2018: THE TURNING POINT FOR QUANT FINANCE?

MACHINES ARE ONLY GETTING SMARTER. WHERE DOES THAT LEAVE QUANTS?

QuantMinds

QUANTMINDS

Victoria Chatterton
Strategy Director

Ed Stapley
Editor-in-Chief

Charlie Burgess
Digital Content Editor

Georgia Hood
Business Development Manager

CONTRIBUTORS

Matthew Dixon
Co-founder of the Thalesians and Assistant Professor of Finance at the Illinois Institute of Technology, Chicago

John Hull
Financial Professor Of Derivatives & Risk Management at Joseph L. Rotman School of Management at University Of Toronto

Peter Hafez
Chief Data Scientist at Raven Pack

QuantMinds

INTRODUCTION

The field of quantitative finance is changing ever more rapidly, and in 10-15 years time the quant community could be working on quite different problems, using new analytical tools. So, what will it take to get there?

The impact of technology on finance has always been familiar to quants, who are often more adept with handling the evolution than some of their colleagues in other parts of banks. But as computers get faster, and smarter, and tougher regulations come into play, quants need to keep up by integrating themselves with the bigger picture and developing those tech skills.

In our first of the QuantMinds 2018 quarterly eBook series, we ask the experts and deep dive into the areas they believe will have the biggest impact on quants going forward.

We hope you enjoy it.

The QuantMinds Team

CONTENTS

DEEP PROBABILISTIC PROGRAMMING FOR FINANCIAL MODELLING

Matthew Dixon

Deep probabilistic programming (DPP) combines three fields: Bayesian statistics and machine learning, deep learning (DL), and probabilistic programming. Although in its infancy, DPP is a powerful combination of several different probabilistic modelling approaches and inference techniques that have historically been treated as separate mathematical ideas in the financial modelling literature. Modelling approaches includes Bayesian neural networks (BNNs), directed graphical models and the broad availability of a wide range of techniques for efficient and robust model training. DPP can be regarded as a compositional modelling framework, composing previously disparate modelling concepts and algorithms over a scalable computational graph implementation. Just as TensorFlow's (Abadi et al., 2016) computational graphs have made DL synonymous with modelling high dimensional input spaces with very large data sets, so too is Edward (Tran et al., 2017) poised to make DPP a tool of choice alongside STAN (Carpenter et al., 2017) and PyMC3 (Salvatier, Wiecki, & Fonnesbeck, 2016b) for Bayesian modelling and inference of computationally intractable problems.

Gaussian processes

In order to explore some of the key capabilities of DPP for financial modelling and uncertainty quantification, we shall first revisit more familiar territory - Gaussian processes (GPs). The idea of GPs is to, without parameterizing F(X), place a prior P [F (x)] directly on the space of functions (MacKay, 1998). Gaussian processes can hence be viewed as a generalisation of Gaussian distributions from finite dimensional vector spaces to infinite dimensional function spaces. The main advantage of using GPs is modelling co-movements of time series. GPs learn the relationship between the time series directly building up an empirical model of the co-movements rather than assuming linear correlations between stochastic processes.

Bayesian Neural Networks

Gaussian processes have been shown to be a limiting case of Bayesian Neural networks (BNNs). These offer a probabilistic interpretation of neural network models by inferring distributions over the models' weights. To model uncertainty, BNNs assume that model parameters (weights and biases) are random variables. A prior distribution is placed over the weights and biases, which induces a distribution over a parametric set of functions. As the number of weights goes to infinity (Neal, 2012; Williams, 1997) it shows that the Bayesian neural networks converge to GPs when a standard matrix of Gaussian prior distributions is placed over the weights and point estimates used for the priors. Of course, there is no reason to restrict modelling assumptions to these limiting conditions and, in fact, one of the main benefits of either approach is the flexibility to combine different distributions and include interaction effects.

> *A promising research direction is how to use probabilistic graphical programming to build Bayesian network copulas (Elidan, 2010) for contagion and credit risk.*

Deep Learning

DL has shown remarkable success in a wide field of applications, including Artificial Intelligence (AI) (DeepMind, 2016; Kubota, 2017; Esteva et al., 2017), image processing (Simonyan & Zisserman, 2014), learning in games (DeepMind, 2017), neuroscience (Poggio, 2016), energy conservation (DeepMind, 2016), and skin cancer diagnostics (Kubota, 2017; Esteva et al., 2017). Current advances in hardware (GPU/TPU) and software (Theano/TensorFLow) allow for developing scalable deep learners to find a deterministic >> approximation to an unknown nonlinear function. However, in many trading and risk management applications, it is necessary to model uncertainty of a predictive model. Deep learners are typically used to build deterministic approximations to an unknown nonlinear function and are largely unusable when uncertainty needs to be explicitly represented. And so deep probabilistic programming builds on deep learners by providing inferential techniques to estimate the posterior distribution and hence characterize the uncertainty in a model estimate.

Bayesian computations

As a Bayesian approach, the choice of architecture of a deep network, combined with regularisation, can be viewed as defining a prior probability distribution over non-linear functions and deep learning can be viewed as finding the posterior probability distribution of the unknown function over the network weights (Gal, 2016). In general, the posterior distribution is analytically intractable and Bayesian computational methods are required to approximate it. One common class of solution approaches, referred to as variational inference, solves an optimization problem by minimizing the Kullback-Leiber (KL) distance between an approximating variational distribution and the posterior obtained from the original model. Bayes' filtering techniques, such as Kalman filters, HMMs and particle filters, are examples of variational inference techniques common place in trading and risk management, albeit for low dimensional input spaces. The other well-known class of posterior approximation techniques use sampling techniques, rather than variational techniques - a prime example being the use of MCMC simulation. Neither of these inference engines, however, scale well with respect to model and data size.

Recent advances in variational inference techniques and software can now represent probabilistic models as a computational graph (Blundell, Cornebise, Kavukcuoglu, & Wierstra, 2015; Tran et al., 2017; Salvatier, Wiecki, & Fonnesbeck, 2016a). The key benefit here is the ability to build scalable probabilistic deep learners, without having to perform testing (forward propagation) or inference (gradient-based optimization, with back propagation and automatic differentiation) from scratch. Alternatives to variational and MCMC algorithms were recently proposed by (Gal, 2016) and build on efficient Dropout regularization techniques, a variable selection approach which has fuelled the popularity of DL.

The combination of Edward and TensorFlow is the first time that we have seen large-scale Bayesian modelling fused with AI and HPC and that's likely to spawn critical developments in applications where uncertainty and model risk play a central role, such as finance. The remainder of this article briefly discusses some areas which are likely suitable financial modelling applications for DPP.

Financial applications: predictive modeling

Often, an automated decision to trade or a price forecast may be shrouded in uncertainty arising from noisy data or model risk, either through incorrect model assumptions or parameter error. This uncertainty should be accounted for in predictions so that the model output can be more clearly interpreted. DPP thrives in high dimensional input spaces. In trading and risk management applications, high dimensionality may arise naturally where there is spatial structure in the data. Examples include deep portfolios (Heaton, Polson, & Witte, 2017), large scale loan modelling (Sirignano, Tsoukalas, & Giesecke, 2016), predictors for prices using the many depth levels in the limit order book (Dixon, Polson, & Sokolov, 2017; Dixon, 2017; Sirignano, 2016), predictors arising from combining several data sets as in mortgage risk (Sirignano, Sadhwani, & Giesecke, 2016) or consumer lending (Dixon, Giesecke, Sheshardi, & Troha, 2017), or from text mining and natural language processing of news and other documents (Glasserman & Mamaysky, 2015).

> *"DPP could improve existing applications through Bayesian neural networks and robust regularization techniques for generalizing modelling and improving robustness and scalability."*

Portfolio optimization and market risk

For some of these applications, we need look no further than existing applications of Gaussian processes to finance. Da Barrosa, Salles, & de Oliveira Ribeiro (2016) present a spatio-temporal GP method for optimizing financial asset portfolios which allows for approximating the risk surface. Cousin, Maatouk, & Rulliere (2016) propose a new term-structure interpolation method that extends classical spline techniques by quantifying uncertainty. Another example of GPs includes spatial meta-modelling for Expected Shortfall through nested simulation (Liu & Staum, 2012). Spatio-temporal GPs are used to infer portfolio values in a scenario based on inner-level simulation of nearby scenarios. This significantly reduces the required computational effort by avoiding inner-level simulation in every scenario and takes account of the variance that arises from inner-level simulation. DPP could improve these existing applications through Bayesian neural networks and robust regularization techniques for generalizing modelling and improving robustness and scalability.

Contagion and credit risk

Capturing dependencies between financial institutions is a significant aspect of contagion and counterparty credit risk modelling. Examples in contagion modelling include the inter-dependency between banks participating in Federal Reserve System's emergency programs (Battiston, Puliga, Kaushik, Tasca, & Caldarelli, 2012), >> joint dependencies of banks and other institutional investors through transactions in the repurchase agreement market or through various other channels of systemic risk (Bisias, Flood, Lo, & Valavanis, 2012). Other examples in counterparty credit risk modelling include wrong way risk assessment of OTC portfolios. It is well known that the dependence structure between variables can be explicitly model with a directed graph. Given this capability, a promising research direction is how to use probabilistic graphical programming to build Bayesian network copulas (Elidan, 2010) for contagion and credit risk.

Summary

This article provides some background on deep probabilistic programming and discusses its implications for trading and risk management. In particular, we see merit in exploring high dimensional modelling problems arising in large-scale portfolios, systemic risk and contagion in financial networks and where spatial structure is present in the application, such as limit order book modelling. Tools such as Edward lower the barrier to entry and add salt to the fire ignited by TensorFlow. The barrier to entry is deceptively simple however and some discipline is needed. The programming convenience of Keras for TensorFlow is alluring, but the ability to build effective and interpretable financial models is still in its infancy and should be investigated more fully by cross-disciplinary researchers in computational science, statistics and finance, before any definitive and practical conclusions can be drawn.

WHAT IS THE FUTURE FOR QUANTITATIVE FINANCE?

QuantMinds

What do you get when you put John Hull, Damiano Brigo, Jesper Andreasen and Matthew Sargaison in a room? Answer: a very entertaining debate on the future of quants at QuantMinds.

How long will the quant role be around for?

Andreasen, who is Global Head of Quantitative Research at Danske Bank, said that you could be forgiven for thinking that the regulatory drive towards standardisation would leave quants out of a job. But he thought that their jobs were secure – at least for the next five years.

John Hull, Maple Financial Professor of Derivatives & Risk Management at the University of Toronto, agreed that indeed, the times they were a'changin': "It used to be the case that if you had a Ph.D. in maths or physics and you had solved PDs, you could join the front office, earn a high salary and have a job for life. That's no longer the case. Many front office quants have moved to middle office, and there's much more emphasis on other things like regulations," he said.

Ten to 15 years down the road, Hull predicted that the quant community would be working on quite different problems using new analytical tools. "But that's not so bad," he said. "It makes the job interesting."

The key to a successful career was to constantly update your skills, he added. "As a quant you can't expect to coast using the same skillset. It's important not to have techniques to use in search of problems, but to always use them in search of solutions."

Whilst the perspective was different on the buy-side, the message from Matthew Sargaison, Chief Investment Officer at AHL Partners LLP was the same:"I have people working for me that see every problem as something they can throw a Gaussian process at," he said. "That's not what matters. It's having the right level of curiosity and flexibility to think about what tools you use," he said.

Damiano Brigo, Chair and Co-Head of Group, Mathematical Finance at Imperial College, London agreed: "You cannot survive with one mentality all your life, but unfortunately that's not always true in academia – often you can survive by keeping your boundaries narrow."

Technology

What impact would technology, already transforming the entire financial industry, have on quants?

Hull felt that technological skills, encompassing everything from blockchain to machine learning, were now an essential part of the skillset that quants would need. "There is a trend of employers wanting their quants to have good computer science skills, which traditionally was not the case," he said.

Sargaison said that there had definitely been a move towards more technology on the buy side: "It's not necessarily big data, much more the techniques of machine learning, whether its Gaussian process or convolutional neural nets," he said. This, he added, was also where the hiring stress was: "You're no longer in competition with Goldman Sachs, you're in competition with Facebook and Google."

But not all technology now being feted was all that new, noted Brigo. "Artificial neural networks were very popular 20 years ago but the computational power was not good enough so it died. Now you have regression and people are rebranding it as machine learning. These things come and go, sometimes you need to tap into collective knowledge and choose from the many tool boxes available to solve the problem."

> *We are the rats of the system; we will always survive even if they try to outlaw us with all this regulation. We will find a way of making a living in the banking world.*

Regulations

How useful are regulations, according to the panel?

"There are elements that are useful," answered Andreasen. "The whole credit risk in the system, for instance, but I'm not so sure everything is."

"Having said that, in the wake of all these exercises banks will get cleaned up, in their systems and possibly also in their behaviour," he continued. "There's also going to be some good effects, the infrastructure in banks will have to be a lot better than it is now, and from that there will be interesting things for quants to do."

The problem with regulation, though, said Brigo, was that it always ended up as a compromise: "There are so many people to be kept happy, practitioners in banks and politicians, that in the end you end up with something that doesn't respond to the fundamental questions in a satisfactory way," he said.

The role of quants

Another member of the QuantMinds audience sought advice on how to answer her students' complaints that working in quant finance was "boring", compared to the work they had been doing at university.

Andreasen was the first to dispense some very sensible advice: "It's a question of mentality. You're not going to be given a differential question to solve on your first day. "You have to seek the problems. If things are boring in your work, figure out a way to automate; find out how to get the machine to do your boring job, and then move onto something else. This is how you get recognition in the industry and get to the interesting stuff."

Hull added that quants these days could "no longer be isolated from the rest of what's going on in the bank". They had to interact much more with the bank infrastructure and understand the big picture.

Brigo agreed, saying that: "It's our job to let them know that life will not always be exciting. I don't want to say something unpopular, but on planet Earth 99% of the people are not doing the job that they dreamt of."

One thing that Hull emphasised was the importance of communication skills. "The ones who progress fastest are not the ones that got A+, it's the ones that can communicate well."

Looking beyond the next five years, Andreasen ended the session on a cheery note:

"Who knows what the role of the quant will be. But we are the rats of the system; we will always survive even if they try to outlaw us with all this regulation. We will find a way of making a living in the banking world."

We also spoke with Diana Ribeiro, Head of Linear Rates & Inflation, Quantative Research at Lloyds Banking Group, about how quants are adapting to current challenges in the industry.

QuantMinds International

14-18 May 2018
Lisbon Marriott,
Lisbon

> "If you need to choose just one event a year, then this is the one."

Fabio Mercurio, Head of Quant Analytics, Bloomberg L.P.

QuantMinds Americas

September 24 - 26, 2018
Marriot Longwharf, Boston,
Boston

Related events

RiskMinds International

3 - 7 December 2018
Hotel Okura,
Amsterdam

RiskMinds Americas

September 24 - 26, 2018
Marriot Longwharf, Boston,
Boston

RiskMinds Insurance

3 - 7 December 2018
Hotel Okura,
Amsterdam

RiskMinds Asia

October 2018
Singapore

"INPUT GOES IN, THE MAGIC HAPPENS":
PRACTICAL TOOLS & REAL WORLD APPLICATIONS OF QUANTUM COMPUTING

QuantMinds

Quantum computing is receiving a fair amount of hype at the moment. Hundreds of millions of dollars are being poured into research, from both the private sector as well as business, with the belief that this really is the next big thing.

But how close are we, in reality?

At QuantMinds 2017, the very people at the heart of recent developments in the quantum field told us just how close we were to a world in which quantum computing would become the norm. The reason everybody wants to know, of course, is that quantum computing could be a real game-changer.

What is quantum computing?

Professor Helmut Katzgraber from the Department of Physics & Astronomy at Texas A&M University, took a good stab at stealing Justin Trudeau's crown for best concise explanation of quantum computing when he said: "Input goes in, then magic happens."

Vasil Denchev, Software Engineer at the Quantum Artificial Intelligence Lab at Google, had an explanation that illuminated why quantum computing was so important: "There is a lot of evidence that quantum computers will be able to solve certain computational tasks much faster than classical computers. This could have very practical outcomes for different areas of science and business," he theorised. Denchev's day job involves making this theory a reality.

What are the issues surrounding quantum computing?

Quantum computing is not without its problems, though. "We do have optimisation problems in building quantum computers," stated Davide Venturelli, Science Operations Manager at Quantum Artificial Intelligence Laboratory, NASA Ames Research Center. "Perhaps one of the early applications of a quantum computer should be to build the quantum computer," he reasoned.

Different sectors, including finance, could contribute to finding a solution to those optimisation problems, he added. Helmut pointed out that claims of the quantum computer's super speed had to be taken with a grain of salt. "When you talk about speed up there are two important ingredients," he explained. "One, how much faster is it? And two, how does this speed change with a number of variables?"

So far, experiments involving quantum computing have focussed on the D-Wave machine – the only available quantum optimiser that implements a restricted version of quantum computing.

Vasil explained that his experiments had showed that, for a certain crafted computation problem that fits very well on the D-wave machine, they could achieve a speed up to 100 million times faster than classical computers.

> *Google is investing in both analogue and digital: "because we can," said Denchev. "But also because we don't know for sure which one will have the largest pay off in the short term."*

But this came with the caveat that this was a crafted problem, and the purpose of the experiment was just to show that this was possible, rather than any practical implication.

"We haven't been able to do this for a more practical problem yet, but it's important to know that speed like that is possible."

What is the near-term impact of quantum computing?

What might be the first impacts of quantum computing on the finance world? In the fintech space, for instance, would this enhanced computational power make cryptographic protocols, such as those used in blockchain, less safe?

No, said Katzgraber. "It's going to take a while before we have a machine that can implement algorithms that can factor large numbers. This will come, but not in the near future. Having said that, if it happens, there is quantum encryption."

Katzgraber felt that only when digital devices overtook analogue would we see a real impact, not just in science but in applications in industry. "An analogue machine is a device that has intrinsic errors, biases and noise," he explained. “So there is only so much that you can really do. There are many optimisation problems that require high precision, and this precision is not available in analogue,” he explained.

Google was investing in both analogue and digital: “Because we can,” said Denchev. “But also because we don’t know for sure which one will have the largest pay off in the short term.”

“We still have a lot of hope for the analogue approach for quantum annealing. Our current thinking is that a hybrid quantum/classical optimisation approach might be the most successful because that allows us to get the best of both worlds.”

20 years from now, what percentage of the world will have quantum computing?

The use of quantum computing may remain a specialised area for quite some time, felt the panel members.

Katzgraber explained how there were two obstacles to the proliferation of quantum computers.

> *I don’t envision having general purpose quantum computing that can be used for mundane tasks.*

“First, will we have large enough digital devices that can do things that are useful? Second, there is the fact that there are not too many quantum algorithms out there. “In my opinion, there will be small pockets this special hardware will be used in, but where it is used, there will be a very large impact.” Denchev agreed: “I don’t envision having general purpose quantum computing that can be used for mundane tasks. “But I might be wrong,” he said.

And even a two to three percent improvement could be game-changing, interjected Davide.

“There is so much that we need to learn about quantum information science, but the number of scientists is growing, and the investment in it is growing.”

It may be here sooner than we think.

“We’re at the dawn of the computing age”. Hear from Vern Brownell, CEO at D-Wave, as he tells us about what is having an impact now.

TRADING COMMODITIES USING NLP: A CASE STUDY

Peter Hafez

Being a commodity investor can be a rollercoaster ride. This was clear a few years ago when the prices of energy-related commodities collapsed, led by crude oil. Supply and demand dynamics may dictate the long-term evolution of prices, but, as the recent meltdown of energy commodities has shown, there can be prolonged periods of pricing anomalies. Economic indicators, natural disasters and political upheaval all influence prices, as do commodity-specific issues such as oil spills, gas pipeline leakages and mining accidents.

Given the vast amount of unstructured textual content available in today's market, investors often struggle to separate noise from signal and determine which of these events has the potential to impact commodity prices. However, by utilizing the latest advances NLP, market participants have the ability to detect high-impact events and identify commodity price triggers in real-time.

Researchers from Duke University have previously shown how sentiment on geopolitical and fundamental news can impact oil prices, both in the short-term and long-term. However, in our latest research report, “Machine Learning & Event Detection for Trading Energy and Metals Futures”, we took a different approach by demonstrating how to take advantage of commodity specific events to predict next-day returns across two commodities baskets: Energy and Metals.

We train ten different machine learning algorithms across close to 100 different commodity specific events detected from news, including anything from drilling events to import tax, supply, or inventory increase or decrease events. Whenever a novel event category is detected at least once for a given commodity on a given day, we assign a one (1) to the relevant column in our features matrix, and zero (0) otherwise.

To minimize the prediction bias, we take an ensemble approach by equally weighting the ten ML algorithms. Overall, we find that such portfolios outperform the individual models in eight and nine cases out of ten across the energy and metals baskets, respectively.

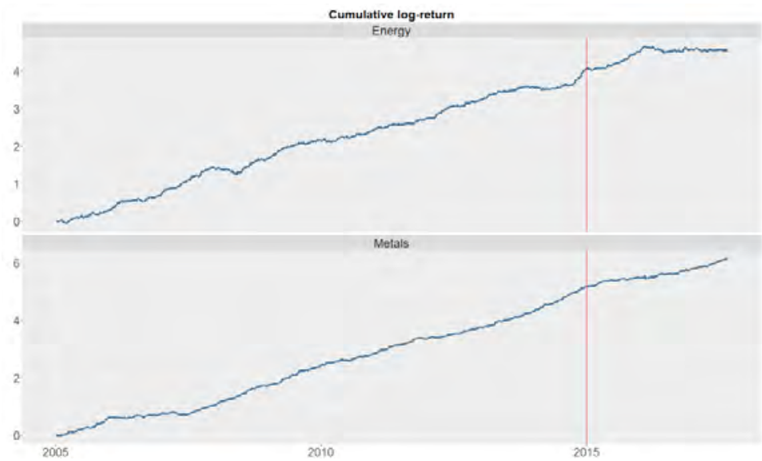

Figure 1: Cumulative log-return

The cumulative return profile, presented in Figure 1, underlines the impressive performance of the metals basket, delivering total returns of 97.8%, while the performance of the energy basket trails off in the second half of the out-of-sample period. Nevertheless, it still returns 48.3%.

There has been a dearth of volatility in many asset classes of late, including energy commodities. This may help explain the lack of performance towards the end of the out-of-sample period. Our strategy does not take volatility into account when placing a trade, resulting in a lot of small wins and losses lacking a clear direction. We take steps to correct this by acknowledging that higher volatility regimes often provide stronger signals for trading, while we should stay on the sidelines during lower-volatility regimes.

We implement this into practice by imposing a very simple rule: only trade the model whenever short-run volatility (10-days) is above medium-term volatility (21-days). Table 1 shows the results of trading only during periods of increasing volatility (“Increasing”) compared to simply trading all signals (“Standard”). For completeness, we also add the results of a system which only trades during periods of declining volatility (“Decreasing”).

Ensemble	Energy			Metals		
Statistics	Increasing	Decreasing	Standard	Increasing	Decreasing	Standard
Annualized Return	20.6%	7.5%	17.8%	39.2%	33.0%	35.7%
Annualized Volatility	17.0%	15.8%	13.7%	16.8%	14.8%	11.5%
Information Ratio	1.21	0.48	1.29	2.34	2.24	3.11
Hit Ratio	55.1%	52.2%	53.6%	56.4%	57.8%	58.7%
Max Drawdown	31.6%	46.8%	17.4%	13.2%	11.5%	9.0%
Per-Trade Return (basis point)	8.17	3.99	6.06	11.06	7.35	9.08
Number of Trades	1,537	1,568	3,105	2,311	2,653	4,964

Table 1: Performance statistics for regime-dependent strategies

Conditioning on increasing volatility results in higher annualized returns for Energy as well as higher per-trade returns (8.2bp vs. 6.1bp). This occurs despite the number of trades being more or less cut in half, resulting in lower diversification and therefore slightly lower IR. We observe the same for Metals: higher returns for both annualized and per trade, though the drop in IR is more pronounced (from 3.11 to 2.34) here. We find that both Increasing portfolios have IRs which are statistically significant at the 1% level (when compared to a random portfolio with the same number of trades), confirming that the improvements we see in annualized returns come from greater predictability. We note that by applying an increasing volatility filter there is no longer a drop-off in performance for Energy towards the end of the out-of-sample period. This underlines that we can add further alpha by incorporating regimes into our trading strategy.

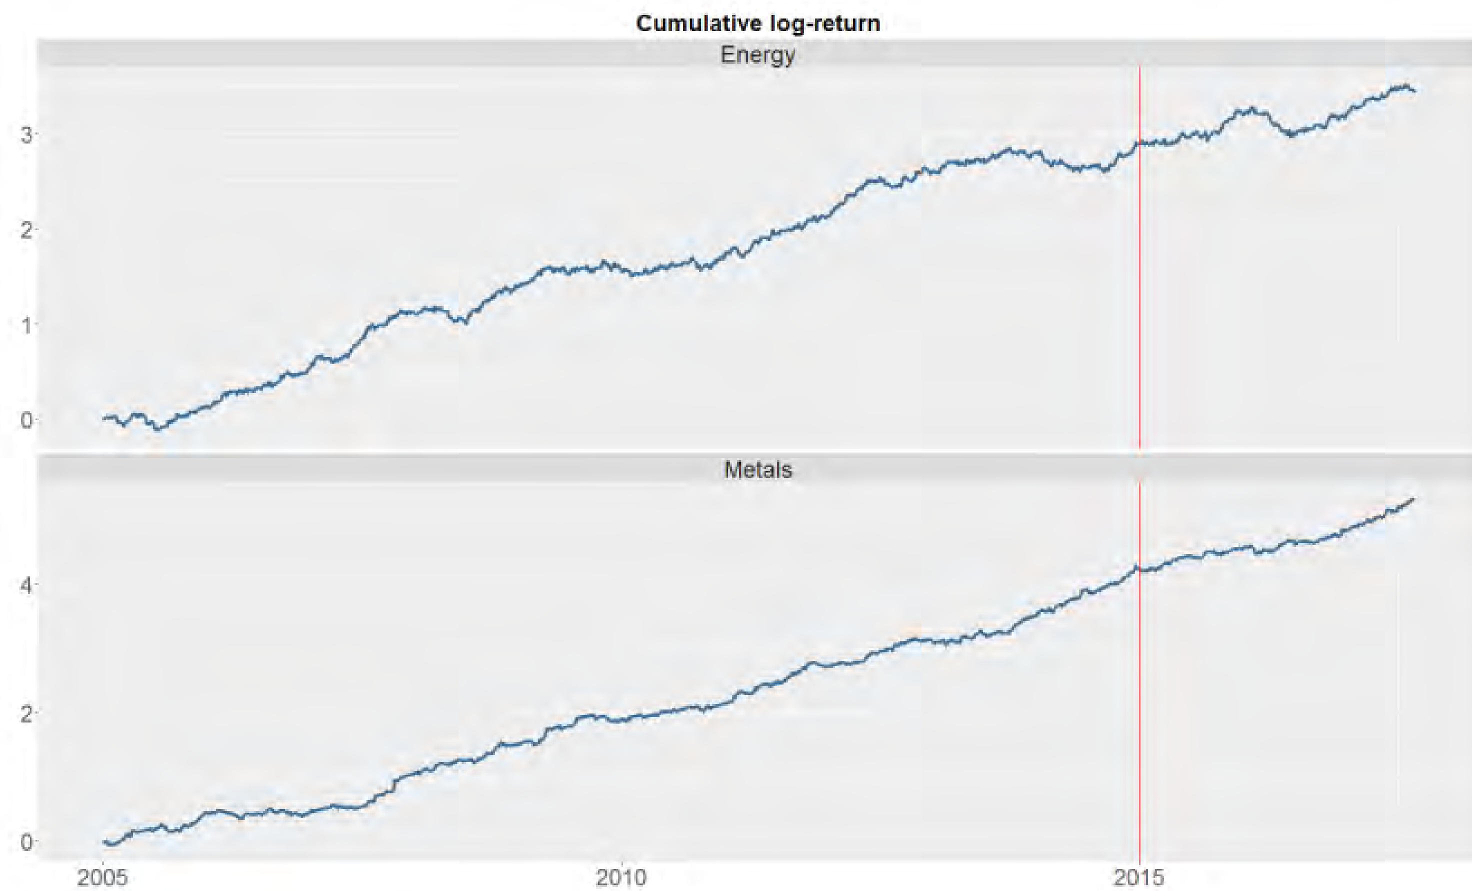

Figure 2: Cumulative log-return for the increasing volatility strategy

THE NEW RULES FOR OTC DERIVATIVES

John Hull

The Leaders' statement issued after the G-20 meetings in Pittsburgh in September 2009 included the following paragraph:

"All standardized OTC derivative contracts should be traded on exchanges or electronic trading platforms, where appropriate, and cleared through central counterparties by end-2012 at the latest. OTC derivative contracts should be reported to trade repositories. Non-centrally cleared contracts should be subject to higher capital requirements. We ask the FSB and its relevant members to assess regularly implementation and whether it is sufficient to improve transparency in the derivatives markets, mitigate systemic risk, and protect against market abuse."

Eight years on, it seems appropriate to review the changes that this has led to.

The key objective immediately following the crisis was to reduce systemic risk by requiring more collateral to be posted when financial institutions trade with each other. This objective has largely been achieved. Standard transactions between financial institutions are cleared through CCPs and attract both initial margin and variation margin. Non-standard transactions between financial institutions continue to be cleared bilaterally. But, following the 2011 G-20 meeting in Cannes, rules requiring initial and variation margin for these transactions are being implemented.

One result of these changes is that there has been a trend away from customized OTC derivatives toward more standard products. This should reduce systemic risk, but there are potential disadvantages. If dealers are less willing to customize transactions, end users may make less use of derivatives for hedging. One of the reasons that end users require non-standard derivative transactions is to make the transaction match some physical transaction so that the derivative and the physical transaction qualify for hedge accounting rules. The end user may not be willing to enter into a standard derivative if it does not qualify for hedge accounting. Also, there is a danger that the new rules will hinder financial innovation by dealers.

There can be little doubt that reporting all OTC derivative transactions to trade repositories such as the Depository Trust and Clearing Corporation (DTCC) is desirable. It gives regulators the opportunity to recognize situations where unacceptable risks are being taken. It also creates more post-trade price transparency.

> *There has been a trend away from customized OTC derivatives toward more standard products. This should reduce systemic risk, but there are potential disadvantages.*

No doubt politicians and regulators were greatly influenced by the AIG fiasco. AIG Financial Products entered into many transactions where it guaranteed the AAA-rated securities created from the securitization and re-securitization of subprime mortgages. The performance of AIG Financial Products was guaranteed by its U.S. parent. It was not required to post collateral on its transactions providing AIG's credit rating remained above AA. In mid-September, AIG's credit rating fell below AA and it was unable to provide the required collateral. Only then did regulators become aware of the risks that had been taken. A massive bailout followed.

A situation similar to AIG should never happen again. First, trade repositories would allow regulators to be more aware of the one-sided risks being taken, making it possible for them to step in earlier. Second, a company entering into trades similar to those of AIG would be required to post so much initial margin and variation margin that its appetite for the trades would be greatly diminished.

The least important, and least defensible, of the new regulations for OTC derivatives is the requirement that standard transactions between financial institutions be traded on electronic platforms. The motivation for this seems to be that, if OTC derivatives are traded like exchange-traded derivatives, there will be more price transparency and problems such as those observed during the crisis will be avoided. In fact, the problems during the crisis were caused by non-standard derivatives and there is no requirement that these be traded on electronic platforms.

> *It is certainly true that CCPs are too big to fail. But arguably they are much easier to regulate than banks and therefore are much less likely to fail.*

There was not a serious problem in the way OTC derivatives were traded pre-crisis. It is not clear that there was a lack of price transparency. Industry participants had access to reliable sources of price quotes. In any case, trade repositories should take care of any transparency issues. Trading OTC derivatives in the same way as exchange traded derivatives is therefore not necessary to achieve price transparency.

In addition, there is a danger in trying to trade OTC derivatives in the same way as exchange-traded derivatives. This is because there are important differences between the two. OTC derivatives trade intermittently whereas exchange-traded derivatives such as futures trade continuously. The size of a typical OTC derivative is much larger than that of a typical exchange-traded derivative. There are fewer market participants in the OTC market, but they are more sophisticated than the average participant in exchange-traded markets.

A final question is whether the effect of the new regulations is to move the too-big-to-fail problem from banks to CCPs. It is certainly true that CCPs are too big to fail. But arguably they are much easier to regulate than banks and therefore are much less likely to fail.

This article is an abbreviated version of an editorial that appeared in Journal of Risk Management and Financial Institutions, Autumn/Fall 2015