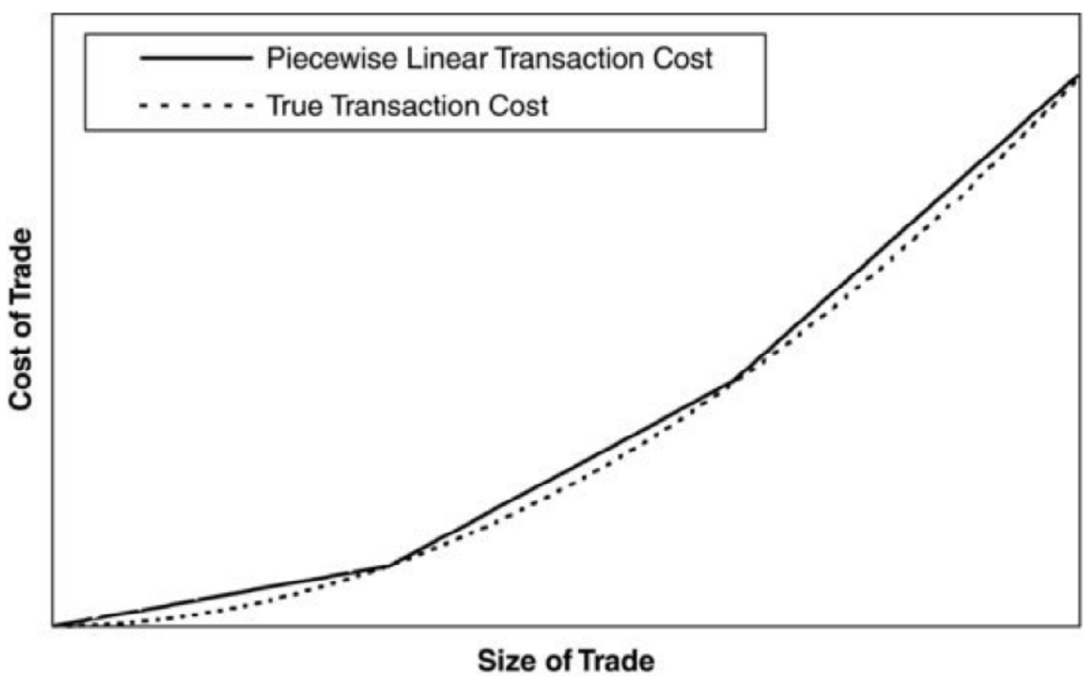

Quadratic Transaction Cost Models

Finally, quants can build *quadratic* models of transaction costs. These are computationally the most intensive because the function involved is not nearly as simple as what is used for a linear model, or even for a piecewise-linear model. It has multiple terms and exponents, and generally is a pain to build.

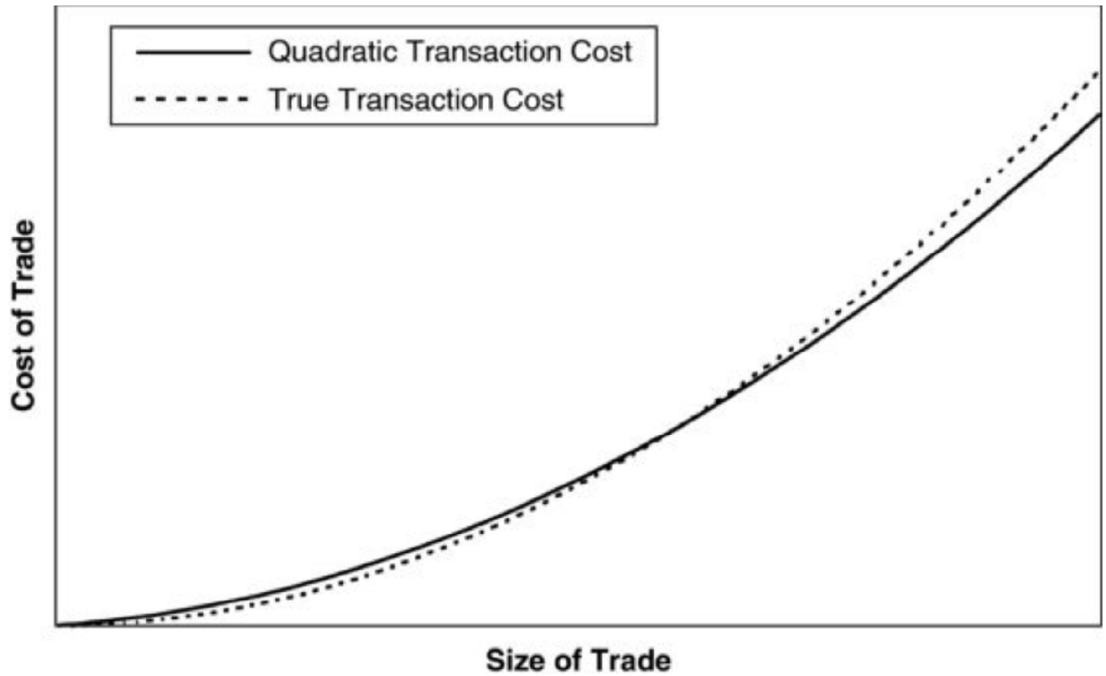

Portfolio Construction Model

The model acts like an arbitrator, hearing the arguments of the optimist (alpha model), the pessimist (risk model), and the cost-conscious accountant (transaction cost model), and then making a decision about how to proceed. The decision to allocate this or that amount to the various holdings in a portfolio is mostly based on a balancing of considerations of expected return, risk, and transaction costs. Too much emphasis on the opportunity can lead to ruin by ignoring risk. Too much emphasis on the risk can lead to underperformance by ignoring the opportunity. Too much emphasis on transaction costs can lead to paralysis because this will tend to cause the trader to hold positions indefinitely instead of taking on the cost of refreshing the portfolio.

RULE-BASED PORTFOLIO CONSTRUCTION MODELS

There are four common types of rule-based portfolio construction models:

1. equal position weighting,
2. equal risk weighting,
3. alpha-driven weighting, and
4. decision-tree weighting.

The first two are the simplest and have at their core a philosophy of equal weighting; they differ only in what specifically is being equally weighted. Alpha-driven portfolio construction models mainly rely on the alpha model for guidance on the correct position sizing and portfolio construction. Decision-tree approaches, which look at a defined set of rules in a particular order to determine position sizing, can be rather simple or amazingly complex.

Portfolio Optimisers

Portfolio optimisation is one of the most important topics in quantitative finance. This is one of the first areas in quant finance to receive the attention of serious academic work; in fact, the case could easily be made that the father of quantitative analysis is Harry Markowitz, who published a landmark paper entitled "Portfolio Selection."1 He invented a technique known as mean variance optimisation, which is still ubiquitous today, though much sophistication has been built around its core. In 1990, he shared a Nobel Prize with William Sharpe for both their contributions to the understanding of the quantitative analysis of portfolio construction.

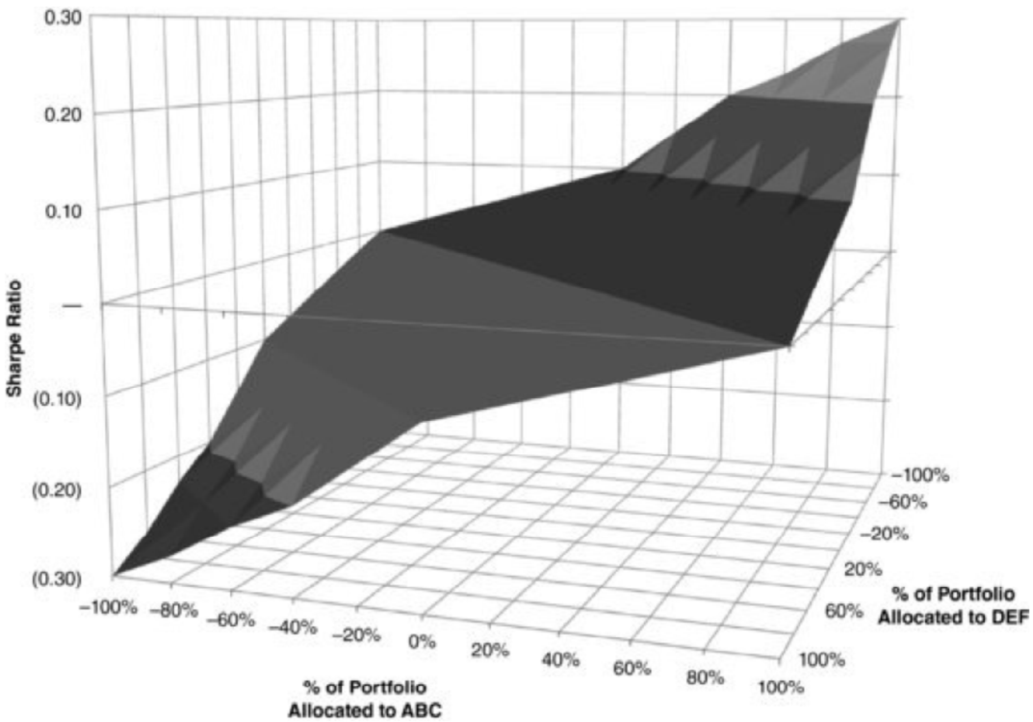

Visual Representation of Constraining the Search Space for an Optimisation

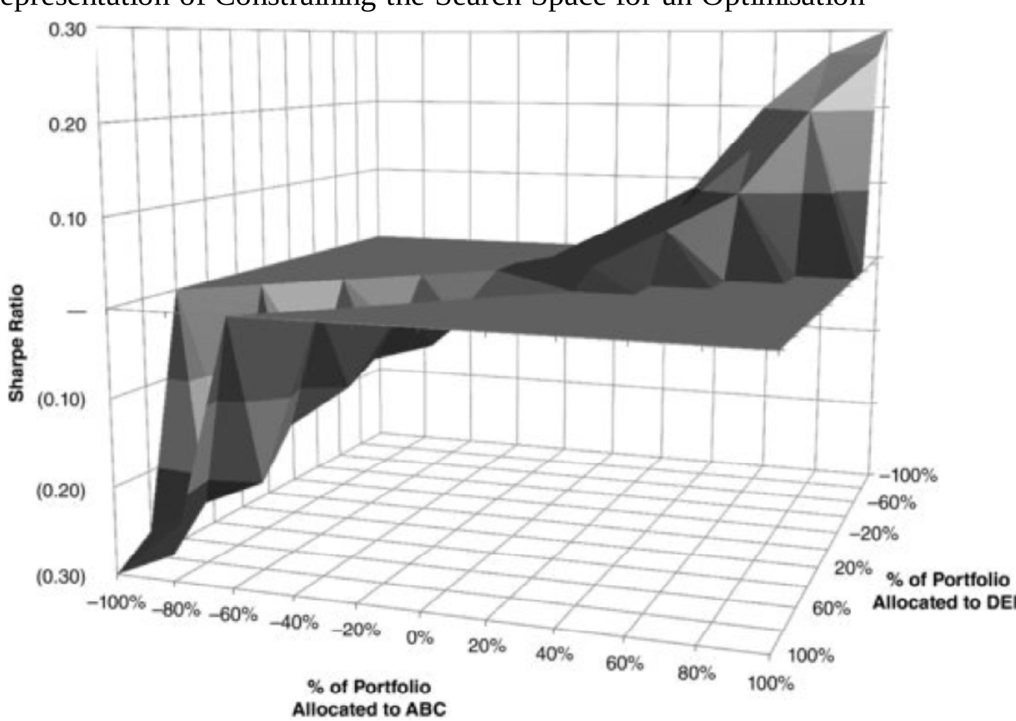

Execution model

These models are fed into a portfolio construction model, which determines a target portfolio. But having a target portfolio on a piece of paper or computer screen is considerably different from actually owning that portfolio. The final part of the black box itself is to implement the portfolio decisions made by the portfolio construction model, which is accomplished by executing the desired trades.

There are two basic ways to execute a trade:

A. electronically or

B. through a human intermediary (e.g., a broker).

Most quants elect to utilise the electronic method, because the number of transactions is frequently so large that it would be unreasonable and unnecessary to expect people to succeed at it. Electronic execution is accomplished through direct market access (DMA),which allows traders to utilise the infrastructure and exchange connectivity of their brokerage firms to trade directly on electronic markets such as ECNs.

Order Execution Algorithms

Order execution algorithms determine the way in which systematic execution of a portfolio is actually done. We can examine the kinds of decisions the algorithms must make in real time in much the same framework in which we'd think about how discretionary traders implement their orders. The kinds of considerations are the same in both cases, we find that quants differ here from their discretionary counterparts principally in the mechanics and not so much in the ideas. The principal goal of execution algorithms, and the function of most execution desks in general, is to minimise the cost of trading into and out of positions.

Aggressive versus Passive Order Execution

There are two general approaches to execution: aggressive and passive. Aggressive orders (most often in the form of market orders) are submitted to the marketplace and are generally unconditional. They can be filled in pieces or in full at whatever price prevails at the market at the time the order's turn to be executed arrives (within reasonable boundaries, and so long as there is a bid or offer resting in the order book to take the other side of the market order). In contrast, passive orders (a subset of all limit orders) allow the trader to control the worst price at which he is willing to transact, but the trader must accept that his order might not get executed at all or that only a part of it might be executed.

CONCLUSION

The role of transaction cost models is simply to advise the portfolio construction model how much it might cost to transact.Risk management is frequently misunderstood to be an exercise designed to reduce risk. It is really about the selection and sizing of exposures, to maximise returns for a given level of risk. After all, reducing risk almost always comes at the cost of reducing return. The two major families of portfolio construction models. Rule-based models take a heuristic approach, whereas portfolio optimisers utilise logic rooted in modern portfolio theory. Within each family are numerous techniques and, along with these, numerous challenges.Execution is where the rubber meets the road for a quant system and how the quant interacts with the rest of the marketplace. This continues to be a fruitful area of research, as it has been ever since markets have begun to become electronic.

REFERENCES

[1] *www.turtletrader.com/trader-seykota.html.*

[2] Larry Hite and Steven Feldman, "Game Theory Applications," Commodity Journal (May–June 1972).

[3] Ginger Szala, “Making a Mint: How a Scientist, Statistician and Businessman Mixed,” Futures, March 1, 1989.

[4] Gregory Zuckerman, “Renaissance Man: James Simons Does the Math on Fund,” Wall Street Journal, July 1, 2005.

[5] Eugene Fama and Kenneth French, “The Cross Section of Expected Stock Returns,” Journal of Finance 47 (June 1992): 427.

[6] Ray Ball, “Anomalies in Relationships Between Securities’ Yields and Yield-Surrogates,” Journal of Financial Economics 6, nos. 2–3 (1978): 103–126.

[7] Nate Silver, The Signal and the Noise: Why Most Predictions Fail but Some Don’t (New York: Penguin Press, 2012).

[8] This was calculated using Bayes’ Theorem.

[9] *https://normaldeviate.wordpress.com/2012/12/04/nate-silver-is-a-frequentist-review-of-the-signal-and-the-noise/*.

[10] Matthew Philips, “Where Has all the Stock Trading Gone?” May 10, 2012, *www.businessweek.com/articles/2012-05-10/where-has-all-the-stock-trading-gon*

[11] Risk concept was formalised in the Kelly criterion, in a paper by John L. Kelly, Jr., in the Bell System Technical Journal in 1956.

[12] Harry Markowitz, “Portfolio Selection,” Journal of Finance 7, no. 1 (March 1952): 77–91.2.

[13] Tim Bollerslev, “Generalized Autoregressive Conditional Heteroskedasticity,” Journal of Econometrics 31 (June 1986): 307–327.3.

[14] Fischer Black and Robert Litterman, “Global Portfolio Optimisation,” Financial Analysts Journal (September–October 1982): 28–43.4.

[15] Richard Grinold and Ronald Kahn, Active Portfolio Management: A Quantitative Approach for Producing Superior Returns and Controlling Risk (New York: McGraw-Hill, 1999).

[16] Richard Michaud, Efficient Asset Management: A Practical Guide to Stock Portfolio Optimisation and Asset Allocation (New York: Oxford University Press, 2001).

[17] Matthew Philips, “Where Has All the Stock Trading Gone?” May 10, 2012, *www.businessweek.com/articles/2012-05-10/where-has-all-the-stock-trading-gone#p1*.