SUPERCOMPUTING: THE QUANT GUIDE

QuantMinds 365

QUANTMINDS

Ed Stapley
Editor-in-Chief

Charlie Burgess
Digital Content Editor

Georgia Hood
Business Development Manager

CONTRIBUTORS

Marcos Lopez de Prado
Research fellow, Cornell University

Dr. Leo Razoumov
Principal Machine Learning Scientist, Amazon

David Leinweber
Founder Centre for Innovation Financial Technology, Lawrence Berkeley Laboratory

Horst Simon
Deputy Lab Director for Research, Lawrence Berkeley National Laboratory

Alejandro Perdomo-Ortiz
Quantum Machine Learning Lead and Founder, NASA, USRA Quantum AI Lab/ Qubtera LLC

Michael Harris
Quant systematic and discretionary trader

Dr. Uwe Naumann
Professor of Computer Science, RWTH Aachen University

Dr. Svetlana Borovkova
Associate Professor of Quantitative Finance, Vrije Universiteit Amsterdam and Head of Quantitative Modelling, Probability and Partners

Dr. Paul Bilokon
Founder, CEO and Chairman, Thalesians Ltd.

INTRODUCTION

When first coming across the term 'supercomputing' you could be forgiven for conjuring up futuristic images which might stretch even into the realms of science fiction. Today's reality though is that supercomputing is steadily advancing and, with the right applications, has the ability to revolutionise a range of industries including finance. Quants find themselves at the forefront of this incredible journey. Their unique skillsets, including solving highly complex mathematical problems, building predictive models and algorithms and working on computational and numerical efficiencies, can be made exponentially more efficient through the improved capabilities of computers.

It is for this reason that the 25th annual QuantMinds conference this year is going to be bringing experts together to cross-examine the opportunities that supercomputing is unlocking. This eMagazine will serve to whet the appetite for those with an interest in this topic and the potential it is realising.

We hope you enjoy it.

The QuantMinds team

CONTENTS

FRONTIERS IN BIG DATA, MACHINE LEARNING AND SUPERCOMPUTING

Big Data, Machine Learning (ML) and supercomputing have become an indispensable part of scientific research, and are making waves in the quant finance industry. At QuantMinds International this May, a panel of thought leaders in these fields will discuss how financial institutions could benefit greatly from the adoption of these tools and approaches. Here, QuantMinds asked the panellists to introduce some of the key ideas and highlight some major developments to watch out for.

The panel:

Marcos Lopez de Prado
Research Fellow at Cornell University (MLP)

Dr. Leo Razoumov
Principal Machine Learning Scientist in Amazon (LR)

David Leinweber
Founder Centre for Innovation Financial Technology at Lawrence Berkeley Laboratory and author (DL)

Horst Simon
Deputy Lab Director for Research at Lawrence Berkeley National Laboratory (HS)

Alejandro Perdomo-Ortiz
Quantum Machine Learning Lead and Founder at NASA, USRA Quantum AI Lab/ Qubtera LLC (APO)

Why have Big Data, ML and supercomputing come to the forefront over the last year?

DL: I think this it is due to some impressive, easily understandable, and well publicised successes. Few people notice when a computer is used to prove a theorem in topology. Beat the reigning Jeopardy champ, and it makes the news. Throw in widely used speech recognisers (Alexa, Siri and the like) and put them on phones, and everyone notices. Artificial Intelligence (AI) has been around, and had impressive, but had much narrower successes, for decades.

MLP: Recent breakthroughs have been achieved before most people expected them. Today, these technologies can accomplish tasks that until recently only expert humans could perform. High performance computing (HPC), Big Data and ML are ubiquitous today in our daily lives. Take airports, people travel in an aircraft where the pilot made a fraction of the decisions he would have made in the past. Upon arrival, screening algorithms conducted a background check, looking for any suspicious pattern. For those who found their luggage, they have an algorithm to thank, and for those who didn't find it, an algorithm will soon recover the suitcase that some human misplaced.

LR: In the last couple of years, usefulness of ML solutions was conclusively demonstrated in several important and diverse domains ranging from Speech Recognition and Natural Language Processing, to beating humans in board games like Go (Google Alpha Go). Late last year ML took a lead in game of Poker -- a non-deterministic game which is very different from types of Chess and Go. And impressively, training of the models was achieved by so-called "self-play", when a machine plays millions of games against itself. This relies on availability of massive computational power. Most of the fundamental Deep Learning algorithms have been known since 80s and 90s of the last century. It is the massive advances in computation power and storage that made those old algorithms perform so well.

APO: Besides the significant progress in ML and AI, there are still tons of problems, of which even the largest supercomputer centers have a hard time in solving, and some of these will not be solved in the next decades. Still, quantum computing has come to the forefront as well as an alternative with the potential of revolutionising the computer industry, with FinTech one of the core applications.

Where do you see these technologies moving to over the next 5 years?

DL: The most dramatic recent advances have been enabled by Moore's Law allowing us to build and run deeper neural networks. Despite rumours of its demise, the growth predicted so accurately by Moore continues. Coupled with advances in effectors, the technologies to interact with the physical world (robots, large and nano-scale, computer manufacturing, autonomous vehicles) there is likely to be major shift in the economic landscape. I don't think this will be another case of horseshoe makers becoming auto mechanics.

LR: Well, predicting immediate future pays my salary at Amazon :-) But, caveat emptor, predicting technological advances for a 5 year horizon is probably impossible. Few speculations in no particular order: Digital assistants and self-driving cars are nearing practical usefulness. Dialog systems will be common place. Wide spread of Machine Learning in Healthcare, specifically in diagnosis and robotically assisted surgeries.

Do you think there are persisting misconceptions about these technologies?

MLP: Indeed. One of them is the belief that these technologies are off-the-shelf, or that they run by themselves. While reporters like to focus on the millions of jobs that may be replaced by ML, they seem to ignore that the implementation of these technologies is non-trivial and it requires substantial investments in human capital. Those who resist progress will be replaced, not by machines, but by those who embrace it and adapt. Another misconception is the belief that these technologies are somewhat in an early stage of development. As David said, the truth is, these technologies are mature and have demonstrated their power for many years in large-scale research programs. It is only now that the public has come into contact with them, through consumer products.

LR: Real successes are accompanied by lots of hype, which I guess is inevitable in such a red-hot field as ML. One persistent misconception which irritates me quite often is misuse of the word AI, I prefer to deal with concrete problems, algorithms and systems leaving philosophical questions to a later stage.

APO: There is also lot of hype about what quantum computing is capable of doing in the near-term. Although we are in an exciting time in history, I think we need to communicate clearly the challenges we face. Understanding and talking openly about these will guide the major developments in the field and it will help to balance the expectations from investors. I strongly agree with Prof. Wilhelm, from whom I recently heard: "quantum computing is really cool, even when we restrict it to the true facts". I totally agree with this line, and I see it as an invitation to take with a grain of salt popular media articles and to follow closely what the experts say.

Read the full interview here >>

QuantMinds International
IN 40 SECONDS

- 350+ global quant experts from banks, buy-side, regulators, silicon valley and academia.
- Fourteen technical quantitative streams. Two summits. Four workshops.
- "The quality of the speakers, the name recognition, the thought-leadership. It's just overwhelming – you don't find it everywhere."

IMPACT OF ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING ON TRADING AND INVESTING

Michael Harris

Very few can ignore the presence of Artificial Intelligence and Machine Learning in today's world, and even less so if you work in quantitative finance. Here, Michael Harris, quant systematic and discretionary trader and bestselling author, discusses the impact these technologies are having on trading and investing.

General impact of AI and ML on trading

AI allows the replacement of humans with machines. In the 1980s, AI research focused primarily on expert systems and fuzzy logic. With computational power becoming cheaper, using machines to solve large-scale optimisation problems became economically feasible. As a result of the advances in hardware and software, nowadays AI focuses on the use of neural networks and other learning methods for identifying and analysing predictors, also known as features, or factors, that have economic value and can be used with classifiers to develop profitable models. This particular application of AI often goes by the name ML

The application of methods for developing trading strategies based on AI, both in short-term time frames and for longer-term investing, is gaining popularity and there are a few hedge funds that are very active in this field. However, broad acceptance of this new technology is slow due to various factors, the most important being that AI requires investment in new tools and human talent. Most funds use fundamental analysis because this is what managers learn in their MBA programs. There are not many hedge funds that rely solely on AI. Application of AI is growing at the retail level but the majority of traders still use methods that were proposed in mid twentieth century, including traditional technical analysis, because they are easy to learn and apply.

Note that AI and ML are not only used to develop trading strategies but also in other areas, for example in developing liquidity searching algos and suggesting portfolios to clients. Therefore, with AI applications gaining ground, the number of humans involved in trading and investment decisions decreases and this obviously affects markets and price action. It is early to speculate on the overall effects this new technology will have on the industry but it is possible that extensive use of AI will result in more efficient markets with lower volatility for extended periods of time followed by occasional volatility spikes due to regime changes. This is possible because the impact of subjective evaluation of information by humans will be minimized and with that the associated noise. But that remains to be seen in practice.

"One problem with trading strategies based on AI is that they can yield models that are worse than random."

Impact of AI and ML on alpha generation

During these initial phases of the adoption of AI technology, there will be opportunities for those who understand it and know how to manage its risks. One problem with trading strategies based on AI is that they can yield models that are worse than random. I will try to explain what I mean by this: traditional technical analysis is an unprofitable method of trading because strategies based on chart patterns and indicators draw their returns from a distribution with zero mean before any transaction costs. Some traders will always be found at the right tail of the distribution and this gives the false impression that these methods have economic value. My research shows that especially in the futures and forex markets, longer-term profitability is hard to achieve no matter which method is used because these markets are designed to benefit market makers. However, in shorter periods of time some traders can realiSe large profits in leveraged markets due to luck. Then, these traders attribute their success to their strategies and skills, rather than to luck.

With AI and ML, there are additional effects, such as the bias-variance trade-off. Data-mining bias can result in strategies that are over-fitted to past data but immediately fail on new data, or strategies that are too simple and do not capture important signals in the data that have economic value. The result of this trade-off is worse-than-random strategies and a negative skew in the distribution of returns of these traders even before transaction cost is added. This presents an opportunity for profit for large funds and investors in the post-quantitative easing era. However, as the worse-than-random AI traders are being removed from the market and only those with robust models remain, the battle for profits will become intense. It is too early to speculate whether AI traders or large investors will win this battle.

I would also like to mention a common misconception in this area: some people believe that the value is in the ML algos used. This is not true. The true value is in the predictors used. ML algorithms cannot find gold where there is none. One problem is that most ML professionals use the same predictors and try to develop models in an iterative fashion that will produce the best results. This process is plagued by data-mining bias and eventually fails. In a nutshell, data-mining bias results from the dangerous practice of using data multiple times with many models until results are acceptable in the training and testing samples. My research in this area indicates that if a simple classifier, such as Binary Logistic Regression, does not work satisfactorily with a given set of predictors, then it is highly likely that there is no economic value. Therefore, success depends on what is called feature engineering, and this is both a science and an art that requires knowledge, experience and imagination to come up with features that have economic value and only a small percentage of professionals can do that.

Impact of artificial intelligence and machine learning on technical analysis

We have to make a distinction between traditional and quantitative technical analysis because all methods that rely on the analysis of price and volume series fall under this subject. Traditional technical analysis, i.e., chart patterns, some simple indicators, certain theories of price action, etc., was not effective to start with. Other than a few incomplete efforts of limited scope and reach, publications that touted these methods never presented their longer-term statistical expectation but offered only promises that if this or that rule is used there would be profit potential. Since profits and losses in the markets follow some statistical distribution, there were always those who attributed their luck to these methods. At the same time, a whole industry developed around these methods because there were easy to learn. Unfortunately, many thought they could profit by being better at using methods known to everyone else and the result was massive wealth transfer from these naïve traders to market makers and other well-informed professionals.

Success depends on what is called feature engineering, and this is both a science and an art that requires knowledge, experience and imagination.

In the early 1990s, some market professionals realized that a large number of retail traders were trading using these naive methods. Some developed algos and AI expert systems to identify the formations in advance and then trade against them, causing in the process volatility that retail traders, also known as weak hands, could not cope with. In a more fundamental way, the failure of traditional technical analysis can be attributed to the disappearance of high serial correlation from the markets starting in the 1990s. It was basically the high serial correlation that offered the wrong impression that these methods worked. Nowadays, with few exceptions, markets are mean-reverting, not leaving room to simple technical analysis methods to work. However, some quantitative technical analysis methods often work well, such as mean-reversion and statistical arbitrage models, including ML algorithms that use features with economic value.

Note that this type of arbitrage is unlikely to be repeated in the case of AI and ML because of the great variety of models and the fact that most are being kept proprietary, but the main problem with this new technology is not confirmation bias, as in the case of traditional technical analysis, but data-mining bias.

In my opinion, observing the market and looking at charts is becoming an obsolete process. The future of trading is about processing information, developing and validating models in real-time. The hedge fund of the future will not rely on chart analysis. Some still do this because they are at the transition boundary where old ways meet with a new era. Many traders not familiar with AI will find it hard to compete in the future and will withdraw.

LESSONS FROM 10+ YEARS OF ALGORITHMIC DIFFERENTIATION IN COMPUTATIONAL FINANCE

Uwe Naumann

It has been more than a decade since Mike Giles and Paul Glasserman introduced Algorithmic Differentiation (AD) and its adjoint mode in particular to Computational Finance. In their seminal paper titled "Smoking Adjoints: fast Monte Carlo Greeks" and published in Risk Magazine in January 2006, they use adjoint AD (AAD) to compute first- and second-order Greeks of a LIBOR market model. They show how AAD allows computation of potentially very large gradients at the cost of only a few function evaluations. This "cheap gradient" result had been an essential ingredient of numerical methods in Computational Science and Engineering for many years. Contributions from data assimilation in meteorology/physical oceanography, shape/topology optimization in automotive and aerospace engineering, or back-propagation in neural networks are collected in an extensive bibliography maintained by the AD community.

In close collaboration with the Numerical Algorithms Group Ltd. I have been providing AAD solutions and software to several tier-1 investment banks for approximately 10 years. Together with Mike Giles and Luca Capriotti I have been presenting several workshops on adjoints in Computational Finance as part of the QuantMinds conference series. I appreciate this great opportunity to interact with some of the leading figures in the field.

The acceptance of AAD has been growing with the increasing number of sensitivities required, for example, by modern computational approaches to XVA and FRTB. Computational Finance has become one of the major target applications for AAD. Relatively simple proof-of-concept prototypes are increasingly replaced by production-grade AAD solutions targeting substantial parts of existing Quant libraries and running on heterogeneous hardware including large clusters and accelerators, for example, GPUs. Consequently, state of the art AAD solutions and software require a level of professionalism which goes beyond academic experiments. In the following I comment briefly on a collection of relevant issues to be considered in the context of what I would refer to as "production-grade AAD solutions."

Memory Requirement

Due to the necessary reversal of the data flow AAD is known to suffer from potentially prohibitive memory requirement when applied naively to nontrivial numerical simulations. Checkpointing and preaccumulation techniques have been developed to facilitate construction of feasible AAD solutions for a given upper bound on the available memory. The combinatorial problem of minimising the overall computational cost is hard. Powerful heuristics exist.

Second-Order Vector AAD

Second-order adjoints are computed as tangents (directional derivatives) of first-order adjoints. Vector tangent AD facilitates high-performance Hessian accumulation through parallelisation/acceleration.

Overloading/Source Transformation/Finite Differences

With C++ dominating the Quant library development landscape operator and function overloading has become the method of choice for AD in Computational Finance. Automatic source transformation is typically applied to simpler (domain-specific) scripting languages or to subsets of general-purpose programming languages.

Hand-coding may prove useful in the context of isolated, run time critical, and stable parts of the target code. Near-optimal AD solutions typically combine local source transformation (for example, applied to scripted payoffs) and selective finite difference approximation (for example, applied to black-box routines) with an overall overloading approach. Template meta-programming in C++ and improved compiler optimisation result in a continuously decreasing performance gap between overloading and source transformation.

Symbolic Differentiation/Implicit Function Theorem

Basic AAD operates at the level of elemental functions consisting of arithmetic operators and intrinsic functions provided by the underlying programming language. For the sake of efficiency this low level of granularity should be avoided whenever possible. Higher-level elementals comprise symbolic derivatives of primal numerical kernels ranging from BLAS and (non-)linear system solvers via numerical integration methods to solvers of (partial) differential equations. Consistency of the resulting tangents and adjoints requires accurate primal solutions as well as careful numerical treatment of the symbolic adjoint formulations.

Non-smoothness/Noise

Conceptually, AD assumes its target code to be continuously differentiable up to the required order. It cannot be expected to "magically" deal with general non-smoothness. Smoothing of the target code is typically required as a preprocessing step for AD. Approximation by finite differences gives reasonable results in some cases. Similar issues arise in the context of noisy/chaotic functions.

Parallelism/GPU

AD solutions have been developed for parallel computers based on shared/distributed/hybrid memory models. Template meta programming in C++ allows for efficient and scalable implementations on massively parallel accelerators such as GPUs. The state of the art includes the efficient evaluation of second-order adjoints on GPUs.

Software Engineering/Modeling and Simulation

AD is "invasive." First steps toward production-grade solutions restrict the use of AD to selected parts of the target code. Once the often substantial benefits are realized, the use of AD is extended to larger code sections with nontrivial consequences for software design, coding guidelines, debugging, and testing. Tight integration of AD results in better code. Knowledge about parameter sensitivities contributes to the design of better models, methods and to the choice of suitable data.

Project Management

Production-grade AD solutions require appropriately trained personnel and a high degree of automation of the software development process to facilitate sustainability and robustness with respect to the ongoing evolution of the hard- and software infrastructure. Corresponding management decisions define the level of success of the adoption of AD as a fundamental ingredient of the computational "tool box" within financial institutions.

In conclusion, there is no viable alternative to adjoints as the method of choice for large-scale sensitivity analysis, error control, calibration, and deep learning. The development, maintenance, and evolution of robust, efficient, and sustainable adjoint-enabled Quant libraries remains a major challenge for modern Computational Finance. The ongoing collaboration between experts from AD and from Computational Finance contributes to more sophisticated AD software tools enabling improved security and predictivity of financial simulations.

“FINANCIAL INSTITUTIONS NEED TO EMBRACE CHANGE”

Svetlana Borovkova

#QuantWomen highlights key women in the quantitative finance industry. Today, we speak to Dr. Svetlana Borovkova, Associate Professor of Quantitative Finance, Vrije Universiteit Amsterdam and Head of Quantitative Modelling, Probability and Partners about her interest in the field and advice to new starters. Svetlana will be at QuantMinds International discussing sustainable and impact investing.

How did you start your career?

Ever since I was very little, I loved numbers and maths. When I was three, I would be standing at a bus stop with my mam in Moscow, and add the bus numbers in my head as they came along. And in Moscow, buses have three-digit numbers: 608, 299 and such. So the choice was made quite early that I should study maths. Moreover, I grew up during the tail end of communism in Soviet Union, so being a journalist or a historian meant complying with the ideology, and maths is obviously beyond ideology.

What was your lightbulb moment?

I graduated from an applied math department at a military high college and then communism collapsed. So I started looking abroad, and entered first a Master and then PhD program in The Netherlands and in US. During my PhD, I worked on chaos theory which was very sexy at that time, but still very theoretical for my taste.

The lightbulb moment came at the time Scholes and Merton received their Nobel Prize for economics, for their famous formula. At the time of the announcement, I was in London, attending a meeting of Royal Society, which, upon the announcement, quickly turned into a discussion (and for uninitiated mathematicians like myself - into a tutorial) about option pricing and quant finance in general. The minute I heard the explanation of the Black Scholes option formula and binomial tree, the penny dropped - I understood the beauty and elegance of it and its awesome usefulness at the same time - concepts usually incompatible in maths (except for cryptography, things in maths are either beautiful and elegant or useful, but not both).

I immediately thought: this is what I want to do. So after receiving my PhD, I departed to London, to Shell trading, and became an analyst for oil trading desk. The first thing I had to do was to develop their very first Value at Risk system. And the second thing happened to be an arbitrage futures trading strategy that generated more money per barrel of oil than any Shell refinery (and continued generating profits for years afterwards). So you could say this was a flying start. Although nowadays I branch into many different areas of quant finance and risk - ranging from derivatives pricing and quantitative risk models to sentiment analysis, financial stability and systemic risk - commodities and in particular energy markets have remained my soft spot - first love never dies I guess.

Why do we see so few women in quantitative finance?

Currently I also run a prestigious MSc honours program in Quantitative Finance and Risk Management in a Dutch university, and every year I try incredibly hard to attract girls into our program. Unfortunately without much success. I am puzzled by this - quant finance offers such amazing opportunities at the moment and for years to come, work is almost never boring (except when you have to do stuff for regulators) and hours are much more human than, say, in corporate finance or M&A. Moreover, it is scientifically proven that women make better investment decisions and have a better understanding of risk, and hence are superior investment and risk managers. There are plenty of girls coming out of quantitative bachelor studies, but they all go to computer science or data analytics and not to quant finance. So, I have no explanation for the lack of women in our field, and my herculean attempts to change things, at least at the level of my program, so far have been completely futile.

What advice do you have for women starting out their career in quant finance?

My advice to women considering career in quant finance is borrowed from Nike: just do it! As I said, work is fun, hours are not too bad, opportunities are plentiful and, as a woman, you will have an incredible edge in this environment, especially as financial firms are pressed really hard to improve their gender diversity. Let alone the unlimited amount of eligible men on fat salaries you get access to.

What will the future of quant finance/risk management look like?

Regarding the future of quant finance and risk management: I think it is quite rosy. The increased regulation means of course, on one hand, more boring work, but on the other hand - simply more work, more jobs and more opportunities. Another area which could lift our profession to the next level is FinTech. So for people with combination of quant, AI and computer skills (who now only strive to go work for Google and such) this will be an excellent career path. However, for this, we need big financial institutions to embrace change, become more agile and light on their feet, think and act quickly. If this happens in the next 5 years, future of our profession is guaranteed. If not, we can still go and work in alternative finance which will then take place of traditional finance.

THE PAST, PRESENT, AND FUTURE OF HIGH-PERFORMANCE COMPUTING IN FINANCE AND BEYOND

Dr. Paul Bilokon

It was Gordon Moore (b. 1929), a co-founder of Intel Corporation, who observed that the number of transistors in an integrated circuit doubled approximately every two years. This observation, since its pronouncement in 1965, has become known as Moore's Law. However, the free lunch is over — that's the title of Herb Sutter's article that appeared in Dr. Dobb's Journal back in 2005. Sutter pointed out that chip designers were under immense pressure to deliver ever faster CPUs. In some places they were approaching physical limits.

Instead of trying to pack more transistors into a single chip and make them run faster, hardware manufacturers started building multicore systems, i.e. systems with multiple independent processing units that read and execute machine code instructions at the same time. Graphical processing units (GPUs) are among such multicore systems.

Programmers no longer enjoy their free lunch. They must write their code differently to take advantage of multicore systems. Their code needs to be multithreaded, a thread being a sequence of instructions that can execute in parallel with others. One particular concurrent design pattern has been taking the machine learning community by storm since its publication by Google's Jeffrey Dean and Sanjay Ghemawat in 2004: MapReduce. This idea has been utilised by Apache's open-source framework Hadoop (2011) for processing big data.

One reason why multithreaded, concurrent, programming is so difficult is that one has to reason about multiple interrelated streams of data arriving asynchronously in time. New programming paradigms have been developed to help build such asynchronous data-flow systems. Functional Reactive Programming (FRP) dates back to the work on animation published by Conal Elliott (Microsoft Research) and Paul Hudak (Yale University) in 1997. FRP has evolved from those early days and has been realised in products such as Flux, React, Immutable, and AngularJS with a focus on user interfaces, but has also appeared in general-purpose libraries, such as Sodium for Haskell and Java.

While GPUs have enabled parallelism by increasing the number of cores and thus adding more structure to the chip, the Field Programmable Gate Arrays (FPGAs) have achieved the same end by going to the other extreme. There are no instruction processing units on FPGAs that would execute commands sequentially, one after the other. Instead, energy travels across the integrated circuit through a sequence of (re-)programmable logic gates and interconnects. The integrated circuit is programmed in a bespoke manner, so that most of the space on it is utilised for a particular task at hand.

Some argue that we now have to move towards heterogeneous computing to maintain the validity of Moore's Law and its corollaries. We are now seeing hardware products that combine elements and ideas from CPUs, GPUs, and FPGAs. In the meantime, the more futuristic work on quantum computing continues. In 2011, D-Wave Systems announced the first commercial quantum annealing system, D-Wave One, with manufactured spins, claiming a 128 qubit processor. Subsequent machines boasted 1,000 and 2,000 qubits. The company remains the world's only commercial seller of quantum computers, with customers including Lockheed Martin, NASA, and Google, and there is budding interest among the computational finance community. Concurrency arises naturally in the studies of programming for distributed quantum computing and much of current research in this field relies on process algebras. More recently, IBM has started to take steps to commercialise their expertise in quantum computing and grant selected customers access to an experimental quantum system.

The free lunch is over for both the hardware and software manufacturers. To keep up with Moore's Law we need to introduce and master not only new technologies, but new ways of thinking. As Edsger Dijkstra (1930– 2002) put it, "the tools we use have a profound (and devious!) influence on our thinking habits, and, therefore, on our thinking abilities". Thankfully, the new generation of electronic engineers, software architects and developers is ready and open to new challenges.


JOIN THE CONVERSATION
#QUANTMINDS
QuantMinds
@QuantMinds
QuantMinds #Trading & #Risk Management is at the helm of the #quant #finance industry, curating cutting edge research, interviews and opinions. #QuantMinds
Global
443 Following
1,381 Followers
Tweets
Tweets & replies
Media
QuantMinds @QuantMin... · 39m
It's problematic to use an un-adjusted level of interest rates when modelling using QE since interest rates have been trending down spr.ly/6016DZ4Pi #Volatility