


March 2, 2012 9:50 pm

School for quants

By Sam Knight

Inside UCL's Financial Computing Centre, the planet's brightest quantitative analysts are now calculating our future



Students look at equations from a PhD thesis that uses Bayesian analysis to examine the relation between bond trades and economic data releases. Got that?

On a recent winter's afternoon, nine computer science students were sitting around a conference table in the engineering faculty at University College London. The room was strip-lit, unadorned, and windowless. On the wall, a formerly white whiteboard was a dirty cloud, tormented by the weight of technical scribbles and rubbings-out upon it. A poster in the corner described the importance of having a heterogeneous experimental network, or Hen.

Six of the students were undergraduates. The other three were PhD researchers from UCL's elite Financial Computing Centre. The only person keeping notes was one of them: a bearded, 30-year-old Polish researcher called Michal Galas. Galas was leading the meeting, a weekly update on the building of a vast new collection of social data, culled from the internet. Under the direction of the PhD students, the undergraduates were writing computer programs to haul millions of pages of publicly available digital chatter – from Facebook, Twitter, blogs and news stories – into a real-time archive

which could be analysed for signs of the public mood, particularly in regard to financial markets. Word of the project, known as SocialSTREAM, had reached the City months ago. The Financial Computing Centre was getting calls most days from companies wanting to know when it would be finished. The Bank of England had been in touch.

During the meeting, Galas asked each undergraduate about a particular corner of the database. Most of the time, the language was computer: impenetrable exchanges about batching, pseudo-storage and the risks of propagating on tiers. “Should that work in a distributed field as well?” an undergraduate asked. “If you have different connectors running on different machines, you might start duplicating...” Galas considered this. “Personally,” he replied, “I would run the connectors outside the cloud machine.”

Every now and again, though, the discussion became comprehensible. The students discussed annoyances – so much data about animals! – and possibilities. One of the PhD students, Ilya Zheludev, talked about “Wikipedia deltas” – records of deleted sections from the online encyclopaedia. Immediately, the students hit on the idea of tracking the Wikipedia entries of large companies and seeing what was deleted, and when.

The mood of the meeting was casual and exacting at the same time. Galas, who is from Gdansk and once had ambitions to be a hacker, is something of a giant at the Financial Computing Centre. One of the first students to enrol in 2009, he has a gift for writing extremely large computer programs. In order to carry out his own research, Galas has built an electronic trading platform that he estimates would satisfy the needs of a small bank. As a result, what he says goes. Galas closed the meeting by giving the undergraduates a hard time about the overall messiness of their programming. “I like beauty!” he declared, staring around the room.

The Financial Computing Centre at UCL, a collaboration with the London School of Economics, the London Business School and 20 leading financial institutions, claims to be the only institute of its kind in Europe. Each year since its establishment in late 2008, between 600 and 800 students have applied for its 12 fully funded PhD places, which each cost the taxpayer £30,000 per year. Dozens more applicants come from the



Philip Treleven, director of the Financial Computing Centre, believes that the meeting of computing power and fine young brains will transform problem-solving

financial industry, where employers are willing to subsidise up to five years of research at the tantalising intersection of computers, data and money.

As of this winter, the centre had about 60 PhD students, of whom 80 per cent were men. Virtually all hailed from such forbiddingly numerate subjects as electrical engineering, computational statistics, pure mathematics and artificial intelligence. These realms of knowledge contain concepts such as data mining, non-linear dynamics and chaos theory that make many of us nervous just to see written down. Philip Treleven, the centre's director, is delighted by this. "Bright buggers," he calls his students. "They want to do great things."

In one sense, the centre is the logical culmination of a relationship between the financial industry and the natural sciences that has been deepening for the past 40 years. The first postgraduate scientists began to crop up on trading floors in the early 1970s, when rising interest rates transformed the previously staid calculations of bond trading into a field of complex mathematics. The most successful financial equation of all time – the Black-Scholes model of options pricing – was published in 1973 (the authors were awarded a Nobel prize in 1997).



Stathi Panayi is trying to devise a way of measuring liquidity in financial markets – this could help regulators intervene early and prevent events such as the 2007 credit crunch

By the mid-1980s, the figure of the "quantitative analyst" or "quant" or "rocket scientist" (most contemporary quants disdain this nickname, pointing out that rocket science is not all that complicated any more) was a rare but not unheard-of species in most investment houses. Twenty years later, the twin explosions of cheap credit and cheap computing power made quants into the banking equivalent of super-charged particles. Given freedom to roam, the best were able – it seemed – to summon ever more refined, risk-free and sophisticated financial products from the edges of the known universe.

Of course it all looks rather different now. Derivatives so fancy you need a degree in calculus to understand them are hardly flavour of the month these days. Proprietary trading desks in banks, the traditional home of quants, have been decimated by losses and attempts at regulation since the start of the financial crisis. There is nothing like the number of jobs there used to be.

Moreover, among some older quants at least, there is a feeling that the era of genuine discovery is over. Pioneering thinking in the 1970s and 1980s has long been programmed into the most idiot-proof trading software. You can download an Excel spreadsheet of the Black-Scholes model in a few seconds. I went to see Piotr Karasinski, a former head of quantitative analysis at HSBC, who has a model named after him: the Black-Karasinski model for short-term interest rates, which he developed with Fischer

Black (of Black-Scholes). “I think that the field is shrinking,” said Karasinski, in his quiet office at the European Bank of Reconstruction and Development. “Very few people do original work.”

You don’t hear that kind of talk at the Financial Computing Centre. And that is mainly down to its founder, Professor Treleaven. Treleaven was the manager of a holiday camp before he went to study computer science at Brunel University in 1964. One of his father’s friends had heard that computers might be the next big thing. “I took a punt on that,” Treleaven says, and for the past 47 years – 30 of them at UCL – he has enjoyed throwing his machines at pretty much any problem he hears about.

Treleaven began working with the financial industry in the late 1980s. His first project – an attempt to use artificial intelligence to forecast markets – was a failure. But he hit pay dirt with his second: an automated fraud-detection system for the London Stock Exchange. Over the years, the occasional phone calls to his office from financial firms became regular, and by the early 2000s, Treleaven and his department was caught up in the wave of innovation sweeping the industry. Until recently, by far the most popular topic for both his students and his City contacts was the apparently limitless world of algorithmic trading.

In its starkest form, algorithmic trading is the replacement of human decision-making in financial transactions with computer programs. An algorithm – a series of instructions (when to start trading, when to stop, how much risk to take) – issues its own orders to buy and sell. In theory, algorithms can do anything, but in practice they work along a spectrum from simply executing trades to coming up with their own ideas to make money. At their most advanced, and Frankensteinian, algorithmic trading strategies independently scour the world’s financial markets, looking for discrepancies, statistical correlations and arbitrage opportunities, trading most of the time against each other. High Frequency Trading – mass automated dealings of this type – now accounts for about 75 per cent of all US equity transactions.

To some, algorithmic trading is a harbinger of a world out of our control. Robert Harris’s recent novel, *The Fear Index*, was inspired in part by the “flash crash” of May 6 2010, in which the New York Stock Exchange plunged crazily and then recovered after algorithms responded to an unexpectedly large order in the electronic futures market. To others, algorithmic trading shows just how far the automation of financial markets has yet to run. “People say about algorithmic trading, ‘They’re just a bunch of cowboys, you know.’” Treleaven shook his head. “No,” he said. “It is industrialisation. It is like putting robots in car factories.”

Treleaven’s excitement stems, at least in part, from the fact that financial disasters are just as interesting to academics as success stories. When I asked him whether the ongoing agony in the economy had thrown up research

opportunities, he said: “Oh yes, absolutely... the mother of invention is all something, you know? Look at what happened in the war, you have loads of scientific breakthroughs.” But the professor’s real animus is that he believes that what has been taking place in the financial industry – a heady meeting of computing power and the finest young scientific brains – is about to break into the rest of the social sciences. That is because what Treleaven’s students do, what quants do, is find patterns in oceans of electronic data. In a hedge fund, that might mean finding relationships between price movements and then trading on them. In public health, it might mean tracking millions of pharmacy transactions and spotting the next outbreak of flu.



Ilya Zheludev has studied 500,000 internal Enron emails to show a spike in emotion in employees in April 1999, just months before the company’s stock took off

“They think of themselves as doing computational finance,” Treleaven said of his students, “but let’s jump forward: people who are interested in politics in the next 10 years will be doing computational politics.” The calculating power and analytical techniques used in finance could also model the impact of public policies, or seek insights in sport and education. This year, for the first time, Treleaven has a psychology graduate among his students and he enjoys telling undergraduates from other faculties – economics, music even – that they should learn how to program computers if they want to stand a chance in the world that is coming.

Most of Treleaven’s students have the most obvious destination in mind, however: the trading floor. One afternoon I met Mahnoosh Mirghaemi, a 29-year-old Iranian who was awarded the centre’s first PhD last October. Mirghaemi brought her thesis with her: “Bayesian Learning in Financial Markets”. Bayesian learning is very vogueish among quants at the moment. It uses a probability theory first devised by Thomas Bayes, an 18th-century English clergyman, to create financial models that learn and adapt to new information.

Mirghaemi spent two years using Bayesian techniques to study how European bond markets responded to 3,077 separate releases of economic data between 2007 and 2008. She studied 1.6 million bond trades and figured out which pieces of news moved the markets more, and which ones analysts and traders were more likely to forecast poorly. “It made my eyesight like a double,” she said. But Mirghaemi’s research should now, in theory, allow traders, and trading algorithms, to position themselves better on an hour-by-hour basis. “It definitely makes money,” she said.

Mirghaemi was hired by BNP Paribas last summer. A few months later, her boss – a trader for 37 years – mentioned that he could never work out the simultaneous price and position of a trade. On the spot, Mirghaemi wrote down a cosine formula from

physics useful for measuring electromagnetic waves. “He was just looking at it,” she said. Mirghaemi emphasised her respect for her seniors at the bank but she said that she felt different. “I think I come from the new generation,” she said, “looking at the finance, the economic, the engineering, the computing altogether.”



Michal Galas, one of the PhD researchers at the centre, has built an electronic trading platform that he estimates would satisfy the needs of a small bank

There was a touch, almost, of sympathy in the way that Mirghaemi described colleagues coming to terms with the changing nature of the markets. “Their minds are like, ‘We know as economists this is what is happening, or should be happening,’” she said. “But the real world says ‘No.’ The computer systems and all these quant people are changing the market much more rapidly than they actually want to.” And not necessarily for the better. When asked whether she thought all these quants made for more stable financial

markets, Mirghaemi looked at me and said: “It is very, very risky and it brings a lot of volatility to the markets and it is out of control.”

Students at the Financial Computing Centre are comfortable making such statements, because they believe they are equipped to handle their implications. When I asked Mirghaemi how this unstable future made her feel, she said: “It puts me in a very good situation.”

...

A willingness to embrace uncertainty, a certain ruthlessness in acquiring, testing and rejecting new ideas is also what employers are looking for. “That’s what it’s like in this area,” said Rafael Molinero, who runs a quant-led hedge fund, Molinero Capital Management, where three students from the centre are currently on work placements. “You always have to reinvent yourself to stay on top of the curve. That is a big driver for us. That is what we want to see in them.” And the best way to keep your head is to listen to your algorithms, rather than your heart. As Molinero put it: “The main idea when you become a quant is that a computer is less prone to pitfalls than a human.”

If Molinero is right, then Michal Galas became a quant a long time ago. For his PhD, Galas is building what he calls an “adaptable algorithm trading portfolio” – a production line of automated trading strategies, from which computers will select the most appropriate one, depending on what is happening in a particular market. Algorithms upon algorithms upon algorithms. Galas imagines it as a hedge fund without employees. “There is no human intervention necessary,” he said.

I told Galas that this degree of trust in machines unnerved me. I told him that, as I understood it, the sheer complexity of some financial products and an over-reliance on

mathematical models had been a major contributor to the financial crisis. To illustrate this, I drew two diverging lines in my notebook and told him that I thought this gap – the difference between what we know and what we think we know – had proved itself dangerous. Galas looked at the widening gap as if he recognised it. “That is an opportunity to use computers,” he said. “So yeah, good for me.”

Of course, not every student at the centre speaks like this. Stathi Panayi spent an unhappy year at an investment bank in 2008 and has vowed never to go back. He returned to Cyprus and worked as an English teacher before starting his PhD in financial computing last year. Panayi is now trying to devise better ways to measure liquidity in financial markets, and coming up with ideas for how regulators might intervene earlier to prevent events such as the credit crunch of 2007. One of his advisers works at the Bank of England.


Even so, Panayi was sanguine about the eagerness – the right – of his fellow students to come up with ever more abstract ways to beat our battered markets. The methods of quants might be difficult to fathom but humanity, and capitalism, has not progressed by putting limits on invention. “Is the answer going down a level of sophistication, making things easier for the layman to understand?” Panayi asked me. “Is there a limit? Who is to say what is the limit of sophistication in the market?” He considered this. “I don’t think the answer is banning everything you don’t understand.”

Panayi suggested that the answer is not to fear quants but to join them. Even better: give them one of your problems to solve.

Technology and analytical thinking, after all, is neutral: what matters is the aim in which it is deployed. That is why SocialSTREAM – the database the students were trying to figure out on that winter’s afternoon – could turn out to be so useful.

I got a preview of what the database might be capable of, shortly before it went live last month. The idea behind SocialSTREAM, and other experiments like it, is to collect reams of live text being published to the internet, and to run it through dictionaries designed to test language for signs of mood. For an individual tweet – “Good morning!” – this might be meaningless. But taken across millions of postings, from the personal to political, a rough indicator of popular sentiment does emerge.

Ilya Zheludev, one of the students from the meeting, showed me his study of 500,000 internal Enron emails, which were released following the collapse of the energy company in 2001. Zheludev’s sentiment analysis showed a spike in emotion among employees – both positive and negative, a massive, contradictory shiver – in April 1999, a few months before the company’s stock began to take off on its exponential (and



Mahnoosh Mirghaemi studied 3,077 separate releases of economic data and 1.6 million bond trades to identify which pieces of news moved the markets more

fraudulent) trajectory.

Picking up on such bubbles of emotion as they emerge (around a company, for instance, or a government) even in such murky waters as Twitter, or Facebook, or the website of the Financial Times, has an obvious allure to individual investors trying to stay ahead of the market. At least one London-based hedge fund, Derwent Capital, now trades purely on social data, mined in this way.

But there are clear civic and academic possibilities as well. Take political polling: while I watched, Zheludev set up SocialSTREAM to trawl the entire current output of Twitter for mentions of “Obama” and to analyse each mention for an approval rating of +1 to -1. Then he put the findings on a graph. A jagged line appeared, and from 4.22pm to 4.27pm, on January 16 2012, Zheludev and I were looking at one crude, real-time measure of the political fortunes of President Obama. “There you have it,” said Zheludev. We stared at the zigzag on the screen, wondering what it might possibly mean.

Deciphering such patterns is what excites collaborators with the Financial Computing Centre who are more interested in stabilising the markets than beating them. Zheludev’s supervisor is David Tuckett, a psychoanalyst at UCL who studies the interplay of emotion and the unconscious in trading decisions. He told me that the database could, if used properly, allow us to see our exaggerated hopes and paranoias for what they are, before they grow to overwhelm us. “If you think about it like the sea,” said Tuckett, of the torrents of digital information that we produce each day, “can we identify narratives when they are not yet at the surface? Can we learn about how they come and go?”

As Tuckett spoke, I began to believe in the idea of quants enabling us to digest the world in more rational ways, to become, in a sense, better versions of ourselves. “We are not interested in a world that is completely without excitement or volatility,” said Tuckett, “But we are interested in getting a handle on things before they get out of hand.” The paradox is that in order to become safer, in order to become better informed, we will have to continue to place ever more faith in brains and machines that we only begin to understand. It is always easy to start. The problem is knowing when to stop.

To comment, please email magazineletters@ft.com

Printed from: <http://www.ft.com/cms/s/2/0664cd92-6277-11e1-872e-00144feabdc0.html>

Print a single copy of this article for personal use. Contact us if you wish to print more to distribute to others.

© THE FINANCIAL TIMES LTD 2012 FT and ‘Financial Times’ are trademarks of The Financial Times Ltd.