

From Maps to Apps: the Power of Machine Learning and Artificial Intelligence for Regulators¹

Speech by Stefan Hunt², Beesley Lecture Series on regulatory economics, 19 October 2017

1. Introduction

We live in a digital world. And our actions create ever-increasing stores of data. Be it our posts or likes on social media, our use of trains or planes, or the water or electricity we consume.

Using these vast stores of data, algorithms are transforming our day-to-day lives. The search results we see on Google, the stories we see on Facebook, or the recommendations we see on Spotify. They directly shape our experience of the world around us.

And algorithms also affect what we do not see. They block emails likely to be spam, optimise a delivery driver's route, or flag potentially fraudulent financial transactions.

But commercial companies are not the only beneficiaries of this deluge of data. The public sector can also use data to help us tackle pressing issues.

Police officers, for example, use data to stop bad guys from committing crime.

It turns out that algorithms – using weather, recent crime and other information – can predict the location of crime better than police officers with decades of experience. In 2008, the Los Angeles Police Department and UCLA started working on Predictive Policing. This is not quite 'Minority Report' – for those who have seen the film – no swooping in just before a specific crime occurs. Police, if not already attending a call, return to one of the predicted high crime areas, using their presence to deter crime before it happens. The impact of this, while modest, is nonetheless impressive. A randomised controlled field trial with the LAPD found a 7.4% reduction in crime compared with surrounding areas.³ At the end of 2016, Predictive Policing was being used in roughly 20 of the largest 50

¹ Originally delivered with the title 'Harnessing the Power of Data Science for Regulators'

² Head of Behavioural Economics and Data Science, Financial Conduct Authority. I thank Jan Spiess, Jon Kolstad and Anthony Niblett for their generous comments and suggestions and Jamie Pickering, Darragh Kelly and Raza Ali for their excellent contributions and assistance. I am also grateful to Joe Perkins, Chris Jenkins, Vian Quitaz and remaining FCA colleagues for their thoughts.

³ 'Randomized Controlled Field Trials of Predictive Policing' G. O. Mohler, M. B. Short, Sean Malinowski, Mark Johnson, G. E. Tita, Andrea L. Bertozzi & P. J. Brantingham, Journal of the American Statistical Association Vol. 110, Iss. 512, 2015

US police departments, with 11 more considering it, and predictive policing is being used in several UK forces.⁴

So if you don't see a policeman regularly walking your streets, you should be pleased. You probably live in a low crime area.

And algorithms do not just help us prevent crime. They can also help economists and businesses answer questions about consumer demand. Let's start with something specific: the demand for potato chips, or crisps.

Patrick Bajari, noted industrial organisation economist and now Chief Economist at Amazon, together with academic co-authors wanted to see how much machine learning could help them estimate consumer demand. They chose to focus on predicting which salty snacks consumers buy in a grocery store.⁵ They found that machine learning models predicted demand much better than traditional econometric models, reducing error by over 20%.

So machine learning allows us to estimate consumer demand much better. For Doritos and no doubt gas, telecoms and transport too.

Machine learning can help economics be more driven by data and so by the real-world. And the economics profession is enthusiastically embracing this technological wave. Machine learning is one of the hottest new areas in the discipline.⁶ At the NBER Summer Institute meetings in 2015, over 250 economics professors – many of them senior tenured economists – packed out the four hour professional development lecture on machine learning.⁷ Attendees at conferences, such as the American Economic Association meetings, have been captivated by new papers, many of interest to regulatory economists.

And there's a reason why this topic is so popular right now. We have moved to a world with an abundance of data. We have already seen the behavioural revolution challenge more traditional economic theory, resulting in a Nobel Prize for Economics for leading thinkers like Daniel Kahneman, and last week, Richard Thaler. Like behavioural economics before it, I believe machine learning heralds the next paradigm shift for economics.

⁴ The Guardian, 31st August 2016. Available [here](#).

⁵ See NBER working paper 20955, February 2015

⁶ For introductions to machine learning for economists see Varian (2014), 'Big Data: New Tricks for Econometrics'. Journal of Economic Perspectives and Mullainathan and Spiess (2017), 'Machine Learning: An Applied Econometric Approach', Journal of Economic Perspectives. For a thorough review of machine learning, see Hastie, Tibshirani and Friedman (2009) 'The Elements of Statistical Learning: Data Mining, Inference, and Prediction'

⁷ Lecture delivered by Susan Athey and Guido Imbens.
www.nber.org/econometrics_minicourse_2015/

For those of us who are regulators, machine learning is starting to make an impact on the tools we use to regulate better, for spotting the bad guys, estimating demand, and for many other regulatory problems.

And machine learning has potentially huge implications for the efficiency and effectiveness of regulators.

Today, my aim is to answer three questions.

- How can we best understand the opportunity to use these techniques in regulation?
- How have regulators started to apply these techniques in practice?
- And, what might be the future consequences of regulators using these techniques?

But first, let me tell you a story. When I was nineteen, I travelled with a friend to Morocco. We arrived in Tangiers and, armed only with a Lonely Planet, got straight onto a bus to the city of Meknes. On arrival, we had only a simple map from the guide book that did not provide much detail. We got totally lost. We could not identify the buildings around us or the patch of wasteland nearby, nor even the area were in.

Can you imagine this happening today? I don't think so. All students nowadays are equipped with a smartphone and Google Maps. It takes seconds to sort out exactly where they are and how to get to their destination.

Why am I revealing tales of my misspent youth? Well, the transition from map to app is our topic today. The leap from guidebook map to smartphone app – moving from a two dimensional piece of paper to multi-dimensional layers of personalised information – is parallel to the transition from traditional regulation to using data science for regulation.

But, to understand this, you need to understand the different components of data science. So let's unpack them. What do smartphone apps actually do?

First, smartphone apps describe the world around us, helping us abstract from the detail and *see the wood for the trees*, sometimes quite literally. And apps simplify and distil so much information: Google Maps has layers for traffic, public transport, cycling, satellite or terrain in addition to the standard map and Street View. It has personalised information on locations likely of interest to us.

Second, smartphone apps help us decide where to go and navigate the path through the trees using their suggested routes. Map technology can predict the quickest routes and provide us easy-to-follow options.

Third, with a smartphone app, you can make better judgements. But human judgement is as necessary as ever for making good decisions. We are not talking about self-driving but rather about using apps with humans still in the driving seat.

So, to recap: the move from maps to apps now allows us to see the wood for the trees, navigate our path through the trees and make better decisions, when humans are in control.

How do we apply this technology to regulation?

Firstly, we need to see the wood for the trees. Using data science, we can understand the markets we regulate, the players within them, and the relationships between those players. We can move from a deluge of data to a nuanced overview.

Secondly, we want to navigate our path through the trees. Data science can do a lot of crunching for us and provide us with options. It helps us prioritise, especially when there is an ocean of complicated data. It's like Netflix, giving you a personalised menu based on your past viewing habits, rather than navigating a gargantuan list of options.

Lastly, when using data science for regulation; human judgement is as necessary as ever for making good decisions. Just as you can't yet follow the Satellite Navigation system blindly, we can't rely blindly on machine learning for regulation. But data science is already helping us make better decisions.

We'll explore each of these three areas in more depth. I will then add one further topic and – cutting through the hype – reflect on what the future holds for data science in regulation.

Before beginning the first stage and our early forays into the forest, let's start with an overview of what data science actually is, and how data science fits with regulatory objectives

2. What is data science?

Data science is a general term for extracting information from data. It is an interdisciplinary field that investigates, develops and uses scientific methods, processes, and systems to wrest knowledge from data. This includes traditional estimation and modelling techniques as well as more modern techniques.⁹

Machine Learning is a part of data science. And it's machine learning we are going to focus on in this talk.

Machine learning is a set of analytical tools developed by mathematicians, statisticians and computer scientists since the 1950s.¹⁰ In brief, machine

⁹ The term Data Science started being used more actively in the late 1990s and became very popular following a 2012 article in the Harvard Business Review, 'Data Scientist: The Sexiest Job of the 21st Century'.

¹⁰ Arthur Samuel, an American pioneer in the field of computer gaming and artificial intelligence, coined the term 'machine learning' in 1959 while at IBM.

learning is the ability to learn without being explicitly programmed, instead learning patterns from many examples. Think of these as heuristics that solve problems such as predicting what you are going to watch on TV this evening given your past viewing and demographic data. They are intuitive ways of solving the problem of picking out patterns.

Artificial intelligence (AI) is a broader field that incorporates machine learning and also other techniques such as automated reasoning, i.e. allowing computer programs to reason completely or nearly completely.¹¹ But there is a lot of hype around artificial intelligence, so I will avoid using the term. For the rest of the talk I will mostly use the plainer and more specific term, machine learning.

A computer algorithm is a process or set of rules to be followed in calculations or other problem-solving operations by a computer. Computer algorithms execute solutions to data science, machine learning or artificial intelligence problems.

So to recap: Data science means extracting information from data.

Machine learning is a part of data science, and is the ability to learn without being explicitly programmed.

3. How does data science fit with regulatory objectives?

In the private sector, we see regulated firms using large datasets and machine learning to make profit. For example, Morgan Stanley is supporting its 16,000 financial advisers in the US on what trades to make, Upside Energy is optimising energy storage between self-consumption and providing power, and telecoms companies are targeting their marketing and decreasing churn.¹² As a result these firms increase revenues, reduce costs and improve their bottom lines.

Of course for regulators our aim is different. The mission of the FCA – as laid out in the recent [Mission document](#) – is to reduce harm, the potential for harm or markets not working as well as they could. Or in economic speak, we seek to increase welfare, especially consumer welfare. The actions that the FCA can take include making policy rules and allocating our resources to detect or deter bad behaviour. But it's not easy. The FCA regulates approximately 56,000 firms. These firms have millions of employees and sell to nearly every adult in the UK. Together these firms contribute just over 7% of the national economy.¹³

¹¹ The term artificial intelligence was coined in 1955 by John McCarthy, a mathematics professor at Dartmouth College

¹² For details on Morgan Stanley see Bloomberg, article available [here](#). For details on Upside energy, see their website available [here](#). For details on telecoms companies see an IBM paper, see [here](#).

¹³ House of Commons, briefing paper number 6193, 31 March 2017 'Financial services: contribution to the UK economy'

As regulators, we are – or can be – immersed in micro-data. Every day we get data on the behaviour of consumers, firms, employees and financial traders.

For example, in the market investigation into retail banking, the Competition and Markets Authority (CMA) received monthly data on over 40 variables for 120,000 consumers for two years. A total of over 200 million data points.

In the FCA's credit card market study, we received information on 74 million credit cards over 5 years. A total of more than 100 billion data points.

To detect insider trading, every day the FCA receives details of over 20 million transactions in equity markets. Over a year that's 150 billion data points.

Similarly Ofgem collects energy trading data daily, that's millions of transactions annually.

The challenge is how can we extract the maximum information from this vast data and turn it into useful insights and predictions, so that we can act efficiently and effectively?

Or in other words, how can we, as regulators, best allocate our resource for maximum impact and efficiently find the needles in these haystacks, with limited staff? With the regulatory needles being mis-selling, misleading advertisements, firms colluding with each other on prices or other issues.

A short answer to this question is that much of what regulators do is ultimately about recognising patterns in the data we have. Machine learning helps us find these patterns efficiently.

And machine learning is now widely available. Everybody can use the same high quality analytics by downloading free software. And with access to the cloud we have immense number crunching capacity at low cost.

4. Seeing the wood for the trees: unsupervised machine learning

To begin to answer the question of how to use data science for regulation in more depth, I'll move to the first of the four main stages of our journey, seeing the wood for the trees.

In the complex world of regulated markets, how can we describe our data and be sure we understand the most relevant patterns and trends?

This takes us to the realm of what is called unsupervised machine learning.

With unsupervised machine learning there is no outcome that trains, or supervises, the machine when it learns. Here, we are not aiming to predict an outcome, but rather to describe entities and their relationships.

A trivial, but amusing, example comes from Google's X lab. A team of computer scientists built a neural network of 16,000 computer processors with one billion connections, and let it browse YouTube. Without any labelled dataset to train the network on, it started to categorise videos according to similarity. What did it discover?

Cats. Lots of cats. It recognised by itself that a major category of video that humans like to watch is cat videos. Most data that we collect does not have labels on it that say, for example, this is a cat or this is a dog or this is a human. Or, for regulators, that this consumer is vulnerable or disengaged. Or this company's is colluding with its competitors. So unsupervised learning can be particularly important in creating insight from the oceans of unlabelled data that surround us.

There are many different types of unsupervised learning. Today we will look at just two.

First, we can see how we can form groups of market participants that are similar to each other – be they firms, consumers or traders. This might be useful, for example, to understand which consumers are actively engaged and which are not.

Second, we can see how to identify what drives the behaviour of the different market participants we observe. This can be useful for identifying drivers of behaviour that worry us, such as the desire to commit fraud.

One important application of this technique is understanding the underlying topics that drive the language we see in documents or in conversations. So I will also speak about natural language processing more generally.

So, first, how can we form groups of market participants that are similar to each other?


We can do this using clustering algorithms, or cluster analysis. These algorithms are a useful set of tools for exploring our data. They create groups using just the available data and no information on what group each participant is in. The really exciting thing about clustering algorithms is that we can use them even when we have limited prior knowledge to guide us. That means we can minimise bias because the algorithms don't rely on human judgement.

One example comes from our recent policy work. The FCA is currently considering rules to help people manage their current accounts. This may include having alerts such as low balance warnings. And we are considering overdraft remedies more broadly.¹⁴

¹⁴ [FCA Feedback Statement FS17/2: High-cost credit and review of the high-cost short-term credit price cap](#)

Bank overdraft charges are concentrated among a small group of individuals. We wanted to understand what was really happening for these individuals and what might be the causes of their situation. To do this, we turned to clustering algorithms.

Figure 1: Clustering of overdraft users



Source: FCA analysis

These figures show a sample from a dataset of 250,000 representative current accounts from each of the big six providers. We see all transactions in 2015 to 2016 for all 1.5 million accounts. Roughly a quarter of people use unarranged overdrafts (shown on the left) and roughly a quarter use arranged overdrafts (shown on the right). These charts show overdrafters only. Each chart shows a representative sample of 5,000 accounts. Each row represents one account and the row is dark when the account is overdrawn. You can see from the chart four clusters of consumers, starting from the bottom cluster: occasional short spell overdraft users, occasional long spell users, more intense user and persistent users.

Now these clusters you see were produced by the clustering algorithm. We told it how many groups to find but did not provide any information on the membership of each group.

The charts show that overdraft usage is highly concentrated. And you can see markedly different patterns of behaviour across consumers. In both charts the cluster at the top shows those consumers with more sustained overdrafts.

The data here is from one calendar year. At the end of each month, on payday, persistent overdraft users tend to go into positive balance. These are the monthly lines you see. Note that there is no clear monthly pattern for occasional users. So the chart helps us to identify possible causes of overdraft use for different people.

You can also see differences between arranged and unarranged use of overdrafts. The persistent user group is much smaller for unarranged overdrafts than for arranged overdrafts. Also persistent unarranged overdraft users show intense overdrafting for approximately six months only.

This raises an important question: what brings them into and out of this pattern? Now, contrast this with those users with persistent arranged overdrafts – they are almost constantly in overdraft, and even in the second cluster they tend to overdraft regularly every month. The distinction between sudden sustained spells (unarranged) and regular use (arranged) seems to be meaningful – it changed how I was thinking about the market failures present here – and so relevant to policy. And we are examining these patterns and the market failures further.

Using clustering algorithms, we were quickly able to analyse large amounts of data and create a deep and useful visual representation to show us what was going on in overdrafting behaviour for different groups of consumers.¹⁵

And, these techniques are not only useful for segmenting consumers. George Michailidis and co-authors clustered financial traders based on Commodity Futures Trading Commission data from a derivatives exchange. Based on buying or selling behaviour and trading intensity, they put traders into five categories: high frequency traders, market makers, opportunistic traders, fundamental traders, and small traders.¹⁶ And they found that these categories were stable over time.

Using this analysis we can understand traders better. Supervisors can categorise the business models of firms in advance of regulatory visits. And our Market Oversight team can clarify the normal behaviour of traders and detect deviations that might flag insider trading.

Clustering algorithms help us understand how market participants differ, be they mobile telephone consumers, gas traders or small business consumers of water. Clustering algorithms are a core tool in the machine learning toolkit. They are often one of the first tools we use to explore a new dataset.

¹⁵ A recent publication from the US Consumer Financial Protection Bureau also used clustering techniques. See Data Point: Frequent Overdrafters, Consumer Financial Protection Bureau, August 2017. Available [here](#).

¹⁶ Mankad, S., Michailidis, G., & Kirilenko, A. (2013). Discovering the ecosystem of an electronic financial market with a dynamic machine-learning method. *Algorithmic Finance*, 2(2), 151-165. Available [here](#).

So, secondly, how can we identify the underlying drivers of behaviour of different market participants?

We can do this using a class of techniques called topic models. These models allow us to discover the abstract topics or motivations that drive the behaviour we see. They were originally developed to discover hidden semantic structures in a text body, but have wider application. I will start with a non-regulatory example.

Brad Love, a psychologist at UCL, and co-authors analysed the shopping baskets of millions of consumers at a major UK supermarket chain. The academics wanted to find the underlying motivations driving shopping behaviour, using only data on the contents of shopping baskets. They found about 30 different underlying motivations. In terms of general trends, they found that some people's shopping could be described as low-cost, or alternatively more upmarket. But other people's shopping could be described more specifically, such as 'produce for making a stir-fry from scratch' or 'pre-Christmas shopping'. The mix of motivations in a given shopper's basket comes direct from the data.

Knowing what a shopper is trying to do – their mission – makes it easier to know what else they might want to buy and therefore what to market to them. And knowing the mix of motivations of shoppers at a particular store can inform the layout of that particular store, placing thematically related items together in the aisle.¹⁷

We can use the same technique to find the underlying drivers of the behaviour of the participants in the markets we regulate. For example, we might want to uncover motives such as committing fraud.

We can also use this technique to find underlying topics in language. It is what it was originally designed for. For example some of our topics this evening are regulation or machine learning or economics. We can extract topics for any media – be it phonecalls to a regulator's call centre – say Ofgem's reserve unit – or social media posts. We can use this information in a multitude of ways.

Topic modelling is one unsupervised technique among many supervised and unsupervised techniques that make up text-mining, or natural language processing. I expect that every one of our organisations will use some text-mining in its work within a couple of years, if not already.

These techniques have begun to be used in regulation, e.g. by the Securities and Exchange Commission, to detect accounting fraud.¹⁸ And they have also been

¹⁷ There are many examples of the use of unsupervised learning in the supermarket space, including the first well known use of it in the retail space: Walmart using frequent pattern mining to identify products frequently sold together and placing them together back in the 1990s.

¹⁸ Gerard Hoberg and Craig Lewis performed sentiment analysis to assess text with a negative tone or tone of obfuscation using corporate filings with the Securities and

used elsewhere in the public sector. The Serious Fraud Office was able to expose large-scale bribery and corruption at Rolls-Royce. They used machine learning to sift through 30 million documents, processing up to 600,000 every day.¹⁹ The robot could assign topics to, index and summarise documents much as a human investigator could do, but much faster. If you face having to sift through enormous piles of documents, you could use similar techniques.

Topics models can help us take reams of data and model the underlying topics or motivations that drive behaviour. Text mining is one important application of these techniques.

As you can imagine, there are many other applications of unsupervised learning. Using unsupervised learning, you can visualise the relationships in your data, such as traders in a financial system, using network analysis or graph models.²⁰ Or you might do this for distributed electricity networks or other networks. You can detect anomalies that might flag fraudulent transactions.²¹ Or signal that a consumer is transitioning into a vulnerable state. You can reduce oceans of data down to the bare necessities using dimension reduction, allowing you to navigate complex data and protect privacy, by removing direct information on people's behaviour.²²

In summary, unsupervised learning techniques help us see the wood for the trees. We regulators can extract more information from data, especially large datasets. While we often use unsupervised techniques as a precursor to predictive analysis – which is where we are heading next – the insights created can also be directly useful in their own right.

5. Navigating our path through the trees: supervised machine learning

We are now moving from using machine learning to describe, to using it to predict – helping us make decisions about whether to do A or do B. This is called

Exchange Commission. They found that fraudulent managers grandstand good performance and disclose fewer details explaining the sources of the firm's performance. The abnormal text predicts fraud in out-of-sample tests. So the predictions of the model can be used for supervision. Gerard Hoberg and Craig Lewis (2015) Do Fraudulent Firms Produce Abnormal Disclosure? Available [here](#).

¹⁹ 'SFO expected to promote Ravn's crime-solving AI robot', Financial Times, February 13, 2017. Available [here](#).

²⁰ For example, see 'Big Data jigsaws for Central Banks – the impact of the Swiss franc de-pegging' at the Bank of England's Bank Underground website. Available [here](#)

²¹ See Bolton RJ, Hand DJ. Unsupervised profiling methods for fraud detection. Proceedings of the Conference on Credit Scoring and Credit Control; 2001; Edinburgh, UK.

²² This chart shows a publicly available dataset of credit card fraud. See www.kaggle.com/cherzy/visualization-on-a-2d-map-with-t-sne

predictive analytics or supervised machine learning. It is called supervised machine learning because we train, or supervise our algorithms with knowledge of whether our predictions are right or not. It is supervised machine learning, mostly, that has had transformational results and received press and TV coverage in abundance.

In this section, first, we'll look at why you should care about prediction in regulation. And we'll see why supervised machine learning promises to be the answer. Secondly, we'll get a flavour of the mechanics of these methods and why these techniques work. Importantly, I want you to understand how they are so different to what we are used to from econometrics. When you understand the mechanics you'll see why the techniques are so powerful. We'll explore how the methods can go wrong and the crucial role of economists in making sure that they don't. Thirdly, I will talk us through examples from the FCA where we are starting to apply the techniques in practice.

Why focus on prediction? And why use supervised machine learning?

You've heard about predictive policing, and how Google Maps or Citymapper provide route suggestions. This is just the beginning. When cars drive themselves, they – mostly – use supervised machine learning to detect what the objects around them are: lines in the road, red traffic lights and pedestrians. When the iPhone X learns to recognise your face, it is doing so based on a training dataset of a billion images. So that's supervised learning too.

The accuracy of prediction has just got better and better. But the GoogleNet, convolutional neural network has little problem, for example, in figuring out which of these [images](#) are Chihuahuas and which are muffins. And it is not that easy.

Or which of these is Chad Smith of the Red Hot Chili Peppers, and which of these (see [here](#)) is Will Ferrell pretending to be Chad Smith, and which is Jimmy Fallon. Note, the machine is identifying Fallon from the side on.

Many regulatory problems are, in essence, prediction problems.

As a regulator, we often have to make choices about how to allocate scarce resource. We need to decide which firms or employees to get more information from, and which to investigate more thoroughly. And, at the FCA, we are often monitoring tens of thousands of firms or financial professionals. These choices about prioritisation could be about providing permission to trade to a financial adviser, supervising a hedge fund or checking that an insurance firm has not engaged in collusive behaviour. The CMA, Ofgem, Ofcom, Ofwat, ORR, CAA and other regulators make similar prioritisation decisions.

Even if a regulator does not have on-going feeds of data from firms, it needs to use the information it can get to prioritise. Be it the detection of cartels or bid

rigging or other collusive behaviour, or using Twitter and other social media to get early warning signals about consumer protection problems.²³

When we choose where to focus, we are saying – implicitly or otherwise – that what we focus on is more likely to have an issue. There is a greater chance of a regulatory problem. We are making a prediction.²⁴

Moreover, detecting existing issues, when we do not have full information, is also a prediction problem. Detecting an adviser mis-selling products, detecting a bad product, or detecting insider trading, spoofing or layering on energy trading platforms all require prediction. This is the same sense in which a self-driving car is predicting – hopefully with perfect accuracy – that the object in front is a pedestrian. So prediction is a general problem of what to do when we do not have enough information. It isn't necessarily about the future.

As regulators, we often need to decide where to investigate more using all the data that we already have. Consider an FCA team tasked with detecting and deterring firms from rule breaches. The team has access to millions of data points from disparate sources. These include multiple datasets on consumer complaints from the FCA, Financial Ombudsman Services and others, balance sheets and income statements, historical firm and individual issues, intelligence data, publicly available information, and also information on the firms' staff, including their training and employment patterns.

These amount to hundreds of variables.

So, how should the team best combine all this information?

In supervised machine learning the algorithm is learning from the data. It sifts through these hundreds of variables figuring out which are useful and which are not, perhaps including subtle combinations of variables.²⁵

We can formalise the FCA team's task as a supervised learning problem: we teach an algorithm to learn from past breaches of regulations and predict new breaches. These predictions then go to the FCA team to help them decide what to do next.

And we can similarly formalise many other regulatory problems as supervised learning problems, for example, again, cartel detection.

²³ On extracting early warnings about consumer problems from Twitter see work from NAO, around p.40 and the appendix: <https://www.nao.org.uk/wp-content/uploads/2015/06/Putting-things-right.pdf>

²⁴ Though for the purposes of deterrence of course we want some probability of focusing on every single entity, including low risk firms. So this does not mean that every firm or employee that the FCA focuses on is high risk.

²⁵ The algorithm might aim for accuracy or for precision – two different statistical concepts – a combination of the two or something else.

Assessments of where to get more information drive a large amount of regulatory activity. Many regulatory problems are, in essence, prediction problems. Supervised machine learning solves prediction problems. And it does so by trawling through all the data we have by itself, looking for patterns.

What is it that is so different to econometrics?

For prediction problems, contrary to most traditional economic problems, we are normally just interested in maximising our predictive power. To prevent bad events, such as product mis-selling, or catch illicit behaviour we want to know where it is likely to occur.

There are many different algorithms – you may have heard of Random Forests or Deep Learning. And they find patterns in different ways. But – for those interested in the statistics a little – at their core these techniques all benefit from the same basic trick: creating a hold-out sample of data – the test set – and splitting it from the training sample. By doing this, we can exploit all kinds of funky and crazy heuristics to spot patterns in the training sample. We can then use the test set to see how well the algorithm actually predicts.

A second statistical trick is crucial to achieving great prediction. All these algorithms run the risk of overfitting the training data. The model can get overly complicated and predict brilliantly well in the training sample.²⁶ But it predicts terribly out of sample because it does not represent the true structure and relationships that exist in the world. To avoid this, we need to ‘regularise’ complexity. An important way of doing this is to create extra test sets within the training sample, allowing us to choose the right level of complexity before moving to the test set.²⁷

These two tricks - the test set and regularisation – allow us to find true and meaningful, rather than spurious, correlations between variables. Armed with these tricks, computer scientists invented a myriad of pattern-hunting heuristics.


²⁶ Overfitting is the problem that as you put more and more parameters into your model when working with the training data you are able to explain more and more of your data. But, after a certain point, this fit is completely spurious. This is similar to say if we want to predict which horses are the fastest at the racing track. Wouldn't that be nice. Let's say that we first consider whether the horse likes the turf soft or firm, or the age of the horse, or the horse's previous achievements. These can all be useful predictors. But let's imagine we throw in some other variables – whether the horse seems to like Wednesdays, whether its name begins with the letter A etc. – the mechanics of prediction with a set amount of data means that more variables can only ever help us get a better fit and prediction. But the problem with these variables is that they are just noise. Any fit we get from them is definitely spurious. We can't really predict which horses are going to win their races, unfortunately. In fact it turns out that if we add lots of variables that seem reasonable – are not clearly garbage – we get the same problem. Too many variables – too much complication in our model – creates spurious correlation.

²⁷ The technique of creating extra test sets in the training sample is called cross-validation and allows us to understand the properties of our model without using the hold-out sample.

Describing a couple of algorithms will give you a sense of what all the fuss is about, and importantly, what the issues are.

Let's go back to the FCA overdrafts data from earlier. Let's imagine that we want to predict who will incur an unarranged overdraft charge next year, based only on a consumer's age and overdraft limit.

Figure 2: Prediction – econometrics versus data science (1)



Source: FCA analysis

The x-axis shows customer unarranged overdraft limits and the y-axis shows consumers' age. As you might be able to see from the figure, consumers who incur unarranged charges over a year (red dots) tend to be younger than those who don't (blue dots).

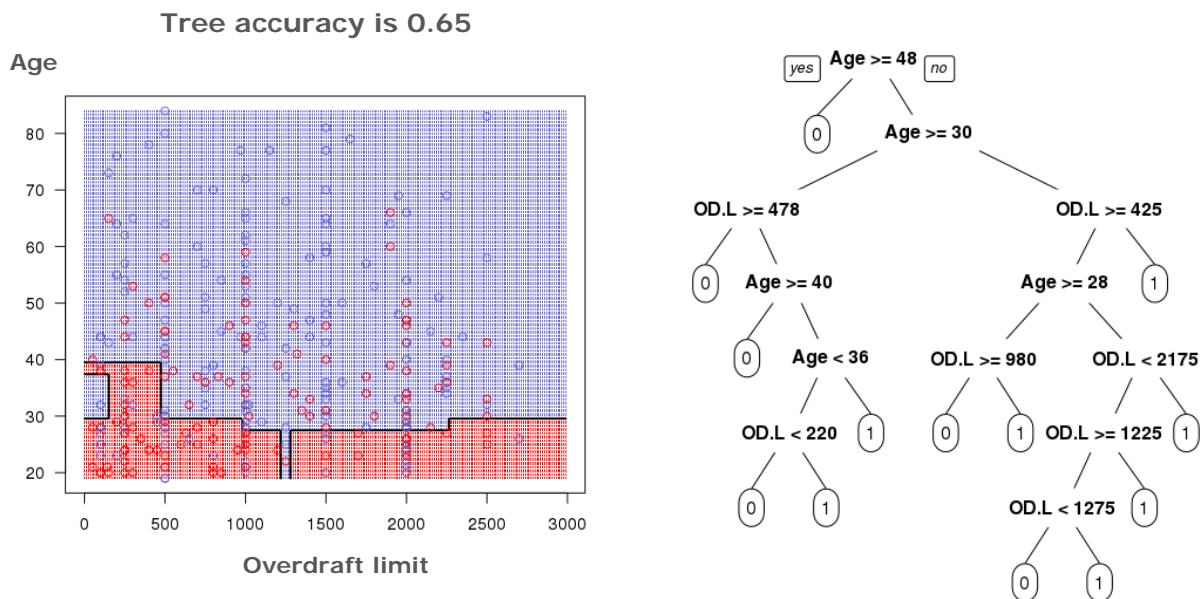
The left hand figure shows how the well-known logit or logistic regression – that you will all be familiar with – predicts. The model draws a line that separates the predictive space as accurately as it can. It predicts that consumers who tend to be younger and have lower overdraft limits are more likely to incur unarranged overdraft charges.

The figure on the right hand side shows how a decision tree – one of the simplest supervised learning algorithms - would separate the predictive space. Note that the model is more flexible, drawing non-linear boundaries to separate the space - and it is also more accurate, boosting accuracy by a couple of percentage points.

The decision tree works by looking at all of the variables available to predict - here age and overdraft limit, but it could be hundreds or thousands of variables. Across all of the points for each of these variables – i.e. age of 20, 25, 30, 35 or

overdraft limit of £1,000, £1,500, £2,000 etc. – the algorithm finds the one point where splitting the data into two allows the model to predict overdrafting as best as possible. It does this by asking a simple yes/no question, in this example ‘is the person younger than 48 or not?’. When this point has been found the algorithm splits the data. It does this again and again until it stops materially improving model performance as measured by regularisation. The tree in this chart represents the figure that I just showed you.

Figure 3: Prediction – econometrics versus data science (2)



Source: FCA analysis

There are three points worth noting. First, the algorithm is crunching huge amounts of data. It is systematically and efficiently sifting through it to identify the signals from all the noise. Second, the algorithm is non-linear. It allows for potentially very complex functional forms – see how it uses age repeatedly – and interactions between variables. Third, it deals quickly and automatically with a whole range of statistical issues that we worry about in econometrics.²⁸

Decision trees have some really strong points, in particular they are easy to understand – they are ‘white boxes’. But there is a problem. Decision trees can overfit on the training sample and not perform so well in the test set, particularly when working with high-dimensional data.

²⁸ For those who have ever spent days worrying about running regressions with too many correlated variables, the assumptions on your regression residuals or whether to use a logarithmic transform, this method deals with all of these issues.

The good news is that a little over twenty years ago, a data scientist from AT&T developed a solution. She grew lots of decision trees! What is now called a random forest.

The algorithm creates hundreds or thousands of somewhat different trees, and the predictions are combined to give one overall prediction. This procedure makes the overall prediction much less sensitive to specific variables.

This slide shows how Random Forests can be used in our overdrafts example. While it looks unusual, it adds another three percentage points in accuracy. And can add far more.

Figure 4: Prediction – econometrics versus data science (3)



Source: FCA analysis

Random Forests come with a drawback: the model is complex, combining hundreds or more non-linear models, maybe tens of thousands of yes/no questions. So you can see why some machine learning can be described as a 'black box'.

But these methods perform really well and are being adopted widely.

I was talking to a senior executive at a credit company a few weeks ago. They have moved not only most of their credit scoring but also their marketing analysis to a type of tree regression. Better prediction is ultimately valuable, both for business and for an efficiency-seeking public sector.

Regression trees are just one class of supervised learning.²⁹ There are many more. One class worth looking at is neural networks.³⁰ I am sure you have heard about them, whether it be Deep Learning techniques and their use in self-driving cars or particular company's models, such as Google's Deep Mind or IBM's Watson. These models are based loosely on the workings of the brain. Deep learning refers to having many layers of artificial neurons. This allows raw data to be fed into the model rather than the data scientist having to decide how to transform it. The machine figures out the transformations. This is the technique that has been responsible for real-time translation of menus using Google Translate or Android's voice recognition.³¹ But understanding how the network is predicting is hard. It is based on the weights of many neurons and the strength of many connections, which is not at all easy to interpret.

Albeit that this kind of model is completely different to Random Forests, it and other successful supervised learning techniques have similar differences to econometrics. They can crunch huge amounts of data, sniff out complex patterns and deal with statistical issues. But some of the best performing algorithms do so while often becoming black boxes, opaque and hard to interpret.

How are regulators starting to apply these techniques?

Where are regulators at in actually using these techniques?

Well, it is early days, but we have made some good progress. I will discuss some FCA examples.

The first example comes from investment management, i.e. asset and wealth managers. These firms play a critical role in our society through allocating their capital. But they can cause a number of problems when things go wrong. One fairly frequent example is a breach of an investment mandate, where a portfolio manager makes investments that are out of scope of what he or she has promised in the fund's literature. This might make the fund more risky than investors had expected, or less risky too, of course. Another type of risk is suspected incidents of insider trading.

The question we set ourselves was: can we develop a categorisation of firms that differentiates between firms likely to generate notable breaches, risk events, and those less likely to do so?³²

²⁹ If you want an introduction to the array of different algorithms and the potential and pitfalls of supervised learning, I recommend the Journal of Economic Perspectives article by Mullainathan and Spiess from earlier this year. See footnote 5

³⁰ Neural nets can also be used for unsupervised machine learning


³¹ Another example comes from the Fake News Challenge – identifying fake news using a labelled dataset – where the winners used a deep neural network coupled with a more conventional neural network to develop the most discriminative model. See Sean Baird, 'Talos Targets Disinformation with Fake News Challenge Victory' Available [here](#).

³² Our challenge here is a strong selection bias. It's too easy to look at firms that the FCA has focused on historically. This sways any model, instead of identifying general

As I mentioned earlier, the FCA has a very large number of internal and external datasets. We first wrote code that pulled these many sources of data together from different systems and cleaned it. Quite some task. And we can now compile all these sources again at the touch of a button!

After initially experimenting with simpler models we used two algorithms to create a model, Random Forests and another algorithm called MetaCost. We trained the model on 2014 and 2015 data and tested it on completely unseen data from 2016.

Figure 5: Investment management model results



Source: FCA analysis

The result of the model is shown in the figure here, with firms ordered by risk on the x-axis and the percentage of firms that generated risk events on the y-axis. We found a small group of firms in the highest risk category, and these had a 54% chance of having risk events, compared to a large group of firms in the lowest category that had just an 8% chance. We used a clustering algorithm to group firms into risk categories.

And we can dig into the model to understand what drives the results. A wide variety of different categories of variables can flag high risk firms, including the

attributes that makes a firm more risky. We dealt with this in two ways. We generated variables less likely to simply identify previously-studied firms. And we also added a cost-function to the evaluation of our model that penalises the model from simply reflecting historic behaviour. Domingos, P. (1999). MetaCost: a general method for making classifiers cost-sensitive. KDD'99, 155-164. Available [here](#).

governance of firms, staff experience or remuneration. We can measure the relative importance of each.

So what did we achieve? The FCA supervision team now has a model that provides a single risk score for each entity. The team can now, layering on its own expert knowledge, decide scientifically which firms are in different risk buckets. And they have new, specific and quantitative, insights about what drives risk. Very importantly, the model potentially allows the supervisors to be more proactive to prevent harms from occurring, rather than resolving problems that have already arisen.

Now we have not put this model into the field yet. But with other models we have. And we have indeed been able to move forward with completely new, proactive supervisory cases.

Using the model – in this and similar projects – we can iterate and improve it. Our latest version suggests we can identify groups of even lower risk firms – with only 3 or 4% chance of risk events.

As we gather rich, new data from firms to check our predictions and take action – e.g. when we request detail on each risk event – we learn much better about the environment we are working in, again quantitatively. And we learn about how we can change what data we collect in the future, to prevent harm more effectively.

A second FCA example is important because of what it illustrates might be possible. The idea is to create an algorithm that can scan new advertising and flag whether it is likely to be misleading. We have scoped out the viability of a project. And we think it is feasible. For example, we analysed samples of promotions for features that a machine learning algorithm could identify and interpret. We found that a significant number of promotions had issues with the risk warning, or lack of it. Algorithms should be able to assess whether a risk warning is missing, is insufficiently prominent or is inadequate.

And more subtle forms of misleading financial promotions might be identified using deep learning based methods.

One reason why the financial promotions project is exciting is that we could run the algorithm over all adverts, rather than just review a sample. The algorithm would flag potential issues for the supervisors to look at. It would change the work of the supervision team. They would spend more time on adverts that require their expert judgement and less time reviewing adverts without issues.

It is clear that in many spheres of life machine learning is having notable successes. These FCA examples illustrate the promise of supervised machine learning for regulation. These models could be built for the many other regulatory problems we discussed earlier, detecting insider trading, cartels or bid rigging, or predicting consumer protection issues from social media.

And it is by no means just the FCA that is interested in using machine learning. The Securities and Exchange Commission³³ and Financial Regulatory Authority in the US, the Bank of England, the Monetary Authority of Singapore, and more, are all investing in or trialling machine learning.

To summarise, supervised learning techniques help us navigate our path through the trees ahead and make choices. Many regulatory problems are, in essence, prediction problems. Supervised learning can provide us with information to help us make prioritisation decisions, efficiently using all the data that we have. These techniques work in a completely different way to those we know from econometrics. They are often highly non-linear models that sniff out every last bit of predictive pattern. But their complexity can leave them as opaque black boxes, at least when it comes to individual predictions.

The examples illustrate the promise. They also provide us with a sense of the pivotal role of humans in the use of this technology as well as provide our first glimpses as to what we can realistically expect in terms of impact.

So these are the two topics left to discuss – human, and what we can expect. I turn now to consider the first of these two topics.

6. Humans in the driving seat

Algorithms can predict leaky pipes in city water network or – in a celebrated example – even select the most perfectly-formed cucumbers for discerning consumers in Japan. These and other successes lead to manual processes being streamlined and workplaces changed.

A natural question is what is the role of humans in regulation as we move to using algorithms?

Humans, as we well know from a slew of behavioural economic examples, have their own biases and inconsistencies. A recent paper by legal scholars Alarie, Niblett and Yoon cited evidence from the use of machine learning for legal disputes and argued that algorithms can streamline operations and provide fast, accurate and consistent judgements with reduced error.³⁴

Now I don't disagree that machines can outperform humans on some tasks. But in regulation, algorithms are far away from making any decisions. Humans are

³³ They went public on their first model – one that used natural language processing to detect accounting fraud – back in early 2013. They now have a whole host of new initiatives analysing previously impenetrable information sets, such as freeform text. These initiatives leverage machine learning to predict bad behaviour, particularly in identifying fraud and misconduct. One of their tools, their Corporate Issuer Risk Assessment program, provides more than 230 custom metrics for SEC staff to use.

³⁴ See Alarie, Niblett and Yoon (2016) Regulation by Machine, 30th Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain. Available [here](#).

needed at every stage of the process. When we are talking about using algorithms in regulation, we are talking 'decision prosthetics', i.e. Citymapper or Google Maps suggesting routes and providing information.³⁵ And us humans still need to take all this information and make a choice between options. We are not talking Robocop or Minority Report. We don't want to be following the sat nav and ending up at a building site. Let me explain why.

First, at the FCA, even if we were to use algorithms to help decide where to focus, we have established procedures, run by humans, for determining whether a firm or person is at fault. Machines cannot substitute for making such an assessment.

Second, human judgement is needed when using models. When deciding which firms to visit, say, predictions may help, but supervisors can also use additional information that is not in the model. They may have important information from the latest visit to the executive team, or knowledge of recent changes in the market or to our regulations.

That said, machine learning models can be improved if humans make better predictions in some areas. For example the credit start-up Aire employs a team of experienced bank underwriters to play against their credit scoring system and help it behave more like an underwriter, taking into account complex personal circumstances.³⁶

Third, creating the model itself involves many choices that require intuition from seasoned modellers and domain experts. There is a lot of discretion. From creating and selecting a subset of relevant variables to enter into the model, to choosing the best outcome variable, to algorithm selection, to regularisation, to comprehending models and making them as interpretable as possible – turning black boxes into white boxes – analysts have to make decisions. And compared to our standard econometric tools, less is known about which algorithm works best in which situation.

This highlights the importance of analyst skill. When making modelling decisions we do not just throw all the variables in and press return.³⁷ For example, knowing which external datasets likely have a strong predictive signal and should be included could be key to a project's success.

There is much art in this science. And in my experience there are more degrees of freedom than with econometrics. The data science way of approaching problems is much more like the tinkering, adjustment and finding-a-way-to-make-things-work of engineering, than the modelling of the econometrician.

³⁵ See 'New Vistas in Risk Profiling' by Greg Davies, published by CFA Institute Research Foundation. Summary available [here](#).

³⁶ See article [here](#)

³⁷ See [this article](#) on learning from winners on Kaggle, a platform for predictive modelling and analytics competitions.

And using machine learning models implies different ways of working for operational staff. They need to use more analytical outputs when making decisions. In some ways this might make decision-making easier. But in other ways it might make decision-making harder – think about trying to interpret multi-dimensional charts. Either way it implies change. In general, in regulation, I think we may find that the quantity of human input may not differ, but the quality, the level of analytical expertise, may need to increase.

Interestingly, for more senior decision-makers, there is also skill needed in the commissioning of these projects by management.

So, humans are crucial and in control, but is there any specific role for us economists? Well, there are a number of reasons that economists can play a vital role in using machine learning.

The first reason is that, econometric techniques can provide solutions to problems such as dealing with biases in our data, e.g. because past data may reflect biased human judgements. For example, when judges are making decisions about whether to grant bail to defendants or whether to keep them in jail, if racial minorities are never granted bail, then we cannot know what their likelihood of reoffending is, even if they actually have a lower likelihood than others. More generally selection in our data creates distortions.

A recent paper introduced a solution to this selection bias. A mixed, superstar team of computer scientists and economists investigated the bail decisions of judges in the US. The clever use of econometric technique – instrumental variables³⁸ - combined with machine learning allowed for a better model and a careful evaluation of the potential impact of using machines to decide on bail. A policy simulation shows crime by defendants let out on bail could be reduced by 24.8% with no change in jailing rates.³⁹

A second reason is that, in general, economists have broadly the right skillset for machine learning. It is easy for empirical economists to understand the techniques and code in R or Python. And economic theory can be powerful in choosing which variables to create. Also, trading off false negative and false positives is, in the end, a cost-benefit analysis.

³⁸ Defendants are assigned to judges in a quasi-random way and some judges are more lenient or more strict than others. So we can use the fixed effects of individual judges as instruments.

³⁹ They also had these thoughts on the fairness of algorithms: 'There are some judges who are very lenient. There are some judges who are very harsh. Why should someone's life and liberty depend on the random toss of coin of which judge they get? How is that fair? You might have one person up for bail that would be treated very differently by two different judges. The algorithm won't do that. The algorithm treats like cases alike.' The authors go on to note that algorithms are far less racially biased than human judges in the US, who tend to treat black defendants more harshly.

A third reason is that, economists now understand human biases quite well and can help to make sure any tools ameliorate these biases.

I should also add that, while the topic of this talk is firmly on regulators using algorithms themselves, developing the capacity to use algorithms will help us understand the implications of firms using algorithms. As companies invest further and regulatory issues arise – e.g. the FCA recently looked into the use of big data by general insurance companies – regulators will need to develop their knowledge of techniques.

To summarise, while machines are likely to perform much better than us on some tasks, we very much need humans in charge. At a higher level, the most important human judgement is in discerning when to use machine learning and when not to use it. That is, when do the expected benefits – the impact on efficiency – outweigh the expected costs, especially the cost of getting the right data. This takes me to a few final thoughts on when we might apply these methods and what kind of impact we should expect.

7. Expectations: keeping our feet on the ground

We now reach the fourth and final stage of the journey.

While machine learning has had some ground-breaking successes – e.g. I was recently blown away watching Google Translate translate a foreign menu into English in real time – there is much hype. What can we reasonably expect from machine learning in the context of regulation?

In principle a large amount of what regulators do could be framed as information problems, either description or prediction problems. To solve problems with machine learning we need enough data points to use statistics. Any set of tasks that are fairly repetitive (and so a large number of data points) and have relatively clear defined outcomes (so we can train the data) should be amenable to supervised machine learning.

Certainly at the FCA, it seems we have many tasks – across supervision, authorisations, market oversight, enforcement and competition – that, at face value, could benefit from data science.

How big can we expect the gains from data science to be? Policing and the criminal system might provide some guide. In the predictive policing example we saw a drop in crime of 7.4% in Los Angeles. And compared to expert judges, according to the model produced by the superstar academics, machines could reduce crime by defendants let out on bail by almost 25%.

I would argue that FCA supervisors – at least – are more like judges, making decisions about individual cases with lots of data. So for the FCA, for deterring or

detecting bad conduct, perhaps we might see closer to the higher level of impact, 25%.

That said, machine learning might have greater impact still, where we can apply it. First, in some cases, such as detecting misleading financial promotions, algorithms allow us to review the whole market – e.g. all advertising and marketing – rather than just a sample. Second, much regulation has not typically been framed as a prediction problem and may not have been addressed particularly quantitatively. So just framing regulation as a prediction problem, even when using simple prediction techniques such as logistic regression, could be powerful and have considerable impact.

Well those are the potential efficiency gains, the benefits. What about the costs?

There are several blockages to making quick progress. First among these is data quality. Often the right data is not available. Or we might need to turn unstructured data, e.g. adverts or tweets, into an analysable format. Estimates from data scientists suggest 80% of a typical project is getting datasets, cleaning them and combining them.⁴⁰ Information that sounds in principle very useful may be very costly indeed to get, completely inaccessible, or the burden on firms may be too large.

In addition regulators need to build up the right skills to use these methods, when these skills are in high demand.

Obviously there needs to be enough of a benefit – the improvement in performance – to justify the costs. That is we need a net benefit.

The answer is clearly to focus on the best cases and grow and learn. We need to pick exemplar projects rather than apply these tools to every issue we tackle. And we need to use methods from start-ups, like the use of agile processes to ensure that we are learning fast.

At the FCA, we are, I believe, ahead of the game. We are working through prototype projects and learning. Moreover these models will get much more powerful as we learn, understand how we can improve and iterate, which takes time. The true potential will not reveal itself for some time yet.

In sum, we do not have enough information to judge the impact of machine learning in regulation. Crime gives us a couple of yardsticks and I argued that a 25% increase in efficiency seems plausible. But this does not account for areas where we could see great impacts, nor potentially large costs in getting usable data. The truth is, nobody yet knows. And we are only going to find out as we get experience and see what happens in practice.

⁴⁰ CrowdFlower, cited in Forbes. Article [here](#).

8. Conclusion

At the beginning of this talk, I told you how I got lost in Morocco at the age of nineteen. It turns out my friend and I were just fine. We were befriended by some local teenagers and got the information we needed.

But it is not always possible to meet locals and find our way. I have argued that machine learning can substitute for a friendly local and in fact, can go much beyond, providing insights that even locals don't know.

First, we want to understand where we are and what the world is like around us – and see the wood for the trees. Much regulation is ultimately about recognising patterns. We can use unsupervised learning to dig deep into data and see these patterns. We saw many applications, from clustering derivative traders, to identifying underlying motivations, to sifting through mounds of documents for evidence.

Second, we want to get to our destination, and navigate a path through the trees ahead. Many regulatory problems are, in essence, prediction problems. Supervised learning is designed to solve these problems. The sophistication of these methods and the great successes in many areas of human endeavour show their great potential. And we have seen early successes in regulatory applications, especially in supervising large populations of small firms.

But, third, we must not use these methods blindly. Just as with using a navigation tool, like Citymapper, we have to take responsibility for making decisions. Humans need to be involved in making all the major judgements, in creating the models, and in using them. And, taking the earlier analogy further, combining the knowledge of locals with a Google Maps app is best of all.

And, fourth, crime provides some benchmarks, which suggest we might obtain efficiencies of 25% or so in regulation. Personally I am optimistic that we can get impact larger than that. But we are only going to find out over time.

For economics, I have no doubt that the shift to incorporating machine learning is a major change, a paradigm shift, compared to previous econometrics. Behavioural economics changed our theory to incorporate real human behaviour. Similarly machine learning is fundamentally changing how much of the world we can explain.

And it is highly complementary to these other economic techniques.

From water to telecoms, to aviation, to road and rail, to energy, to financial services, machine learning will change regulatory economics too. It allows us to leverage the ever-increasing pools of data we have. It allows us to extract more information and insight. And this makes us more efficient and effective.

We can use machine learning to tackle policy problems, like targeting our interventions where they will be most effective or predicting how recipients

might behave in response. We can use it to tackle operational problems, like optimising our resources and identifying trends or bottlenecks. And we can use it to tackle supervisory problems, like identifying 'bad apples' and being in the right place at the right time.

Machine learning can shed light on major problems you are facing today.

Finally, what next for our journey from map to app? I would like you to take away three ideas about the future of regulation.

First, high-quality machine-learning tools have only recently become widely available. If even small businesses can use machine learning – and they do – imagine what regulators could do. This is just the beginning.

Second, this technology really can make a difference. We are becoming much better at framing our regulatory problems as informational problems. In particular, again, many regulatory problems are, in essence, prediction problems. We can now predict much better.

Third – the big question for us all – what will be the role of economists in all this? Well, economists have the data skills, market knowledge and the analytical mindset to harness this new technology. Economists are well-placed to lead, or be heavily involved in, this charge.

These are exciting times for economists and exciting times for regulators.

Algorithms will be transforming not only our day-to-day personal lives, but our day-to-day work lives too.

Thank you.