

The challenge of algorithmic economics

Economics is far from being an empty box. Indeed, it contains what is probably the only general law across the whole of the social sciences. Namely, agents react to incentives.

Exactly how they might react is not always easy to gauge in advance, for humans can be very creative in their responses. But without this concept, many outcomes in both the economy and in society more generally can be difficult to understand.

In practical contexts, economists occupy commanding positions. Central banks, finance ministries, regulatory bodies are packed with economists. Much of public policy has to pass through the filter of economics.

It is therefore easy to appreciate why economics often stands aloof from other disciplines. Many economists feel that other scientific endeavours have little to teach them about how the economy operates.

A simple illustration of this insularity is the statistical package R. Free to download, it spans a far wider range of functionality than is contained in econometric packages. R has become the research tool of choice across many disciplines. Yet most economists remain ignorant of even its existence.

The opinions of Nobel Laureates on this matter may be rather more persuasive. Daniel Kahneman, Richard Thaler and Robert Shiller all struggled for decades to gain acceptance of the need to enrich economics with the insights of psychology. Thaler describes this process at length in his book *Misbehaving* (2015).

The injunction of Vernon Smith in his 2002 Nobel lecture needs to be taken far more seriously. He writes: "I importune students to read narrowly within economics, but widely in science. Within economics there is essentially only one model to be adapted to every application: optimization subject to constraints due to resource limitations.... the economic literature is not the best place to find new inspiration beyond these traditional technical methods of modelling" (p.510).

A central challenge for economics at the moment is how to deal with the issues raised by what we can think of as "cyber society", with its stupendous increases in both the volume of information which is being created on a daily basis, and the connectivity between agents.

I would go so far as to say that we need a new, major branch of our discipline to address this challenge, one which we might term "algorithmic economics".

What, for example, are economists contributing to the current debate about fake news? This is a topic of major concern to policy makers. Yet economists are conspicuous, and paradoxically this seems the appropriate phrase to use in this context, by the very low profile of the discipline in the discussions.

The journal *Science* is one of the top two scientific publications in the world, the other being *Nature*. On 9 March 2018, *Science* carried a piece entitled “The science of fake news”. There are 16 authors, a real multi-disciplinary team. But none of them holds a full-time post in a university economics department. One is employed as an economist by Microsoft, and Richard Thaler’s collaborator, Cass Sunstein is an author, but he is in the law school at Harvard.

It is not as though economics has failed to think about how cascades of behaviour might spread across groups of agents. The famous paper by Bikhchandani and colleagues on this – it has nearly 7,000 citations - was published 25 years ago. Their model is based on Bayesian principles. It describes how information cascades can grow through rational herding in a sequential social learning process, with each agent balancing what he or she already knows against what others can be seen to be doing.

In the same issue of *Science* a team of data scientists at MIT published the results of the largest ever study of fake news. Over 100,000 stories tweeted by some 3 million users were analysed over a ten year period.

There are two key ways to measure the spread of a tweet. The first is, quite simply, the number of users who retweet it. The second is the length of the link the tweet passes through. Most tweets are never retweeted at all. But if your tweet is retweeted by a friend, and in turn someone retweets your friend’s retweet, its “length” is two.

The conclusion of the MIT research is rather depressing. Fake news and rumours spread much faster and reach more people than accurate stories, using both measures of the spread of a tweet.

It is not immediately apparent that the economic theory of rational information cascades helps explain these findings.

Social media analysis, of course, might be thought to offer a potentially misleading picture because its user base is obviously not representative of the population as a whole in terms of standard socio-economic classifiers such as age and income.

However, there is increasing evidence that it provides a representative indication of the distribution of emotions and attitudes across the population.

A key building block of economic theory is revealed preference. We traditionally attach relatively little weight to surveys which ask for opinions on the available alternatives in any given situation, preferring to rely on the preferences revealed in actions.

In the same way, social media reveals emotion and attitude, in ways which are very hard to systematically disguise.

In the January 2018 edition of this newsletter, for example, Alan Kirman describes the study carried out on the French Presidential election by a team led by David Chavalarias at the Institut pour la Complexité in Paris. This was based on a data base of tweets and re-tweets and analysed the evolution of the groups supporting the various candidates.

Rather puzzlingly, Kirman states that the research was “difficult to implement because of the sheer quantity of information involved”. In fact, algorithms which are readily available in the public domain – in the package R for example – will comfortably handle very large amounts of unstructured text data on social media.

By coincidence, Rickard Nyman, a computer scientist at UCL, and I carried out a real time analysis of the 2017 UK General Election for a commercial client with Twitter data, using a variety of machine learning algorithms. Using only the 1 per cent random sample of tweets which is made available and applying filters to ensure the tweets were actually about the election, we obtained 8.1 million tweets by 1.2 million users during the campaign. This was analysed on a laptop using algorithms available in R.

The two main parties, Conservative and Labour, achieved their highest combined share of the vote at any election since 1970. In the 2015 election for example the share was 67 per cent, compared to the 82 per cent in the 2017 vote. There was a major squeeze on the minor parties. The analysis identified at an early stage that this was the likely outcome.

The topics which were being discussed was a key focus of the analysis. Brexit was by far the most frequently discussed topic. The British Election Survey, carried out at every election and based upon an expensive conventional survey of tens of thousands of respondents, came to the same conclusion – published in early August rather than in real time!

Until the last week of the campaign, there was a strong correlation between the Conservative lead in the opinion polls and the proportion of tweets which were about Brexit. In the last week, the polls moved in the government’s favour, and the widespread expectation was that it would be returned with a small increase in its majority. In contrast, the proportion of tweets about Brexit declined still further, so the actual result did not come as a surprise.

Robert Shiller’s Presidential Address to the AEA in 2017 was on “Narratives in Economics”. He poses the challenge to the profession that “The field of economics should be expanded to include serious quantitative study of changing popular narratives” (p.967).

Keynes emphasised the importance of sentiment and narrative, writing for example of the “waves of irrational psychology” which drive the business cycle. But he lacked the tools to make these ideas operational.

Brian Arthur (McKinsey Quarterly October 2017) notes the huge success of algorithms in diverse areas such as digital language translation, face recognition, voice recognition and inductive inference. He goes on to say that “What came as a surprise was that these intelligent algorithms were not designed from symbolic logic, with rules and grammar and getting all the exceptions correct. Instead they were put together by using masses of data to form associations”.

In other words, the algorithmic approach works by using what we might describe as a different type of intelligence. It is far from the purpose of this piece to enter into debates about what is meant by the philosophical concept of intelligence. What I mean here is simply that machine learning algorithms succeed by discovering and matching patterns in data. They have a comparative advantage over humans in this respect. The approach is different from the way in which people usually try and address these types of problem.

Shiller’s vision for the direction in which research should go is correct. But his idea that we need to learn from literary theory in order to identify narratives is not. The algorithmic approach does it for us, finding associations by clever statistical methods using large amounts of data.

The algorithmic approach to the analysis of text data has advanced very rapidly. A few years ago, a popular way of gauging sentiment, for example, was based on a count of specific words whose emotional content had been established by surveys or experimental work separate to the text being studied.

This whole approach has now been overtaken in machine learning analysis. Algorithms can learn the sentiment of a document directly from its content. They learn in the context of the overall set of documents being examined which words have positive or negative emotional content. Further, it is not just individual words which are identified, but phrases and groups of words. All this is done without reference to literary or linguistic theory.

In the context of machine learning algorithms applied to more conventional, quantitative data sets, the profession is responding rather more positively to scientific developments. Mullainathan and Spiess, for example, have a very nice paper on machine learning and econometrics in the Spring 2017 issue of the Journal of Economic Perspectives.

What needs to be appreciated is that the best machine learning algorithms are considerably more powerful than the econometric tools we have at our disposal. Manuel Fernandez-Delgado and colleagues compared the performance of 179 algorithms on 121 challenging data sets, in a paper published in the Journal of Machine Learning Research, as long ago (!) as 2014.

They found that two machine learning algorithms, random forests and support vector machines, were decisively the best. Generalised linear models and logistic regression were “simply not competitive at all” (p.3195). The research was of course carried out before the development of the new generation of deep learning neural network algorithms, though these of course do appear to need substantial amounts of data.

The dramatic rise of cyber society raises further key questions for economics. What does it mean, for example, to exercise rational choice in a world in which there is such a massive abundance of data that it is not possible to gather and process anything other than a tiny fraction of the total amount available in any given context? Can we reasonably maintain the assumptions of stable and transitive preferences when agents are bombarded with the choices, opinions and behaviours of others?

But these are broader and deeper questions. What economists can do quite readily is to embrace [the concept of algorithmic economics](#) – modelling and analysis based on AI machine learning and computational statistics - to extend our understanding of the modern world.

Paul Ormerod is a Visiting Professor at University College London (UCL). His latest book, “Against the Grain: Insights from an Economic Contrarian”, will be published this spring by the Institute of Economic Affairs in conjunction with City AM newspaper.

