CAPCO

Operational

Digital Finance: At the Cusp of Revolutionizing Portfolio Optimization and Risk Assessment Systems

Blu Putnam, Graham McDannel, Veenit Shah

Journal

THE CAPCO INSTITUTE JOURNAL OF FINANCIAL TRANSFORMATION

APEX 2016 AWARD WINNER

FINANCIAL TECHNOLOGY

#44
11.2016

EMPOWERING THE [FINANCIAL] WORLD

Pushing the pace of Financial Technology, together we'll help our clients solve technology challenges for their business – whether it's capital markets in Mumbai or community banking in Macon.

We leverage knowledge and insights from our clients around the world:

20,000 clients in towns everywhere are becoming more efficient, modern and scalable.

27 billion transactions processed help solve clients' challenges — big and small.

$9 trillion moved across the globe in a single year empowers our clients' communities to build storefronts, homes and careers.

55,000 hearts and minds have joined forces to bring you greater capabilities in even the smallest places.

Empowering the Financial World

FISGLOBAL.COM

FIS

Journal

The Capco Institute Journal of Financial Transformation

Recipient of the Apex Award for Publication Excellence

WHAT ARE THE DRIVERS AND DISRUPTIONS THAT DETERMINE INNOVATION AND PROSPERITY?

CAN EVERY PROBLEM BE SOLVED WITH A QUESTION? YES, BUT NOT EVERY QUESTION HAS A SINGLE ANSWER.

The Munk School's Master of Global Affairs program is developing a new class of innovators and problem solvers tackling the world's most pressing challenges.

> Tailor-made, inter-disciplinary curriculum delivering the best of both an academic and a professional degree.

> Access to world-leading research in innovation, economic policy and global affairs.

> International internships with top-tier institutions, agencies and companies that ensure students gain essential global experience.

COME EXPLORE WITH US

BE A MASTER OF GLOBAL AFFAIRS

MUNKSCHOOL.UTORONTO.CA
MGA@UTORONTO.CA

MUNK SCHOOL OF GLOBAL AFFAIRS

UNIVERSITY OF TORONTO

Financial Technology

Operational

Transformational

Digital Finance: At the Cusp of Revolutionizing Portfolio Optimization and Risk Assessment Systems[1]

Blu Putnam – Chief Economist, CME Group
Graham McDannel – Director, Data Science, CME Group
Veenit Shah – Data Scientist, CME Group

Abstract
Advances in quantum computing and machine learning are likely to change the face of quantitative portfolio construction and risk management as we know it today, and the focal point will be optimization processes. While financial optimization theory is highly sophisticated and complex, the current state of practice leaves much to be desired and may best be described as a patchwork quilt held together by band-aids and duct tape. On the horizon, however, are potential improvements in the analytical techniques underpinning how optimization methods are used, including the promise of exhaustive searches using quantum computers and advances in pattern recognition available through structured machine learning. To understand the importance and promise of the new developments in technology for financial optimization, it is imperative to appreciate the state of current practice. Critical challenges exist in the internal consistency of volatility and correlation estimates given the mixed methods used in many quantitative practices. With the heightened occurrence of event risk coming from politics, policy, and disruptive innovation, common assumptions concerning the stability of volatility regimes and correlation estimates are in question. Moreover, event risk can create short periods when bimodal expected return distributions dominate, often resulting in underestimation of the potential for pricing gaps and volatility regime shifts. Future progress with exhaustive search optimization using quantum computers and structured machine learning offers the possibility of a much deeper assessment of the probabilities surrounding event risk, improved analysis of the potential presence of bimodal and other non-normal return distributions, and the construction of more robust portfolios to handle the extreme (or fat-tailed) risks that seem to be happening more and more often than traditional approaches tend to predict.

1 Disclaimer: All examples in this report are hypothetical interpretations of situations and are used for explanation purposes only. The views in this report reflect solely those of the authors and not necessarily those of CME Group or its affiliated institutions. This article and the information herein should not be considered investment advice or the results of actual market experience.

Acknowledgement: We would like to offer our thanks and express our gratitude to Professor D. Sykes Wilford, Dr. Jose M. Quintana, and the late Professor Arnold Zellner for their practical applications of and contributions to using financial optimization in the real world despite the many challenges.

INTRODUCTION

Challenges to optimization abound in the world of portfolio construction and financial risk assessment. While financial optimization theory is highly sophisticated, with detailed theoretical attention paid to model construction and critical assumptions, the current state of practice leaves much to be desired, and may best be described as a patchwork quilt held together by band-aids or the ubiquitous duct tape. On the horizon, however, are some potential improvements in the analytical techniques underpinning how optimization methods are used in both portfolio construction and financial risk management. From the promise of exhaustive searches using quantum computers to the advances in pattern recognition available through structured machine learning, financial optimization methods are about to get a major makeover. Change may be coming, and it's about time!

To understand the importance and promise of the new developments in technology for financial optimization, however, it is imperative to appreciate the state of current practice faced by portfolio managers and risk officers. Critical challenges exist in the internal consistency of volatility and correlation estimates given the mixed methods used in many quantitative practices. With the heightened occurrence of event risk coming from politics, policy, and disruptive innovation, common assumptions concerning the stability of volatility regimes and correlation estimates are in question. Moreover, event risk can create short periods when bimodal expected return distributions dominate, often resulting in underestimation of the potential for pricing gaps and volatility regime shifts. Future progress with exhaustive search optimization using quantum computers and structured machine learning offers the possibility of a much deeper assessment of the probabilities surrounding event risk, improved analysis of the potential presence of bimodal and other non-normal return distributions, and the construction of more robust portfolios to handle the extreme (or fat-tailed) risks that seem to be happening more and more often than traditional approaches tend to predict.

Our research is divided into three sections. First, we go back to the father of Modern Portfolio Theory (MPT), Professor Harry Markowitz, and provide some perspective on his contributions. Second, we take a closer look at a few of the all too common practical approaches to financial optimization that fly in the face of critical assumptions embedded in the Markowitz approach. In our analysis of the common challenges to financial optimization that often lead to vast underestimations of risk and the construction of highly sub-optimal portfolios, we draw heavily from examples and illustrations taken from the U.K.'s June 2016 referendum to leave the European Union or "Brexit." Lastly, we come back to our key themes of how two major technical advances – quantum computing and machine learning – are likely to change financial optimization practices for the better.

HARRY MARKOWITZ AND THE ASSUMPTIONS UNDERLYING MEAN-VARIANCE OPTIMIZATION

The pioneer of modern financial optimization for portfolio construction and risk assessment is without a doubt Professor Harry Markowitz, winner of Nobel Prize in Economics in 1990. What is amazing is that over 65 years after the Markowitz mean-variance optimization came into the financial world back in the early 1950s, most practically applied financial optimization problems are addressed with the creative use of band-aids and duct tape (including some especially sophisticated mathematical methods) to handle known challenges that were embedded in the key assumptions chosen by Professor Markowitz in his doctoral dissertation at the University of Chicago to make the optimization problem tractable and available for real world use.

While there is a large and highly sophisticated body of literature involving the use of mean-variance optimization in finance, we will spare the reader both the mathematics and a recitation of the academic literature in favor of an intuitive review of some of the key challenges that scholars and practitioners have spent decades addressing. Our perspective is that an appreciation of the challenges of working with optimization methods in the real world effectively makes the case as to why a revolution in optimization methods finally is on the horizon.

The brilliance of Professor Markowitz's seminal work [Markowitz (1952)] in the 1950s was to recognize the role played by risk assessment in valuing stock and analyzing portfolios, since investors were effectively constructing portfolios with considerable uncertainty about the future. Indeed, MPT effectively embraced the approach set forth by Professor Markowitz, as a key element in security analysis.

As Professor D. Sykes Wilford noted in his insightful review of the contribution of Professor Markowitz to MPT [Wilford (2012)]: "In fact, MPT is ubiquitous to all financial theory and practice. By the same token, often the implementations of MPT break many of the basic assumptions behind MPT (and Markowitz) thereby making the conclusions derived from these actions extremely misleading, and in many cases completely incorrect."

Professor Wilford's contribution was to underscore the need to take a challenging look at how practical applications of financial optimization techniques handle the sometimes heroic assumptions embedded in the basic theory. This will be our approach here as well, and in so doing, we hope to set the stage for an appreciation of how quantum computing and machine learning are going to change the practice of portfolio construction and risk assessment – taking the real world closer to the theoretical world of Professor Markowitz.

THEORY TO PRACTICE WITH FINANCIAL OPTIMIZATION TECHNIQUES

While less appreciated, one of the more important research philosophies of Professor Markowitz was his focus on practical, applicable versions of portfolio optimization. There was in the 1950s and 1960s, a controversy in academic circles over whether economics should be seeking precise and general solutions or whether good approximations were acceptable. In his Nobel Lecture in 1990, Professor Markowitz commented on his approach and this debate [Markowitz (1991)]: "We seek a set of rules which investors can follow in fact - at least investors with sufficient computational resources. Thus, **we prefer an approximate method which is computationally feasible** to a precise one which cannot be computed. I believe that this is the point at which Kenneth Arrow's work on the economics of uncertainty diverges from mine. He sought a precise and general solution. I sought as good an approximation as could be implemented. I believe that both lines of inquiry are valuable" (bold added).

The practical approach of Professor Markowitz is where we start in our intuitive analysis of the challenges of portfolio optimization. We will focus on just a few critical assumptions commonly used in the current state of practice as we set up the case for the advances that will follow from quantum computing and machine learning. The critical assumptions we will review here include: (1) use of historical data to compute estimates for expected volatility and correlations while using a forward-looking method of creating expected returns; (2) use of the standard deviation as the common measurement for volatility; and (3) instability of the correlation matrix and existence of non-normal expected return distributions. All of these challenges are exposed in rather dramatic fashion with the presence of event risk. These intuitive discussions then lead us to illustrate our analysis with examples taken from the study of the "Brexit" referendum in June 2016.

Dangers and challenges of relying on history

To implement a Markowitz mean-variance optimization system, one needs expected values – that is, expected returns, expected volatilities, and expected correlations – that are used to describe aspects of the subjective probability distribution representing the risks faced by investors. When it comes to expected returns, there is no shortage of forward-looking quantitative and qualitative approaches. When turning to the expected volatilities and correlations, however, history is often used as a guide. There is a rarely used yet profound comment by Professor Markowitz on using history as a guide that bears remembering [Markowitz (1991)]: "The calculations . . . are the same as historical returns. **It is not that we recommend this as a way of forming beliefs**; rather, we use this as an example of distributions of returns which occur in fact" (bold added).

Using history as a guide for expected volatilities and correlations absolves the risk manager of any forecasting duties, yet subjects the owners of the underlying portfolio to very large error risk. There are good empirical reasons why many financial regulators require the disclaimer that "past performance is not necessarily a guide to future performance." History is always informative, however, every episode is different, so history is simply not always a good guide for developing expectations. There are serious questions about what period of history to use, how far back to look, to what degree is it appropriate to give older observations less weight and recent observations more weight. These are all quantitative questions on the surface that require subjective analysis, and they are beyond the scope of this research. We chose to place the focus on another challenge that is less well appreciated and yet potentially very dangerous. That is, the optimization problems get worse and the likelihood of risk underestimation gets much larger when the use of a forward-looking expected return method is attached to using history for volatility and correlation estimations.

A common refrain in the computer world is "GIGO" or "garbage in, garbage out." With optimization, the so-called garbage coming into the method bounces around the system in a highly networked manner determined by the expected correlation matrix, and one is quite likely to observe "garbage in, and a landfill of waste coming out the other end" – in effect, mean-variance optimization takes GIGO to an exponentially higher power. The problem is the inconsistencies involving three types of inputs – expected returns, expected volatilities, and expected correlations.

For example, if one has an aggressive expected return assumption for a given security, coupled to a historical set of data that do not reflect very much volatility, then this is asking for trouble in the mean-variance optimization space. The challenge arises from an interesting attribute of mean-variance computer systems – they actually believe what one tells them about expectations. Hence, if one provides an aggressive expected return with an expectation of little volatility, the mean-variance optimizer is going to produce a very large recommended exposure for the security. And then, the portfolio manager or risk officer will look at the output of the mean-variance optimization, remark that the output fails the real-world smell test, and either discount the method or add a set of constraints designed to create a more reasonable looking output.

This latter idea of adding constraints to optimization systems to achieve reasonable looking results is a very bad approach. Effectively, the unreasonable output has been caused by the inconsistency in the expected return and expected volatilities input into the optimizer. Rather than fix the inputs by adjusting expectations to make them more internally consistent, the common solution is to add constraints until the portfolio output passes the real world smell test.

This is like diagnosing the patient as a crazy man, and then resorting to putting the patient in a straitjacket to get the desired behavior. The much better approach, in psychoanalysis and in optimization, is to address the source of the problems directly.

One approach is to use the implied volatility in options pricing. However, efficient and useful options markets may well not exist, and some options-pricing models have built-in assumptions related to stable or flat future returns. Another, simpler band-aid is to incorporate information from the return expectations into the expected volatilities. That is, start with a measure of expected volatility, and then augment the volatility expectation based on the degree of aggressiveness of the expected return. With this approach, the mean-variance optimizer will see the aggressive return forecast, yet it will be coupled to a much larger expected volatility, so the exposure that is recommended in the optimized output will be much smaller and make more sense to the portfolio manager and risk officer.

Take the case of the U.K.'s June 2016 referendum on remaining in the European Union (E.U.) or leaving, known as "Brexit" (Figure 1). Prior to the vote on 23 June, the U.S. dollar (USD) was trading at around 1.42 against the British pound (GBP). If one thought the U.K. was going to vote to "leave," a typical forecast for the USD per GBP was 1.32 or lower. And by contrast, the "remain" camp expected a relief rally and a rise in the pound toward 1.52 (USD per GBP) or higher. The historical volatility in the three weeks before the vote was only an annualized 9.8% (standard deviation), even though market participants were looking for a one-day 7% or so move in one direction or the other depending on the outcome of the vote (i.e., a 5+ standard deviation event, one in a million event). As this case illustrates, and as the aggressiveness of the expected moves in the pound given the outcome of the vote suggested, a risk system or a

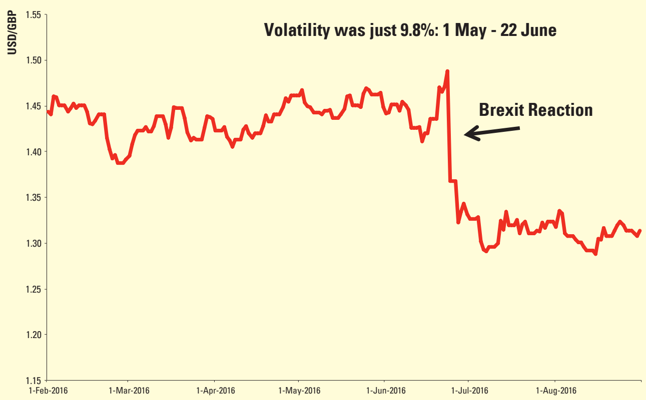

Source: Bloomberg Professional (GBP)

Figure 1 – The impact of Brexit on USD:GBP exchange rate

portfolio construction system needed to augment the recent historical volatility to capture the risks appropriately.

Standard deviation may underestimate volatility and potential skewness

The previous intuition, augmenting expected volatilities with information from the expected returns, raises another challenge. Is the standard deviation the appropriate proxy for the risk of the security returns in the first place? Again, and interestingly, the use of the standard deviation was chosen by Professor Markowitz back in the 1950s to represent risk because of its practical attributes. The standard deviation was straightforward to calculate from historical data and the standard deviation fit neatly into the mathematics of mean-variance optimization. There were other important side-effects of this choice. The standard deviation easily leads to embedding into the closed-form mean-variance optimization method the assumption of a normal or log-normal distribution of expected returns. Thus, we focus on at least two challenges here, (1) the standard deviation as often calculated from historical data may underestimate future volatility, and (2) the probability distribution of returns may well have considerable skewness (that is, fat-tailed event or "black swan" potential).

There are a couple of duct tape solutions available. First, the risk officer can embrace the need to take a forward-looking view of potential risks and incorporate them into the quantitative inputs for expected volatility. That is, when the future looks especially risky, despite the current calm state of markets, risk managers may choose to qualitatively augment their estimates of future volatility. We highly recommend this approach, as risk officers should not be able to hide behind historical calculations when such approaches are well known to underestimate risk and to understate the probability and frequency of highly skewed market events.

Second, one can look at alternative approaches for volatility measurement, such as looking at intra-period swings in prices. For example, if one is willing to assume a normal distribution of returns, then there is a deterministic mathematical relationship between the intra-period high/low price spread and the period-to-period standard deviation [Garman and Klass (1980); Parkinson (1980)]. If these two measures start to deviate in a meaningful way, then a market indicator can be constructed which incorporates the information from intra-period trading activity that may point to market participants worrying about more future volatility potential than the standard deviation suggests.

Again, by illustration, "Brexit" provides an interesting case study. In the weeks and months leading up to the "Brexit" referendum, as already noted, volatility, as measured by the standard deviation of daily percent changes in the USD:GBP exchange rate, suggested only modest risks more typical of "business as usual" activity.

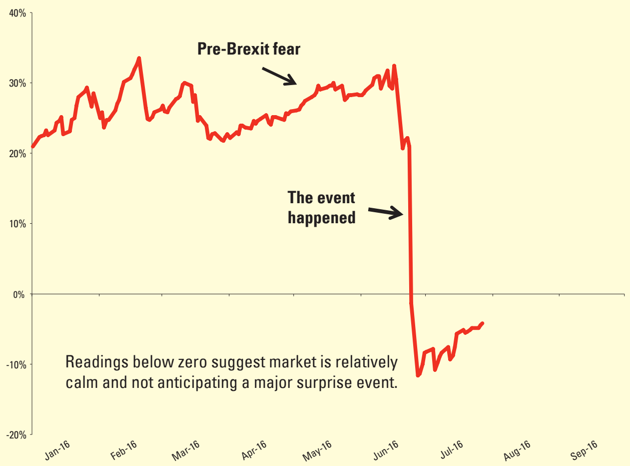

Values over 20% indicate increased worries about the possibility of future fat-tailed events.
Note: worries are often caused by a recent past event that was not expected
Source: CME Economics calculations

Figure 2 – British pound - market worry indicator (considers intra-day dynamics)

By contrast, in the pre-vote period, the intra-day price swings, as measured by the daily high and lows recorded in the nearby British pound futures contract price as traded on CME Group's Globex® electronic platform, suggested much higher risk. And, when the adjusted intra-day high-low price spread[2] is well-above the volatility estimate given by the standard deviation of closing price changes, then one has an indication that market participants are worried about a skewed or fat-tailed event occurring.

Interestingly, once the vote occurred and the outcome was known, the difference in volatility measures from these two techniques disappeared (Figure 2). Essentially, market activity reflected the fact the event had occurred and that another similar event was not expected. That is, the storm was a big one, but once it had passed by, the "worry" indicator slipped into neutral.

Instability of correlations and possibility of non-normal return distributions

Market participants quite often have to deal with the prospects of event risk. For example, corporation A makes a bid to acquire corporation B. However, the bid, even after being accepted by corporation B, needs regulatory approval, which may well be quite controversial. The event of the regulatory decision may be binary and result in the termination or consummation of the announced deal. Before the regulatory decision is announced, the stock prices of corporations A and B will reflect the probabilities of the deal terminating or consummating, meaning that the market price of the stock before the deal will not fully reflect the announced deal price if the probability of termination is greater than zero. After the regulatory decision, the stock price moves instantly to reflect whether the deal is going through or ending. Political event risk can look much the same, as it did with the binary "Brexit" vote. What we are describing here is the likelihood that event risk creates the possibility of bimodal return probability distributions [Putnam (2012)]. A distribution with two modes, where one mode is usually lower and far away from the higher mode, is a strikingly different subjective probability distribution than the normal distribution which is embedded in many risk assessment and portfolio construction systems.

During the pre-event stage, market prices of securities likely to be impacted by the event will move when expected probabilities of the binary outcomes shift. This means that the typical drivers of market prices, and thus observed correlations, may be highly distorted by the very different drivers of the shifts in subjective probabilities related to the event in question. That is, in more typical times, earnings expectations might drive the prices of stocks A and B. Once the acquisition is announced, the earnings matter much less, and the ebb and flow of news and views about the regulatory process that will approve or deny the acquisition take precedent.

As can be appreciated, the apparent increasing frequency of event risk, especially related to political events and policy decisions, is complicating the challenges of portfolio construction and risk assessment. A common practical solution, and one we endorse, is stress-testing with various scenarios reflecting the nature of the event risk about which one is worried. Critically though, the scenarios should be assigned subjective probabilities [Karagiannidis and Wilford (2015)]. It is pathetically easy to ask 20 questions or develop some interesting scenarios, but stress-testing has no meaning or useful application if subjective probabilities are not attached to the scenarios. Again, we see that the risk officer has to be forward-looking and probabilistic.

In addition, some market participants may be drawn to adopt options strategies to manage risk related to upcoming events. Options are favored in this regard because they embed a view of volatility in their price. We are strong supporters of options as a tool to manage event risk. However, we note that some additional sophistication may be required when event risk is present. Options behave differently when confronted with event risk than one might suspect if using an options pricing model derived from the basic Black-Scholes approach. We mention this because it highlights one of our key themes – namely, watch out for embedded assumptions. The Black-Scholes options

2 Adjusted for the difference in volatility measurement between standard deviation and high-low swings.

pricing method [Black and Scholes (1973); Merton (1973)] in its original and basic form makes a number of heroic assumptions designed to simplify the mathematics and allow one to use an options-replicating approach to value the option.

When event risk is present, two critical assumptions are likely to be violated and both have profound implications for the price of the option and the implied volatility expectation embedded in the option price. Event risk raises the prospect of both an instantaneous price jump and a major shift in the volatility regime after the event occurs. That is, one can sometimes observe deceptively calm markets as they wait on the event to happen, such as the release of an important piece of economic data, a merger-and-acquisition regulatory decision, a political election, or referendum. Once the outcome is known, though, the price jumps with no intervening trading to its new equilibrium, reflecting the new reality based on the event outcome, and the volatility regime also shifts to reflect the new post-event reality. Basic Black-Scholes assumes no price jumps (i.e., continuous trading) and no volatility shifts (i.e., homoscedasticity). When these two assumptions are violated, traditional delta hedging strategies will fail miserably and basic options models will underestimate volatility. Fortunately, there are many options pricing models available, although quite complex, that deal with these known challenges [Cox et al. (1979)]. Unfortunately, many risk assessment systems do not use these complex option pricing models and instead embed assumptions of normal distributions, no price jumps, no volatility shifts, and stable correlation structures. No wonder these systems are "surprised" by how many "100-year" floods seem to occur in just one or two decades, instead of the expectation of one per century.

As an aside, relating to previous discussions, price jumps are especially confusing for volatility measurement systems that only look backwards. The price jump creates a one or two-day period where the standard deviation calculation will be extreme; sometimes four or five standard deviations from previous history, and then it settles into a new pattern that is elevated from previous history but not off the charts. From a behavioral finance perspective, what market participants appear to do is to start to discount the event – meaning that its impact on expectations of future volatility starts to diminish, and sometimes rather quickly unless there is good reason to think lightning will strike twice in the same place. Any historically-based volatility measurement system needs to consider whether older data should be more-heavily discounted, or be given equal weight. For example, if one uses a fixed time period for the look back, say three months, then there will be a spike upward when the event occurs in the volatility measure, followed by an "unexplained" reversal when the three-month period ends and the price-gap day drops out of the backward-looking volatility calculation. Bayesian techniques easily handle time decay parameters, as do exponentially-lagged time decay systems. We highly recommend them.

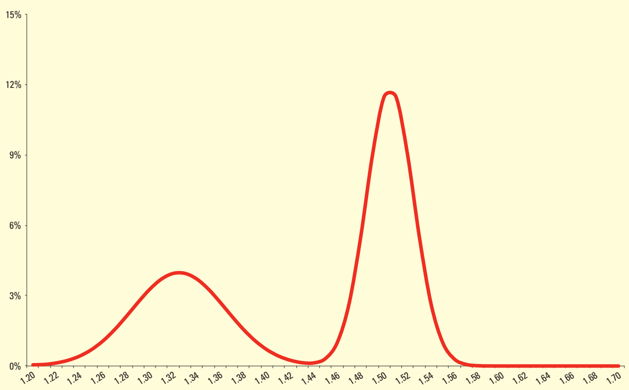

Probability of specific expected exchange rate. Total area under the curve equal 100%.
Source: CME Group Economics.

Figure 3 – Pre-Brexit vote: USD per GBP hypothetical expected probability distribution

Back to considering bimodal distributions and their challenges, and again, "Brexit" is a good example of the potential for a bimodal expected return distribution prior to the vote date (Figure 3). As noted earlier, a "leave" vote was expected to weaken the British pound and "remain" vote was expected to lead to a relief rally and a strengthening pound. What market participants were trying to do was gauge the probabilities of one outcome versus the other. Since the range of probabilities ran more or less from a coin flip to about 60/40, this was a classic case of a bimodal expected return distribution. Of course, once the vote occurred and the outcome was known, the new expected return distribution collapsed almost instantly back into a typical single-mode probability distribution.

Moreover, the process of collapsing back into a single-mode expected return probability distribution had the ability to disturb correlations for a few days. On the 24th and 27th of June 2016, the Friday and Monday after the UK's vote to leave the European Union, the British pound fell 7% and 2%, respectively, while other risky assets, such as equities, also declined, with even the U.S. S&P500® Index falling 3% and 1%, respectively, while most European equity indexes had sharper falls on the 24th. In the weeks afterwards, though, U.S. equities resumed their climb to new highs, while the British pound did not recover, although it stopped falling and traded in a relatively narrow range. In effect, during the disruption, correlations between the British pound and equity indexes were sharply positive, and then fell back toward zero in the weeks after the referendum. Portfolio construction or risk analysis that failed to consider the possibility of a bimodal expected return distribution collapsing back into a single-mode distribution after the event would have underestimated potential volatility, not necessarily have anticipated a gap or price jump as the outcome was announced, and would have missed some very critical correlation shifts.

FUTURE OF FINANCIAL OPTIMIZATION

Two evolving techniques for data analysis are likely to greatly improve risk assessment and portfolio construction – namely, exhaustive search using quantum computers and advances in pattern recognition available through structured machine learning. We will start with a discussion of optimization with quantum computers, although this approach is going to take another five years or so before the computers move from the experimentation phase to being large enough for operational use. Machine learning is already here and gaining ground fast on traditional risk assessment techniques.

Quantum computing is on the way

Quantum computers can be purpose built, and there are a number of experiments on-going in academic labs. To move from the lab to the real world, there is a commercially available quantum computer using an annealing process to solve optimization problems offered by D-Wave Systems of Vancouver, Canada. 1QBit, another Vancouver-based company, is creating software that allows one to utilize the new quantum computers without having to be a quantum computing expert to leverage the best known methods for interacting with quantum hardware. Their software development kit (SDK) enables the rapid and systematic development of higher-level applications that are compatible with both classical and quantum processors. In additional, major computing companies, such as Google, Microsoft, and IBM are known to be experimenting in various ways with quantum computing.

The difference in how quantum computers work compared to classical computing is quite amazing and fascinating. Classical computers have bytes that hold a zero or a one. Quantum computers have qubits that hold a zero or a one as well as a second piece of information that can be intuitively thought of as a probability that the information is a zero or a one. To solve an optimization problem, the quantum computer does not add, subtract, multiply, and divide like a classical computer; instead it uses a process known as quantum annealing to seek the lowest energy state based on how the information in the qubits is arranged. That is, the second piece of information in the qubits allows for quantum effects, including tunneling, not possible in classical computers. Tunneling is the concept in quantum physics of a particle moving through a barrier that would not be possible in a classical system. Suffice it to say, explaining quantum computing is well past the scope of this research, however, for optimization, the demonstration of quantum effects represents a huge step forward.

Optimization with quantum computers offers the promise of solving certain problems that have traditionally been challenging for classical computers using a process that exhaustively searches problems known as "quadratic unconstrained binary optimizations," or qubos. In a classical computer, a complex optimization problem such as a qubo is solved by way of iteration to achieve a close, but estimated answer. In a quantum computer, exhaustive search finds the exact answer. For many uses, the estimated optimal solution from a classical computer may work fine, if the practitioner is artful in how the problem is set up and how the embedded assumptions are handled. However, the promise of quantum computing is to free the researcher from having to make some difficult and often wrong simplifying assumptions. In finance, these difficult optimization problems appear in areas such as asset clustering, cash flow modeling, taxation, and portfolio risk decomposition. We should caution, though, that appreciating the characteristics of the return distribution and how it changes will remain critical to developing robust, forward-looking risk assessments. Quantum computing is going to offer some incredibly important new tools for risk analysis and portfolio construction; however, it is unlikely to provide good answers without an expert at the helm.

Machine learning is here

Machine-learning techniques are essentially a highly sophisticated and advanced pattern recognition system. They constitute methods that involve cleaning (harmonizing) the data, building the model on known data (also known as "training" phase), optimizing the model, and then applying the model on unseen data (often called "testing" phase). The beauty of these algorithms is that they need not be programmed for all the data out there. They learn as and when they see new datasets and evolve. All the machine learning algorithms are categorized into one of these two categories:

- **Supervised learning:** the datasets that belong to supervised learning techniques already have a "label" (outcome/prediction variable) attached to them. Most of the classification and regression problems are categorized as supervised learning techniques.
- **Unsupervised learning:** these algorithms aim at the descriptive nature of the data rather than classifying them. Data exhibits certain characteristics and patterns over a period of time (in case of time-series data) and techniques like clustering and association rules help identify them.

One can develop algorithms for machine learning that are unstructured or structured. The unstructured systems are essentially "frequentist" methods, where the data is asked to speak for itself without expert advice. The unstructured methods are likely to be most popular; simply because they are easy to use and open-source software is available. Unstructured machine learning is great for descriptive analytics; however, as one moves into the world of predictive systems, the unstructured methods are likely to appear extremely successful in back-testing and suffer from a myriad of problems in actual practice – not unlike the challenges facing current practices in financial optimization when history is not necessarily a good guide.

Machine learning has been heavily linked with "big data." Initially, much of the research in finance is aimed at discerning new trends and augmenting security returns forecasts with all kinds of new information not previously available – hence, the term "big data." Data is growing at an enormous rate. "Big data" is usually characterized by the three basic Vs – volume, variety, and velocity. (There are of course other Vs added over time – value, veracity, etc.) The datasets can be from different sources (i.e., variety), can be in motion (real-time data demonstrating velocity), can use different data architecture, and they can still inform a machine learning process. Apache has a lot of open-source projects that have gained popularity in recent years. Apache Spark, an in-memory distributed computing platform is worth mentioning. Spark can scale financial modeling and optimization which includes calculating Value-at-Risk (VaR) to fit models, run simulations, store, and analyze results in the cloud.[3]

Structured machine learning methods allow for different types of expert information to guide the learning process. The combination of expert advice and sophisticated pattern recognition systems offers tremendous process for forecasting financial variables – from returns to volatilities to correlations and beyond. And, machine learning is not necessarily tied to the straitjacket of time series data, so pattern recognition processes can be much more creative in how the historical data is interpreted.

Pattern recognition with financial data does come with some special challenges, and one of the biggest is that the data is exceptionally noisy. With classical statistical regression techniques, one observes the noisy data by finding only relatively weak fits for the modeling of daily returns. With machine learning, the existence of relatively noisy data will put a greater premium on how one sets the various parameters that filter the pattern or how one adds expert advice to the system. This will be essential for the forward-looking results to add substantial value, and it will not be easy.

The advances from machine learning for quantitative finance are already making themselves felt in sales forecasting and marketing techniques; however, this is just the beginning of a revolution. For financial optimization, structured machine learning promises more robust forecasting tools, for expected returns, and using more diverse measures of volatility for risk assessment, while allowing for very creative assessments of stylized (structured) correlation patterns. The era of parallel and distributed computing is here, which makes it possible for computations to scale and provides the ability to make predictions at a granular level. Hence, financial optimization will look totally different in just a few years as the new tools permeate the industry and change an age-old mindset about portfolio construction and risk assessment.

REFERENCES

- Black, F. and M. Scholes, 1973, "The pricing of options and corporate liabilities," The Journal of Political Economy 81:3, 637-654
- Cox, J. C., S. A. Ross, and M. Rubinstein, 1979, "Option pricing: a simplified approach," Journal of Financial Economics 7:3, 229-263
- Garman, M. B., and M. J. Klass, 1980, "On the estimation of security price volatilities from historical data," Journal of Business 53:1, 67-78
- Karagiannidis, I., and D. S. Wilford, 2015, "Modeling fund and portfolio risk: a bi-modal approach to analyzing risk in turbulent markets," Review of Financial Economics 25, 19-26
- Markowitz, H. M., 1952, "Portfolio selection," The Journal of Finance 7:1, 77–91
- Markowitz, H. M., 1991, "Foundations of portfolio theory," The Journal of Finance 46:2, 469-477
- Merton, R., 1973, "Theory of rational option pricing," Bell Journal of Economics and Management Science 4, 141-183
- Parkinson, M., 1980, "The extreme value method for estimating the variance of the rate of return," Journal of Business 53:1 61-65
- Putnam, B. H., 2012, "Volatility expectations in an era of dissonance," Review of Futures Markets, March
- Wilford, D. S., 2012, "True Markowitz or assumptions we break and why it matters," Review of Financial Economics 21, 93 – 101

3 Reference - https://www.infoq.com/presentations/spark-financial-modeling

FINANCIAL COMPUTING & ANALYTICS

STUDENTSHIPS

Four-Year Masters & PhD

for Final Year Undergraduates and Masters Students

As leading banks and funds become more scientific, the demand for excellent PhD students in **computer science, mathematics, statistics, economics, finance** and **physics** is soaring.

In the first major collaboration between the financial services industry and academia, **University College London, London School of Economics,** and **Imperial College London** have established a national PhD training centre in Financial Computing & Analytics with £8m backing from the UK Government and support from twenty leading financial institutions. The Centre covers financial IT, computational finance, financial engineering and business analytics.

The PhD programme is four years with each student following a masters programme in the first year. During years two to four students work on applied research, with support from industry advisors. Financial computing and analytics encompasses a wide range of research areas including mathematical modeling in finance, computational finance, financial IT, quantitative risk management and financial engineering. PhD research areas include stochastic processes, quantitative risk models, financial econometrics, software engineering for financial applications, computational statistics and machine learning, network, high performance computing and statistical signal processing.

The PhD Centre can provide full or fees-only scholarships for UK/EU students, and will endeavour to assist non-UK students in obtaining financial support.

UCL LSE Imperial College London

financialcomputing.org

INDUSTRY PARTNERS

Financial:

Barclays
Bank of America
Bank of England
BNP Paribas
Citi
Credit Suisse
Deutsche Bank
HSBC
LloydsTSB
Merrill Lynch
Morgan Stanley
Nomura
RBS
Thomson Reuters
UBS

Analytics:

BUPA
dunnhumby
SAS
Tesco

MORE INFORMATION

Prof. Philip Treleaven
Centre Director
p.treleaven@ucl.ac.uk

Yonita Carter
Centre Manager
y.carter@ucl.ac.uk

+44 20 7679 0359

Layout, production and coordination: Cypres – Daniel Brandt, Kris Van de Vijver and Pieter Vereertbrugghen

Imperial College Business School | Centre for Global Finance and Technology

Centre for Global Finance and Technology

The Centre for Global Finance and Technology at Imperial College Business School will serve as a hub for multidisciplinary research, business education and global outreach, bringing together leading academics to investigate the impact of technology on finance, business and society.

This interdisciplinary, quantitative research will then feed into new courses and executive education programmes at the Business School and help foster a new generation of fintech experts as well as re-educate existing talent in new financial technologies.

The Centre will also work on providing intellectual guidance to key policymakers and regulators.

"I look forward to the ground-breaking research we will undertake at this new centre, and the challenges and opportunities posed by this new area of research."
– Andrei Kirilenko, Director of the Centre for Global Finance and Technology

Find out more here:
imperial.ac.uk/business-school/research/finance/centre-for-global-finance-and-technology/

CAPCO

BANGALORE
BRATISLAVA
BRUSSELS
CHICAGO
DALLAS
DÜSSELDORF
EDINBURGH
FRANKFURT
GENEVA
HONG KONG
HOUSTON
KUALA LUMPUR
LONDON
NEW YORK
ORLANDO
PARIS
SINGAPORE
TORONTO
VIENNA
ZÜRICH