CONTROLLED BY CALCULATIONS?

POWER AND ACCOUNTABILITY IN THE DIGITAL ECONOMY

PART 3: THE RISE OF ALGORITHMS

Written by: Duncan McCann, Miranda Hall and Robbie Warin

New Economics Foundation
www.neweconomics.org
info@neweconomics.org
+44 (0)20 7820 6300
@NEF

Registered charity number 1055254

CONTENTS

PREFACE

A new economy is emerging. And this new economy is powered by a new type of fuel: data. As the data economy becomes increasingly prominent, there are troubling signs that it is worsening existing power imbalances, and creating new problems of domination and lack of accountability. But it would be wrong simply to draw dystopian visions from our current situation. Technological change does not determine social change, and there is a whole range of potential futures – both emancipatory and discriminatory – open to us. We must decide for ourselves which one we want.

This is the third of four papers exploring power and accountability in the data economy. These will set the stage for future interventions to ensure power becomes more evenly distributed. This paper explores the rise of algorithms to make sense of big data, while other papers examine: the impact of the mass collection of data; the companies built on data, that mediate our interface with the digital world; and the labour market dynamics that they are disrupting.

Our research so far has identified a range of overarching themes around how power and accountability is changing as a result of the rise of the digital economy. These can be summarised into four key points:

- Although the broader digital economy has both concentrated and dispersed power, data is very much a concentrating force.
- A mutually reinforcing government-corporation surveillance architecture – or data panopticon – is being built, that seeks to capture every data trail that we create.
- We are over-collecting and under-protecting data.
- The data economy is changing our approach to accountability from one based on direct causation to one based on correlation, with profound moral and political consequences.

This four-part series explores these areas by reviewing the existing literature and conducting interviews with respected experts from around the world.

Algorithms have become essential to managing our digital lives and navigating the digital world. Without them we could not make sense of the huge mass of digital information available. But now that algorithms weild such influence, we have a responsibility not to misuse them.

- **Algorithms are actively shaping our lives**: They have morphed from curating online content to curating and influencing our lives.

- **We are entering the scored society**: Access to services, both public and private, are increasingly being mediated through algorithms which analyse our data and decide whether or not our digital profile matches the requirements for access.
- **Lazily programmed algorithms entrench discrimination and bias**: This happens especially when data is used uncritically.
- **We need to rethink accountability in age of algorithms**: How we hold these systems accountable, as well as how we use these systems to hold people accountable is shifting.
- **Data cannot be permanently anonymous**: Although diligent preparation can make it harder to re-identify data, with the right tools and techniques it is always possible.

1. INTRODUCTION

The Economist claimed in May 2017 that "the world's most valuable resource is no longer oil, but data."[1] Both have existed throughout human history, yet the rush to extract them did not coincide with their discovery but with the development of refining techniques and real world applications. The digital economy has only become obsessed with extracting ever larger swathes of data now that data refineries can turn the relatively worthless raw material into a potentially valuable commodity. The economy is also finding an ever increasing number of uses for the refined data, just as the economy found a multitude of uses for the various outputs of the oil refining process.

And just as the development of techniques for extracting and refining oil altered the distribution of power in society, so too does the development of techniques to extract actionable insights from big data. Big data analytics has created a "new wellspring of power in society"[2] that is amplifying the power of business and government over ordinary people.

Given modern society's propensity to over-collect data (as seen in Part 1: The rise of the data oligarchs, and in the huge increase in digitally mediated services) tools and techniques are constantly being developed to find and harness meaning from vast amounts of data. This refining process is carried out using algorithms. The Oxford English Dictionary defines an algorithm as "a process or set of rules to be followed in calculations or other problem-solving operations, especially by a computer."[3]

Algorithms are part of our every interaction with the internet, from deciding what search results to show us, to recommending potential friends or deciding what adverts we see. They will form part of the decision about whether we get a loan or the price of an insurance policy. They are responsible for almost 70% of stock trades in the US, fly planes for over 95% of the time in the air and may soon be driving our cars too. And they are expanding into decisions about whether to offer someone a job interview, whether someone will reoffend or what social care provision a person needs. They may come to impact literally every aspect of our lives, especially as they are increasingly mediated through digital platforms and 'internet of things'-enabled devices.

> *"Hopes of feeling in control of these systems are dashed by their hiddenness, their ubiquity, their opacity, and the lack of obvious means to challenge them when they produce unexpected, damaging, unfair or discriminatory results."*[4]

This is born out by a 2015 poll of EU citizens (some of the few citizens with a legal framework for data gathering and processing) which found that over 30% felt they had no control over the data they provided online, with a further 50% stating that they only had partial control.[5] Writers such as Diakopoulos have made clear that those responsible for keeping a check on the powerful, such as journalists, need to ensure that they adapt to include a consideration of algorithms and their impact on wider social and economic issues.

As well as considering power, the issue of accountability is vital when thinking about algorithms. Accountability has two distinct dimensions. Firstly, as we saw in Part 1, the fact that data may be stored forever starts to change how individuals are held accountable for their past actions, or at least their past data trails. Traditionally and legally we have been held responsible for the actions we undertake and for their direct and foreseeable consequences. The algorithmic focus on correlation over causation risks disrupting this assumption.

Secondly, we need to explore how to hold algorithms accountable for their effects on individuals and society. We need to examine the roles and responsibilities of companies, as well as the individuals who design, operate and manage the algorithms. We need to understand how those subjected to algorithmic decision making can seek redress or appeal decisions. Kings puts this well when he states that “when big data analytics are increasingly being used to make decisions about individual people, those people have a right to know on what basis those decisions are made.”[6]

2. ISSUES

2.1 FROM CURATING CONTENT TO CURATING OUR LIVES?

Today, 2.5 quintillion[7] bytes of data are created every day.[8] This includes 500 million tweets,[9] 3.5 billion Google searches,[10] 350 million photos uploaded to Facebook,[11] and 4 million new blog posts every day spread across over 1.3 billion websites.[12] There is now too much digital information available for any one person to navigate through alone.

It has therefore become impossible to use the internet without the content being curated for us by automated systems: algorithms. Today, the majority of content you view has been specifically crafted for you by algorithms, whether it be the search results you return from Google, the newsfeed that you see on Facebook, or the ads that are shown to you online.

Algorithms sort, filter and manipulate everything we encounter online. They define what is visible to us and therefore have the power to shape and reinforce our tastes and interests. Ultimately, in an effort to give us what we want (our desire for content being measured by clicking and sharing rates) these algorithms are starting to influence who we are and how we interact with the world. Taken together, the sum of lots of harmless nudges – a recommended TV programme here, a new friend suggestion there – add up to huge amounts of power that can change people's understanding of reality.

Whereas these algorithms were originally designed almost purely to solve the technical problem of trying to curate the increasing quantity of information as best as possible, recently people have become concerned about their power to actively shape our internal and external lives. Research has shown that we can be influenced and impacted by the results of algorithms.[13,14,15] This passage from The Filter Bubble gives a perfect overview of positive and negative implications of the curated digital world that we live in:

> *"The filter bubble invisibly transforms the world we experience by controlling what we see and don't see. It interferes with the interplay between our mental processes and our external environment. In some ways, it can act like a magnifying glass, helpfully expanding our view of a niche area of knowledge. But at the same time, personalized filters limit what we are exposed to and therefore affect the way we think and learn. They can upset the delicate cognitive balance that helps us make good decisions and come up with new ideas. And because creativity is also a result of this interplay between mind and environment, they can get in the way of innovation. If we want to know what*

the world really looks like, we have to understand how filters share and skew our view of it."[16]

All this happens without most people being aware of it. As with data collection, the algorithmic management of our online existence happens without many users being aware of the nature of what is happening.

A famous experiment conducted in 2012 by researchers at Facebook and Cornell University looked at whether the emotional state of users could be altered by tweaking the algorithm that decides what to show on a user's newsfeed. The research involved almost 700,000 unwitting Facebook users, and demonstrated that "emotional states can be transferred to others via emotional contagion, leading people to experience the same emotions without their awareness."[17] There has been criticism of the experiment as it did not seek informed consent from participants, instead relying on the general conditions that people sign up to when joining Facebook.[18]

Algorithms are actively curating the internet for us by giving us more of what we like, and obscuring content that we do not want to see. This seems like it could be beneficial, as algorithms can distil the endless reams of online material into something manageable and tailored to our interests. But scratch the surface and we find a number of issues. Firstly, the algorithm decides what we like and don't like based on stated and associated preferences, on our clicks and shares, and on the corresponding correlations that follow from those – not on our actual selves. Secondly, the algorithm's propensity to give us more of what we want prevents us from getting a diverse perspective on current events, which in turn makes us more likely to believe false information. Although there are many factors at play it is no coincidence that the rise of the algorithmic curation of our online lives correlates highly with the rise in fake news.[19,20]

Algorithms that seek to curate content are designed to increase our 'engagement' with content on a particular site. Many are concerned that this is driving people to consume and share increasingly 'extreme' material online. This is driven neither by a sudden shift in people's desire to view this content nor a slow desensitisation after prolonged exposure to milder material. For example, we have a tendency towards moribund curiosity, and algorithms trained to maximise engagement would learn that humans are deeply engaged by accidents, and would promote that content as a result. Research from Zeynep Tufecki, insightful critic of our ever more algorithmised world, looked at YouTube recommendations. What it noticed was that regardless of a video's subject matter, the algorithm's recommendations steadily veer you towards ever more extreme content.[21] This is not always sinister: starting with a video on vegetarianism leads to ones about veganism, while jogging turns into ultramarathon running. But the same dynamic plays out for politics or any other subject. Ultimately this matters because more and more people, especially young people, get their news from sites like YouTube that focus on curating their content and optimise for 'engagement'.

As more of us get to grips with the propensity towards extreme content, there are growing calls for digital platforms to ensure that all they remove material that is deemed unacceptable. This may seem hard to object to – but the power wielded by these platforms is worrying.

2.2 WELCOME TO THE SCORED SOCIETY

The Chinese government has announced its intention to create a social credit score for every citizen. Large scale voluntary pilots, conducted by companies rather than the government themselves, are already underway. The biggest are run by the Chinese digital economy giants Alibaba and Tencent and seek to integrate financial and social data into a single score. Ed Jefferson outlines some examples:

> *"If someone is tracked playing video games, they're probably lazy. Decrease their score! If they buy nappies, they're probably a parent, and so probably responsible. Level up! Similarly, interacting with people deemed trustworthy by the system raises your score, while presumably spending too much time with ne'er-do-well gamers will crash it."*[22]

And we are already starting to see the direct consequences, with The Telegraph reporting in March 2018 that over 12 million people in China were placed on domestic travel bans, effectively banning them from buying plane or train tickets.[23] Although the US and Europe are unlikely to adopt the Chinese system any time soon, many similar systems are already being used by the state and big corporations.

In the private sector, algorithms are used to decide whether people should get access to crucial opportunities, including the ability to obtain loans, work, housing, and insurance, and are used to "assess desirable employees, reliable tenants, valuable customers – or deadbeats, shirkers, menaces, and wastes of time."[24] Historically this has been done by combining ever increasing amounts of financial data with other publically available data such as legal proceedings. Today, a new specialised service industry has emerged that enables companies to buy access to constructed digital profiles of people. In a quest for ever more information some companies have now started to look at alternative sources of data. Social media data has been identified as a source of data which could reflect individuals' 'true' personality and it is being used to decide people's likelihood to repay a loan, their potential risk for insurance, or whether they will make a responsible tenant.

The 'fintech' (financial technology) industry has been somewhat of a trailblazer in this area, providing 'alternative credit scores'.[25] For these algorithms, whether you organise your phone contacts by first or last names, and whether you call your mother regularly will help determine your eligibility for a loan. In Kenya, where there is little data protection, a startup called Tala now offers loans based on an Android app which uploads almost the entire contents of a prospective customer's phone to US based servers for data analysis. Loan decisions are based on an analysis of data from the customer's address book, phone records, location data and SMS messages. Alternative credit scoring is targeted at 'no-file' or 'thin-file' customers with little or no existing credit history, particularly within the Global South. Without data protection legislation, companies keep the data for as long as they want, able to use it for any purpose and, if the company is sold, the data becomes another asset in the sale.

The 'proptech' industry is taking a lot of the same tools to the rental market.[26] A new generation of apps offer to 'digitise the renting transaction', while firms such as Blackstone and Colony Capital have developed new software platforms enabling them to invest in and manage massive portfolios of geographically dispersed homes and residents. Everything from maintenance requests to rent payments can be processed through cloud based platforms like TaskEasy and FixFlo. Even eviction has a software solution. 'Trust scores' based on your job, credit score and other personal information allow landlords to screen tenants. Your score could determine not just your access to homes but how you are treated within them. The Waypoints platform gamifies renting by giving tenants points for behaviours aligned with the interest of landlords. These entitle you to new appliances, smart home technologies and general home upgrades.

Scoring systems, in particular credit scores, do not just evaluate behaviour but also ultimately modify it.[27] In the past, these scoring systems were limited to the financial sphere and heavily regulated. The spread of alternative credit scoring, and people's awareness of it could start to have ripple effects on our online activities, such as reducing political and non-conformist activity on social media. A loan provider in India used Twitter profiles to determine eligibility, avoiding those involved in political campaigns.[28]

The public sector is also aggressively rolling out algorithms. Around the world we are seeing a simultaneous increase in poverty and a decrease in resources allocated for those in need. Instead of relying on caseworkers to determine who is eligible for these chronically underfunded services, public authorities around the world are investing in algorithms to rank and rate which families deserve access to housing, food and healthcare, and which do not. This "reframes questions of justice as questions of efficiency."[29]

A graphic American example demonstrates the damaging effects relying on algorithms can have. Tammy Dobbs in Arkansas had extensive care needs due to cerebral palsy. The nurse originally assessing her decided that she required 56 hours of home care visits per week. A few years later she was made to fill out an online form on how often she needed to use the bathroom, her eating patterns, her emotional state and more. An algorithm used this to determine how many hours help she would receive: it was cut to 32 hours. There was no way to challenge this judgement and local authorities defended the assessment as 'objective'. Other automated systems have been seriously flawed and their errors difficult to correct. In Idaho, a similar system resulted in funds for those receiving home care dropping by as much as 42%. When people tried to understand why this had happened, the state refused to share the calculations, saying it qualified as a trade secret.

There are also plans for algorithms to be used to screen visa applicants in the US based on all their available data, including social media. This 'extreme vetting initiative' should be able to predict the chances of a visa applicant becoming a terrorist versus becoming a contributing member of society. Cathy O'Neil has called the programme "a pseudo-scientific excuse to prevent a lot of perfectly good people from coming into our country as immigrants."[30]

Similar systems are being implemented in the UK. One algorithm seeks to identify 'at risk' children and families by looking at school and health records and other sources.[31] Having been in operation for three years, eight out of 10 families assessed as requiring intervention are accurately identified, and the councils using the system put their savings at close to £1m. Although saving money can be a good outcome there is a concern about algorithms being specifically to produce efficiency savings at the expense of those they are casting judgement over.

Since 2013 Durham have been using the Harm Assessment Risk Tool (HART) used to help identify people suitable for specific rehabilitation services, often as an alternative to custodial sentences.[32] The programme offers those who are scored as being at medium risk of re-offending the opportunity to participate in a programme and, if completed successfully, leave without any criminal record. As this programme has potentially life-changing consequences, it is vital that the algorithm make correct decisions, but the risk of discrimination and bias is acute especially where historical data is used uncritically.

2.3 BIAS AND DISCRIMINATION IN ALGORITHMS

The huge increase in the use of algorithms has been driven by advances in machine learning techniques. Whereas historically algorithms would be programmed to complete a task through the input of clearly defined instructions, modern algorithms based on machine learning allow computer systems to create their own instructions based on detecting correlations in huge data sets, learning a multitude of ways to complete a task and creating relationships between inputs and outcomes. The development of machine learning has allowed for computers to take on tasks previously only available to humans, and complete them with extraordinary accuracy. One image labelling algorithm which performed at 72% accuracy in 2010 actually overtook humans to 96% by 2015.[33] The dystopian image of a jobless future where robots have taken over should resemble less an army of Terminators than many lines of code.

Algorithms are increasingly used as tools to facilitate evidence based decision making, utilising large data sets to draw statistically significant relationships between different factors. As the use of these tools expands into more and more areas of our lives, they are bringing with them new ways in which bias and inequalities are reproduced in our society. This bias can emerge from a range of different sources, including who develops software systems, and what goals are prioritised as a system is developed.[34] In addition, bias can emerge depending on what training data is used in developing the algorithm. Algorithms can make decisions on the basis of protected attributes like race, income, or gender even when those attributes are not referenced explicitly, because other information can be used as a proxy for these things – for example, location data or patterns of consumption. As a result, algorithms can unfairly limit opportunities, restrict services, and produce 'technological redlining'[35] - a form of digital data discrimination that uses digital identities and activities to bolster inequality and oppression. Even though human programmed algorithms can also produce biased and discriminatory decisions, identifying these can be simpler because the conditions were physically and consciously programmed.

Modern algorithms are trained by data – lots of it. Whether algorithms are used to recognise faces, find the quickest route across a city, or decide how likely a criminal is to reoffend, they are trained on relevant data which allows them to learn to perform this task to the highest degree of accuracy. But data is not neutral. Contained within the reams of digits, data represents a record of the social interactions and conditions present within everyday life. Whilst it may claim to represent 'the facts', it is blind to the human factors which influence what the facts are.

For example, a data set may show that women within a particular institution are paid less than their male counterparts. But what isn't contained within this data set is an understanding of the history of patriarchy and the systematic undervaluing of women's work. The data set merely contains the outcomes of these processes. Google's 'autocomplete' function which anticipates the rest of a search query based on its first word or two has often demonstrated racist and sexist stereotypes.[36] Its image search has also generated biased results, including categorising photos of Black people as 'gorillas'. Google's solution to the problem was to censor 'gorillas' as a search term and image tag, and even extended the ban to 'chimp', 'chimps' and 'monkey', effectively "rendering these animals unsearchable and in some sense invisible to the artificial intelligence (AI) that powers Google's image searching capabilities."[37]

So how does the application of machine learning work to recreate the biases contained within the data sets used to train them? One example from the US and Canadian judicial systems may provide some clues. The COMPAS software system is used to provide a guide for judges during the sentencing of convicts, through generating a score which indicates the risk that a convict will reoffend, based on a range of criteria and the results of a questionnaire. However, whilst race was not included anywhere within the questionnaire, the system was found to provide significantly higher scores for Black offenders than for white.[38] The machine learning process was found to reproduce the existing social inequalities contained within the data, including varying sentences for different ethnic groups regardless of background and crime.

Machine learning is also increasingly applied to the health sector, creating tools to aid diagnosis and prescription. This involves processes such as using algorithms to analyse images and scans for cancer tissues, allowing for higher accuracy diagnoses.[39] Training data used for healthcare algorithms tends to come from clinical trials, which have been historically dominated by white men. This means that diagnoses emerging from the machine learning process are more likely to be wrong for women and/or people of colour, including an increased likelihood of sickle cell anaemia being misdiagnosed as diabetes amongst groups of African heritage.[40].

A number of questions are being asked about the potential negative social impacts of using predictive policing (PP) algorithms. Central to this concern is the potential bias located within data used to train PP algorithms and the potential impacts this will create in how different areas and communities are policed.[41,42] One study comparing the rate of drug arrests with drug use across the city of Oakland found that, whilst drug use is fairly evenly distributed, drug arrests are concentrated in a few distinct communities. When this data was used to train existing PP systems, this existing bias was replicated.[43] The authors hypothesise that this has the potential to create a feedback loop whereby the presence of more officers in an area will lead to a greater number of arrests, triggering a confirmation bias.[44]

Addressing discrimination within machine learning systems is complicated due to the ability of algorithms to make decisions which discriminate against a particular group via proxy data points. Sensitive data, such as that surrounding ethnicity, gender, voting preference, religion and sexuality is, in most western economies, legally protected from use within decision making. However, the role that factors such as race, ethnicity, gender and age, play in determining factors which *can* be included in decision making, means that even if this protected data is not included in training data, socially unacceptable discrimination can still occur.[45] This is demonstrated in Chicago, where, after accounting for all other factors, minority ethnic neighbourhoods were found to pay 30% more for car insurance than majority white neighbourhoods.[46] The distribution in the coverage of Amazon Prime across America is determined by automated decision making, and has been shown be significantly less prominent in minority ethnic neighbourhoods.[47] Race and ethnicity were not factors included in the training data for either of these cases, but this did not manage existing inequalities being reproduced.

The growing recognition of the impact which algorithms have on our lives, and the potential for them to discriminate against particular groups, has led to a number of proposed solutions. Danielle Citron's call for 'technological due process' is particularly important for big data, and it should apply to both government and corporate decisions. A bill, recently put forward by the sitting mayor of New York, takes a different approach by creating a task force with the aim to scrutinise the use of algorithms within the city's agencies for bias and discrimination. Whilst this is a positive step, we are a long way off seeing whether this is an effective strategy for mitigating the worst impacts of discrimination in algorithmic decision making.

2.4 ACCOUNTABILITY IN THE AGE OF ALGORITHMS

Algorithms are making more and more of the decisions that affect our lives, including our access to goods and services. Although algorithms held out the promise of a more neutral decision making process, in reality it is more accurate to think of them as "an opinion embedded in mathematics."[48] Increasingly, our digital footprint (ie: the data that has been collected, inferred, refined and assigned to our profile) will dictate what we can access and the conditions of that access. In the next section we will examine the implications of algorithmic decision making for the accountability of human data subjects, whereas here we will focus on the accountability of the algorithms themselves and the companies that operate and build them.

Algorithms are powerful and can be very commercially profitable. When Amazon started up in the mid-1990s, one of its competitive advantages was the team of professional reviewers and editors who curated the 'Amazon Voice', named by the New York Times as some of the most influential people in the book trade. However Jeff Bezos, the CEO of Amazon, was not satisfied and wanted to deliver personalised recommendations, with the ultimate goal to show just one book to prospective shoppers: the one they wanted to buy. Today, one third of all book sales on Amazon result from its recommendation algorithm, and the Amazon Voice team has long been disbanded.

What this process demonstrated was that when the goal is stimulating sales, what purchases correlate with other purchases is the most important thing to know. Knowing why these choices are linked is fairly irrelevant: "'In this case knowing *what*, not *why*, is good enough."[49] A potential strategy for making algorithms accountable for their decisions is 'the right to an explanation', enshrined in the General Data Protection Regulation (GDPR) which came into effect in May 2018. This could, in theory, provide a means by which any citizen would be able to gain an explanation for why an automated decision which impacts upon their life has been made. There remain questions to be answered, such as: how can explanations of machine learning algorithms' decisions be made intelligible to humans? And when has a decision been made solely by an algorithm? Recent debates around the GDPR highlight the complexity of the issues around algorithmic accountability, but also demonstrate that discourse has shifted towards finding strategies to tackle them.

2.5 HUMAN AGENCY, CORRELATION AND OUR DIGITAL SELVES

The key feature of algorithms is the ability to discern statistically significant relationships between data points, enabling them to make a decision or prediction. But even a very strong correlation may just be the result of chance. Big data tells us nothing about causation, although it may suggest interesting avenues for further research into causal links.

Humans have always searched for causal links to help them understand the world around them. We have a hard wired predilection for seeing causation.[50] One of the risks of the big data revolution is its emphasis on relying increasingly on correlation between data points. We need to be very wary of the effects of conflating correlation between data points with causation.

Increased use of algorithms to process more and more data can provide ever more detailed and complex correlations that are being used to make decisions about us. As we increase the amount of data that we are able to analyse we also increase the likelihood of chance correlations being interpreted as causal relationships. The volume of US crude oil imports from Norway, and the number of US drivers killed in a collision with a train correlate to a staggering 95.45%, which will be interpreted by algorithms as statistically significant, despite having no causal relationship.[51]

Andresen famously argued that this spells the end of the scientific method and that this new approach to data analysis means that correlation is 'enough'.[52] Although he has since backtracked from this position, Andresen's provocation does force us to consider the limits of using correlation to make real world decisions. Figure 1 provides one set of guidelines.[53] It weighs confidence in the reliability of the correlation against the trade-off between the risk and reward when acting on the correlation. The result is that, when trying to prove causation, when risks are higher, the burden of a strong causal hypothesis should also be higher.

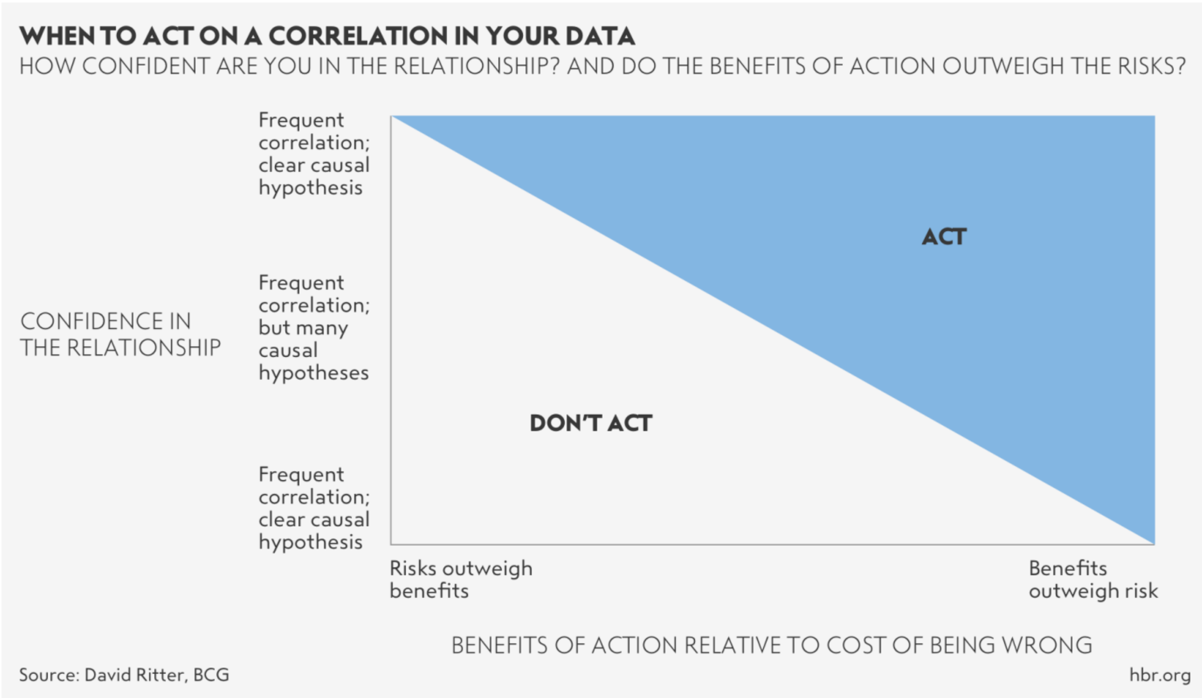

Figure 1. When to act on a correlation in your data[54]

So in fact big data does not spell the end of the scientific method, but it does start a potentially fundamental transformation in the way we make sense of the world. Big data advocates Viktor Myer Schonberger and Kenneth Culkier posit that although this change is very challenging, "the tremendous value that it unleashes will make it not only a worthwhile trade-off, but an inevitable one.'[55]

As well as addressing when it is suitable to use algorithms to make decisions, we also need to consider who they are making decisions about. Because these algorithms are making decisions not based on our actual lives but through combining and analysing our data footprint to create a profile about us. Much of the attributes assigned to our profiles are inferred through correlation.

> *"Who we are is composed of an almost innumerable collection of interpretative layers, of hundred of different companies and agencies identifying us in thousands of competing ways."*[56]

Because these digital identities are dynamic and constantly updating based on new data, it can be hard for us to maintain knowledge and control over them. The act of classification itself bestows power on those companies assembling digital profiles, as they become effectively the gatekeepers deciding what goods and services someone can access.

The use of digital profiles to make decisions about us will also make it harder to break with previously associated attributes, and "tends to reproduce and reinforce assessments and decisions made in the past."[57]

Credit scores are a highly regulated area where companies have collected data to form profiles which determine whether an individual can access certain financial products. Until recently credit scores were based on information that credit score companies such as Equifax or Experian had collected. In order to see your profile score you needed to pay. Recently, companies have been required to provide free access to the information about an individual, as well as a way for individuals to contest incorrect information. We must learn from our experience with credit scores and ensure that people can view and amend digital profiles held about them. Many companies such as Google and Facebook already allow users to download all data held about them and review attributes that have been assigned to them. But it is hard to see how this is a scalable solution, since there are hundreds or thousands of digital profiles about each individual, being used at various times. Their dynamic nature makes it unrealistic to expect people to maintain and correct their own profiles.

2.6 NO SUCH THING AS ANONYMOUS

In order to realise the full potential of the data revolution we need to be able to share and analyse the data that we collect. Big data techniques offer us a way to fundamentally transform how we fight disease, design cities, or undertake research. Our health data, for instance, holds great potential to save lives and reduce suffering, yet most would not want identifiable data on their personal health available to the general public.

The potential benefit of using big data analytics forces us to get to grapple with a dilemma that "people want both perfect privacy and all the benefits of openness. But they cannot have both."[58] Or can they? Could anonymising the data by removing all personally identifiable information enable more valuable big data sets?

It is of course technically possible to scrub all personally identifying information out of a dataset. Doing it properly is a complex and sometimes laborious process which means that it is not always done as well as possible. It is these anonymised data sets that are most frequently traded, allowing them to be shared while retaining people's privacy. Or does it?

Although the data has been de-identified, researchers have found that when analysed alongside other datasets re-identification becomes possible. A 2015 study looked at three months' worth of credit card data for over 1 million people, which had been scrubbed of all personally identifiable data such as name, card number or address. The research found that just "four spatiotemporal [location and date] points are enough to uniquely re-identify 90% of individuals."[59] Studies such as this suggest that it is never possible to render data truly incapable of being de-anonymised.[60,61,62]

However, as with so many things, the devil is in the detail. A major paper in 2014 sought to reassure critics that "de-identification does work."[63] It argued that de-identification is a complex process that, done well, can significantly reduce the chance of re-identification and that re-identification requires a high grade of technical skills. But this rebuttal has not been enough to combat claims that "outside of the controlled environment of academic research, both anonymity and privacy are essentially dead."[64] In the real world "de-identification is hard, and re-identification is forever."[65] Some of the most valuable data points, like location, cannot be de-identified, meaning they it can never be shared anonymously.[66]

But it is important that we do not throw out the baby with the bathwater. The use of de-identified data, especially health data, has "important social benefits … as well as business, commercial, educational benefits and innovation opportunities."[67] The UK government recently committed to making re-identification of anonymised data a crime,[68] but ultimately backtracked because of worries about the potential impact on privacy and security researchers.[69] Even if they had made re-identification a crime, it would not have resolved the issue, as some people would be willing to risk breaking the law. And since the consequences of loss of anonymity through re-identification are potentially permanent because once data is out there it is hard to remove completely. This will be one of the most complicated and potentially intractable issues in trying to maximise the potential positive uses of anonymised big data sets to enable society to get the most from this potentially transformative resource.

CONCLUSION

If data is the 'new oil', then algorithms are the 'new refineries', transforming the raw material into a valuable resource. A collection of data points has little value on its own, but algorithms allow us to process it into information which now powers our most prominent digital services.

The role of algorithms has morphed over time as the scale of the digital economy has expanded. The innocent helpers of the early internet who helped make sense of this new frontier now have taken on a role of actively shaping our lives and interests. We now have an industry serving us more of what we 'want' while blocking out alternative views. This makes it harder for us to make informed decisions, while efforts to make us 'engage' more lead us to consume ever more extreme content.

At the same time, the use of algorithms is expanding rapidly in all areas of society. In the private sector, algorithms are used to decide whether people should get access to crucial opportunities, including the ability to obtain loans, work, housing, and insurance, while in the public sector they are being used to determine who is eligible for essential services.

This shift has been presented as a positive innovation, which removes decision making from the messy human realm of bias and discrimination. However the way in which most algorithms are trained on historical data has in fact algorithms to replicate these biases more deeply.

Who should be responsible for the decisions of algorithms? Recent debates around GDPR highlight the complexity of the issues around algorithmic accountability, but also demonstrate that that we are at least making steps towardstackling them. We also need to consider the nature of human accountability in a post-algorithmic age. Ubiquitous data held permanently and correlated with other data means that we may never be able to escape our past actions nor the attributes which correlate with our digital selves.

Finally, as more data is generated and shared the re-identification of anonymous data becomes possible. The most useful and valuable data, socially and economically, is also the most sensitive, like health or financial data. As we face an increasingly algorithmicly mediated world, we urgently need a better understanding of how to protect people's privacy and anonymity.

ENDNOTES

[1] https://www.economist.com/news/leaders/21721656-data-economy-demands-new-approach-antitrust-rules-worlds-most-valuable-resource
[2] Diakopoulos N, "Algorithmic Accountability: journalistic investigation of computational power structures", Digital Journalism, 3:3, 398-415 p.398
[3] https://en.oxforddictionaries.com/definition/algorithm
[4] Edwards & Veale, "Slave to the Algorithm? Why a 'right to an explanation' is probably not the remedy you are looking for", p.3
[5] https://ec.europa.eu/info/strategy/justice-and-fundamental-rights/data-protection_en
[6] https://review.law.stanford.edu/wp-content/uploads/sites/3/2016/08/66_StanLRevOnline_41_RichardsKing.pdf
[7] Quintillion is 10^{18} or 1,000,000,000,000,000,000
[8] https://www.ibm.com/blogs/insights-on-business/consumer-products/2-5-quintillion-bytes-of-data-created-every-day-how-does-cpg-retail-manage-it/
[9] http://www.internetlivestats.com/twitter-statistics/
[10] http://www.internetlivestats.com/google-search-statistics/
[11] http://www.businessinsider.com/facebook-350-million-photos-each-day-2013-9?IR=T
[12] http://www.internetlivestats.com/
[13] https://www.theguardian.com/technology/2012/dec/16/networker-algorithms-john-naughton
[14] https://www.theneweconomy.com/technology/the-troubling-influence-algorithms-have-on-how-we-make-decisions
[15] Michele Willson (2016) Algorithms (and the) everyday, Information, Communication & Society, 20:1, 137-150, DOI: 10.1080/1369118X.2016.1200645
[16] Pariser E (2011) The Filter Bubble: What the internet is hiding from you p.53
[17] Kramer A, Guillroy J & Hancock J, "Experimental evidence of massive-scale emotional contagion through social networks", PNAS 2014 June, 111 (24) 8788-8790. https://doi.org/10.1073/pnas.1320040111
[18] Flick C, "Informed consent and the Facebook emotional manipulation study", Research Ethics, Vol 12, Issue 1, 2016, http://journals.sagepub.com/doi/full/10.1177/1747016115599568
[19] https://tech.co/google-algorithm-magnifying-fake-news-2017-01
[20] http://observer.com/2016/12/the-problem-isnt-fake-news-its-bad-algorithms-heres-why/
[21] https://www.nytimes.com/2018/03/10/opinion/sunday/youtube-politics-radical.html
[22] https://www.newstatesman.com/world/asia/2018/04/no-china-isn-t-black-mirror-social-credit-scores-are-more-complex-and-sinister
[23] https://www.telegraph.co.uk/news/2018/03/24/chinas-social-credit-system-bans-millions-travelling/
[24] Frank Pasquale
[25] Tom Fisher 'Data Dystopia: The Perils of Fintech' in *Stir to Action*
[26] Desiree Fields 'The Automated Landlord' in *Our Digital Rights to the City*
[27] Kear M (2018) Playing the credit score game: algorithms, 'positive' data and the personification of financial objects, Economy and Society, 46:3-4, 346-368, DOI: 10.1080/03085147.2017.1412642
[28] https://medium.com/privacy-international/social-media-intelligence-and-profiling-in-the-insurance-industry-4958fd11f86f
[29] Eubanks V. (2018) Automating Inequality: How high tech tools profile, police and punish the poor
[30] Cathy O'Neil.
[31] https://apolitical.co/solution_article/london-uses-data-predict-which-children-abuse/
[32] https://www.wired.co.uk/article/police-ai-uk-durham-hart-checkpoint-algorithm-edit
[33] https://royalsociety.org/~/media/policy/projects/machine-learning/publications/machine-learning-report.pdf
[34] AI Now 2017 Report : https://ainowinstitute.org/AI_Now_2017_Report.pdf
[35] Noble S. (2018) *Algorithms of Oppression: How Search Engines Reinforce Racism*
[36] Noble S. (2018) *Algorithms of Oppression: How Search Engines Reinforce Racism*
[37] https://boingboing.net/2018/01/11/gorilla-chimp-monkey-unpersone.html
[38] https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing

[39] https://royalsociety.org/~/media/policy/projects/machine-learning/publications/machine-learning-report.pdf
[40] AI Now 2017 Report : https://ainowinstitute.org/AI_Now_2017_Report.pdf
[41] Shapiro, A., 2017. Reform predictive policing. [Online] Available at: https://www.nature.com/news/reform-predictive-policing-1.21338
[Accessed 1 December 2017].
[42] Lum, K. & Isaac, W., 2016. To predict and serve?. *Significance* , October, 13(5), pp. 14-19.
[43] Ibid
[44] Ibid
[45] Fairer machine learning in the real world: Mitigating discrimination without collecting sensitive data . Michael Veale and Reuben Binns
[46] https://www.propublica.org/article/minority-neighborhoods-higher-car-insurance-premiums-white-areas-same-risk
[47] https://www.bloomberg.com/graphics/2016-amazon-same-day/
[48] Cathy O'Neil
[49] Mayer-Schoberger V & Cukier K (2013) Big Data p.53
[50] Matute, H., Blanco, F., Yarritu, I., Díaz-Lago, M., Vadillo, M. A., & Barberia, I. (2015). Illusions of causality: how they bias our everyday thinking and how they could be reduced. *Frontiers in Psychology, 6*, 888. http://doi.org/10.3389/fpsyg.2015.00888
[51] http://www.tylervigen.com/spurious-correlations
[52] Andresen C, "The end of Theory", Wired Magazine, 23rd June 2008, https://www.wired.com/2008/06/pb-theory/
[53] Ritter D, "When to act on a correlation and when not to", Harvard Business Review, 19th March 2014, https://hbr.org/2014/03/when-to-act-on-a-correlation-and-when-not-to
[54] Ritter D, "When to act on a correlation and when not to", Harvard Business Review, 19th March 2014, https://hbr.org/2014/03/when-to-act-on-a-correlation-and-when-not-to
[55] Mayer-Schoberger V & Cukier K (2013) Big Data p.72
[56] Cheney-Lippold J., (2018) We are Data, p.6
[57] Gandy 'Exploring identity' p.1101
[58] https://www.economist.com/news/science-and-technology/21660966-can-big-databases-be-kept-both-anonymous-and-useful-well-see-you-anon
[59] de Montjoy Y, "Unique in the shopping mall: On the re-identifiability of credit card metadata", Science 30 Jan 2015:Vol. 347, Issue 6221, pp. 536-539
[60] https://qz.com/232838/heres-why-you-should-stay-wary-of-anonymous-data-but-dont-go-full-luddite/
[61] https://hbr.org/2015/02/theres-no-such-thing-as-anonymous-data
[62] https://bits.blogs.nytimes.com/2015/01/29/with-a-few-bits-of-data-researchers-identify-anonymous-people/
[63] http://www2.itif.org/2014-big-data-deidentification.pdf
[64] Pam Dixon quoted in https://www.csoonline.com/article/2902655/big-data-security/big-data-collection-makes-it-hard-for-you-to-remain-anonymous.html
[65] https://qz.com/232838/heres-why-you-should-stay-wary-of-anonymous-data-but-dont-go-full-luddite/
[66] https://qz.com/232838/heres-why-you-should-stay-wary-of-anonymous-data-but-dont-go-full-luddite/
[67] https://datafloq.com/read/re-identifying-anonymous-people-with-big-data/228
[68] https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/635900/2017-08-07_DP_Bill_-_Statement_of_Intent.pdf
[69] https://www.theregister.co.uk/2018/01/10/uk_data_protection_bill_tweaked_to_protect_security_researchers/