So what? (and for that matter "how"?)

Richard Hamblin
Leadership Development Centre 2018 Fellow
December 2018

Acknowledgements and disclaimer

I have been privileged to talk to a wide range of experts here and abroad over the last 12 months about this topic. I am grateful for their generosity of expertise, time and wisdom. I am struck, as ever, that I am fortunate to live in a world where the mantra “knowledge is power” is held by no-one.

I would like to thank leaders and experts from Stats NZ, the Social Investment Agency, and the NZ Transport Agency among others who gave generously of their time. I would also like to thank in particular Dr Tom Smith and team at the Data Science Campus, Wales for generously hosting me, Steve Skelton of the Greater Manchester Analysts Network for his time and networks and Dr Chris Russell of the Turing Institute for his insights into the ethical issues of algorithms.

That said, what follows are my opinions and conclusions, any errors, misrepresentations and controversies are mine alone.

Contents

Introduction

My earlier paper[1] suggested why we need to change how we view data in public services. Data should be seen as a source of insight rather than an instrument of control, and that where monitoring "targets" exist these should be locally relevant and mutually agreed rather than centrally imposed.

What does that mean in practice, and how do we go about doing it? Widespread reading of literature, and conversations with experts in New Zealand suggested that there are five key elements to get right in order to achieve this shift. Since then, the LDC have been good enough to support me in visiting leading overseas organisations. The first of these were two very different bodies in the UK.

The Data Science Campus (DSC) is the national centre of excellence for data science, set up by the Office for National Statistics in their major office, but physically separate and architecturally distinct from the rest of the site. Around 50 staff work on two basic aims under the rubric "data science for the public good": undertaking specific data science projects for the public good in collaboration with stakeholders across government and wider UK; and building data science capacity for government.

The Greater Manchester Analyst network by contrast is a more organic (though still intentional) shift towards using data more imaginatively across the ten local governments of Greater Manchester (population c 3 million), by organisations with a direct role in providing and commissioning services for clients. This has been done against a background of fiscal austerity in local government and devolution of social budgets from central to local government.

Observing how they work, talking to them about their agendas, and reviewing their outputs has confirmed to me that the five key elements are broadly plausible as a categorisation of what needs to be done to make the shift in using data. In short, they are:

1 Establishing the higher purpose in our use of data

2 Developing new skills in statistics and programming which are not currently widely held in public sector analytic teams

3 Establishing the roles that need to be played in the data teams of the future, and how these relate to operation of organisations in general, and interaction with the public

4 What might be described as the social positioning of the teams inside organisation: in other words what uniquely does the organisation do in practice to achieve their higher purpose, and thus what uniquely is the role of the data team within this. A subset of this are the ethical issues around use of data and how this directs both policy and implementation of services.

5 Exercise of influence – how do you get organisations to use the insights that you generate to inform decision making.

This paper considers each of these in more depth and tries to provide at least some answers as to how to achieve these.

[1] 30 Years down the wrong rabbit hole: how we got there and how we get out, Richard Hamblin, July 2018

The importance of higher purpose

Of all the "inspirational" little vignettes that disfigure management literature and LinkedIn streams, few are as hackneyed as the tale of JFK touring NASA headquarters in 1961 and on asking a janitor what he did for NASA receiving the reply "I'm helping to put a man on the moon". So far, so Dilbert. Yet the general principle of workers in an enterprise sharing a view of why the enterprise exists, and believing that their efforts contribute to a goal is surely a valuable one. At the very least, the failure of organisations where there is no commonly shared and understood goal is a commonplace, and the alienation of workers who see no connection between their activities and a meaningful end goal has been understood since Marx (at least2).

However, what I have noted in both the literature and in visiting exemplar organisations is that organisations and networks who are using data in innovative ways to tend to have a quite clearly articulated higher purpose. For example, the New Zealand Social Investment Agency state "by analysing data, we're supporting the social system to better understand how we can best invest in the social wellbeing of New Zealanders", and the various products and projects that they make available have an obvious link to this. The UK ONS Data Science Campus speaks of "Data Science for the Public Good". Its own work and its work in developing the skills in the rest of the UK public sector are dedicated to this, and, critically, all its staff who I spoke to (including support staff) understand their work in the context of achieving this goal. In some ways more interesting still was the commonality of response from the Greater Manchester Analysts network. This group were spread across organisations, and indeed between sectors, and yet all had a common purpose which reflected a response to local circumstances of austerity and devolution. Again, the local, real, activity reflected the (almost unstated) common purpose.

And this is really quite interesting. The "higher purpose" mantra of organisations is often untrue (because the real higher purpose is hidden and about securing resources, profit or power), so seeing an organisation and a network where colleagues from all levels clearly articulate a shared higher purpose is unusual. I found no cynicism (although there were frank admissions about the difficulties and challenges of the work). There was also universally a clear "line of sight" between their own agendas and the organisational goal.

My reflection is that the reason that this has happened is that actions occurred in response of the higher goal. For example, the process for prioritising the projects that the Data Science Campus was willing to take on included scoping and discovery phases that explicitly and rigorously questioned how the project would create public good. To progress, projects needed to: add value to the UK; have a committed partner (i.e. this is someone's day job); and show the likelihood of building new skills or providing new data and so forth. In other words, commitment to the higher goal helps determine what actions are taken, and the actions themselves are likely to further the goal.

Making higher purpose meaningful

As a thought experiment, I set myself the challenge of defining a higher purpose for the shift in use of data that I advocate and the changed behaviour and from that determining the types of actions that need to be taken.

Based on the position that I set out in the first paper, the higher purpose for achieving a transformed use of data might be that,

[2] Qoheleth in Ecclesiastes arguably made the same point the best part of 3,000 years earlier

"By 2025, the New Zealand public sector sees data as a source of insight rather than a means of control"

To achieve this we need, there are five things which should act as a guide of behaviour, a test as to whether we are doing the right things. The good news is that many of the approaches are already in place somewhere or in development.

1 Stop arguing about the validity of individual data sets and use an agreed set based upon an agreed set of standards. The proposed Open Data Action Plan would seem to have the capacity to put much of what is needed in place on this[3]. It is critical for data owners to rigorously and consistently outline what the data set contains, and what it can and can't be used to show.

What this prerequisite doesn't imply is that there is one "approved" set of measures or indicators. There must remain the capacity to exploit the data sets to measure what needs to measured in a way that is appropriate to policy needs. There is not "one" measure of mortality, income, exam success and so forth that covers all needs.

2 Within each public sector agree the "monitoring" measures and automate them in as close to real time as possible. Success here might be that no more than 30 per cent of analyst time within core ministries, crown entities and public bodies is taken up with tasks associated routine performance monitoring. This seems to me essential. Nothing gives a clearer signal of the purpose of a role than how much time within it is spent doing what. If a majority of analyst time is spent using data for accountability and judgement, then they cannot be said to be using it for insight rather than control.

For this to happen implies a revised approach, where routine performance monitoring, whether internal as part of an organisation's management, or external as part of its accountability is automated as much as possible. This is discussed under the next section around the new skills that are required, but it is worth saying that this is not easy, and requires considerable investment in up-skilling together with the next of the requirements, which is...

3 Open source the measures and methods. One of the major emphases that I saw from all the Data Science Campus, and to a lesser extent also in Manchester, was a commitment to sharing and making open the various code and methods that they use. This is a critical building block for automating routine reporting. This builds confidence through transparency, replicability and comparability of method, and reduces time spent building from scratch. Indeed, if well documented and described, both code and statistical method can easily be adjusted to apply to similar problems in different data sets. There appear many ways to do this, GitHub, Slack channels, loomio groups. We already have the SIA data exchange; wouldn't it be fantastic if there was an all NZ public sector code sharing source?

4 Decision making in line with the higher purpose. This has to be rigorous. Organisations actually need to show the work that they undertake, the projects that they choose to pursue, and the analysis that is undertaken furthers the purpose. I suspect that this needs to be quite explicit. The approach adopted by the Data Science Campus seems a good example of this. The idea of a scoping phase which considers both the plausibility of the idea and the potential value of its implementation, followed by a discovery phase which considers how the work can actually be done (and whether when considered in greater detail whether the work remains viable), and only then

[3] https://data.govt.nz/assets/Uploads/NZ-Open-Data-Action-Plan.pdf

moving fully into delivery has a number of advantages. Obviously, it means that there are appropriate “break points” where a project can be stopped if it is unworkable or not providing the expected value. However, it also allows specific points where a project’s likelihood of advancing the higher purpose can be assessed and reviewed. The process demands of us that we ask not only “can we do this?”, but “should we?”

5 Evaluation – is what we have done really making a difference? The courage to question whether work is advancing the higher purpose is essential. However, our ability to do this can be questionable. In some instances, it’s very hard to know whether our actions are achieving the desired effect or not, or even if we can see the desired effect occurring. Was it what we did, some other, unconsidered, effect, some combination of both or even some indirect effects that proceeded from our intervention but weren’t the actual intervention itself?

To some extent this may not matter. The purpose of our evaluation is not so much to produce a peer reviewed journal quality demonstration of attribution, but to force us to clearly articulate what we are doing and what we expect to happen and ensure that we do what we say and check whether the expected consequences follow.

“Better stats than a programmer, better programming than a statistician” – the skills upgrade

So far, most of this writing has concentrated on the need for a change in mindset about how we view data. This is essential. However, to get down to brass tacks, we need a skills upgrade (or at least a more consistent spread and share of the skills we already have) to fully exploit the opportunities of the data we have.

Deriving insight from large data sets essentially takes us into the world of *data science* [ref]. This is variously described and understood, but a useful and common explanation is given in figure 1. Essentially when mathematical, statistical skills are combined with computer science skills *and* a deep understanding of the subject area, data science can occur. The last of these three pre-requisites is, of course, knowledge rather than a skill per se, and it is an interesting question as to whether one person can hold all of these skills and knowledge (a question considered in the next section).

Figure 1: Defining data science

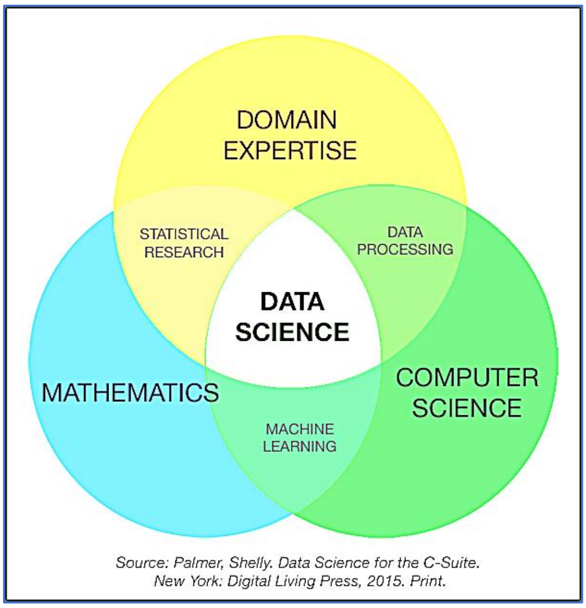

But if we return to the skills required across maths and computer science we quickly come to the somewhat glib but useful definition[4] of a person who is “better at statistics than most programmers and better at programming than most statisticians”.

For reasons that I expand on below about the importance of communication and influence I would add in data visualisation (which could be considered a subset of programming, although I would argue that it is distinct enough to be considered as a separate skill in and of itself). Talking with colleagues in the UK there was commonality about how people talked about their technical skills (basically R *plus* Python *plus* Tableau *plus* GIS) but these can be stripped down into a range of skills which have very specific purposes[5].

Programming skills

Programming skills were recognised widely as essential to support three

1 Management of large datasets to allow automation of routine reporting (to create time to use other data for insight)

Interestingly, on the Data Science Campus (where all manner of innovative and exciting cutting edge work was taking place) a number of interviewees identified that for the Office for National Statistics as a whole, one of the most critical things that could be done was replace manually intensive nested spreadsheets with robust coding for routine reports. This could literally save weeks of work each time the report was run (and ensure that the report was correct each time).

In talking with councils in Greater Manchester it’s clear that while this approach was their plan, it’s also hard to do. Issues with robust IT systems, the pace of public sector “reform” since 2010, and the need to achieve stability have made it hard to automate reporting so that more analysts could work in the research and insight space, have all created barriers to this transformation. The consequence is that the majority of insight work is carried out by a small specialist team.

This point is extended by the concept of robotic process automation, the idea of which is that routine, predictable processes are automated programmatically which has several advantages. Above all it frees up time for the types of high value add insight generation, but it also reduces error, and increases uniformity of process and output and should allow faster processing times for routine interactions with the public. An additional insight is that this “lower” end of the continuum of automation has fewer of the ethical issues associated with Artificial Intelligence and machine learning discussed below.

2 Complex querying and analysis of (often linked) structured data sets

This describes much of the work that is already undertaken through IDI. In particular, the ability to create cohorts of people with particular characteristics and study the pathway of their life experiences following interactions with the state (and compare these with similar cohorts with different interactions, or different cohorts with similar interactions) has tremendous power in understanding the effects of interventions (including unintended consequences).

[4] Variously attributed

[5] It goes without saying that there are a range of skills around data governance, data management, problem definition, collaboration and teamwork which are *sine qua non* for this to be successful. In general, my experience is that these are present inside the NZ public sector.

Reviewing IDI projects[6] (and being involved in a couple myself) shows that this is a common use of the IDI sets, so clearly these skills are not unusual in New Zealand, however, it is not clear that they are as widespread, or as widely used, as they need to be.

3 Harvesting and categorising of unstructured data

This is quite a different way of going about things. It recognises that there are enormous amounts of data which is unstructured and can be exploited if only it can be gathered and organised. Programming to do this, often through Python is essential to make this viable. Examples include the creation of Consumer Prices Index based on web prices for 33 food items from three different supermarkets in the United Kingdom; and scraping and sentiment analysis of Twitter comments about NHS hospitals in order to predict poor outcomes[7].

For all three of these purposes, programming skills in packages such as SAS, R or Python are critical. There is an element of 'horses for courses' in choosing which to use. SAS is already common in the New Zealand public sector and handles very large data sets well, while R, in particular, has the advantage of being open source, which is not only economical, but also provides for a huge public library of reusable code.

Statistical skills

When using the phrase statistical skills in this paper I use it in its proper sense referring to the particular skills used in the science of statistics, the inferential skills associated with the profession of statistics rather than the vernacular meaning of "pertaining to data". In fact, I would argue that the conflation of the two meanings has had a relatively pernicious effect, downplaying both the complexity of the skills needed and their relative scarcity.

In this section I concentrate on what I perceive to be a chain of interrelated sets of statistical skills and knowledge that are required. These are not comprehensive, and they specifically don't cover aspects of mindset such as logic, rigour and scepticism that underpin good analysis. Neither do they cover basic statistical skill.

However, the chain of skills that have been commonly discussed throughout the conversations that I have had can be described as follows

1 Identifying and analysing of relationships between features – typical techniques are various forms of regression and classification techniques such as discriminant analysis

2 Describing and explaining complex systems and interactions – typical techniques include decision tree methods, simulation techniques, and synthetic population generation

3 Predictive analytics – various forms of predictive modelling, including machine learning at its more sophisticated

Visualisation

I would argue that the ability to visualise data effectively, while almost always requiring knowledge of software packages, is in and of itself a separate skill that needs to be learned. As discussed below, the role of story-telling and effective influence is essential and effective visualisation of data can be very useful tool for this, for at least two reasons. The dashboard approach which allows common

[6] https://cdm20045.contentdm.oclc.org/digital/collection/p20045coll17/search/

[7] A Griffiths, M Leaver, Wisdom of patients: predicting the quality of care using aggregated patient feedback, BMJ Quality and Safety Online First, Sept 2017 doi:10.1136/bmjqs-2017-006847

reports to sliced and diced according to different populations in order to gain insight from routine data has tremendous managerial power and allows non-analysts to ask analytic questions ("who does this refer to?". "how is this changing?", "is it the same for everyone?", "how do we compare with the average") without needing analytic support to do this[8]. This helps target managerial attention and action effectively. Less noticed though is that the right visualisation may itself shift thinking about problems. For example, given that many policy problems are associated with progression through systems of populations and individuals, something like a Sankey diagram[9] can be invaluable to understanding the dynamics of flow, and yet these are very infrequently used in policy analysis.

Some of the skills required here are about use of Tableau or Power BI or R shiny or whatever, but an understanding of the principles of visual display of quantitative data is also essential. For this the work of Edward Tufte[10] is perhaps the set text.

Building the cadre

One of the strengths of the New Zealand public sector (and indeed many national public services) is transferability of skills so that an individual can work across different sectors drawing on common base of skills. However, the shadow side of this flexibility is an unstated but quite common public sector belief that clever people can turn their hand to anything with minimal training. The level of sophistication of skills required as described above are unlikely to come from "on the job training", certainly at the scale that we need.

One consequence of the public sector using data primarily as a mechanism for control, is that it has been necessary to populate public sector analytic departments with people with intelligence and talent but not the specific training that the skill set requires. In central agencies this tends to mean policy analysts with a numerate bent[11], while in delivery organisations the equivalents are often IT professionals with an interest in the organisation's end goal or professionals and administrators with good excel skills. In and of itself this is fine, but it does imply a need for some focused development to be able to deploy the skills listed above.

The question then is "how to create the cadre?"

My conversations with the Data Science Campus were particularly of value here, as they have as a corporate goal a role of producing data scientists at scale and speed[12]. To support this, they have a number of approaches in place specifically designed to build the workforce.

These include formal academic qualifications such as an MSc, but there are two innovations which I think would have particular applicability in a New Zealand context. The first is creation of apprentice roles for school leavers, an innovation now being extended to a three-year program that leads to a bachelor's degree with an expectation of 80 percent of learning being "on the job". I was profoundly impressed by this an approach, and indeed profoundly impressed by the apprentices I met. It is a useful challenge to the unthinking bias that a degree is a pre-requisite to starting a professional career, but it also means that skills get taught in the context of a role rather than in academic

[8] An excellent recent example of this capacity to dig down into quite complex data is in the realm of falls prevention https://public.tableau.com/profile/hqi2803#!/vizhome/FallsFracturesOutcomesFramework/Landing

[9] https://en.wikipedia.org/wiki/Sankey_diagram

[10] Edward R. Tufte: The Visual Display of Quantitative Information, Graphics Press, Connecticut, 1983

[11] Full disclosure: this is, more or less, my career path

[12] "Through the Data Science Campus, produce 500 qualified data analysts for government by 2021"

abstraction. For example, the very first skill that the data scientist apprenticeship recognises is the ability to identify and clarify organisational problems and reformulate them as data science problems. The real-world grip implicit in this statement is striking.

The second approach is that of the "Accelerator" programmes[13]. These are essentially free mentoring programs to the broader public service where the participants bring a data science project with them and are provided with tools (a MacBook, access to open source software for natural language processing, machine learning, visualisation, geospatial analysis etc), mentorship from a campus data scientist and three months to deliver. The approach builds skills in real world situations (as well as addressing local problems) and leaves behind a group of champions (in some cases those who've been through the program have then created their own, local programs).

These approaches, practical, focused in real world problems and relatively low cost seem to me to be relatively transportable to our context.

Defining the roles – because unicorns don't exist

One of the corollaries of the Venn diagram of data science above is that the mix of roles required to extract the value from data is wide ranging, and the likelihood is that one person is not going to be able to play all of them. Indeed, the reflection of how hard it is to find one person who can has given rise to a widely shared concept of the "data science unicorn"[14] - a mythical but much sought after creature.

As it is unlikely that one person will naturally be able to do everything necessary to exploit data for insight, determining the roles that need to be played to do this is an urgent task. Across conversations with colleagues in NZ and the UK and more broadly in the literature, there are various categorisations that people have used, and in most instances how to apply domain expertise is the central concern of these.

In essence, there are a range of roles around:

- translating specific policy problems into the language of data and analysis,
- interpreting the meaning of and implications of results,
- iterating between these two modes of operation, and finally,
- being able to tell compelling stories with the results

An alternative categorisation is made by Open Data Manchester[15] (see figure 2) and provides some particular insight into the roles that need to be played, as follows:

Scout – the person who finds the data necessary to provide insight into a problem

Storyteller – the person who draws out the meaning of the data for a wider audience

Analyst –analogous to the statistician skills above

Engineer - analogous to the programming skills above

Designer – analogous to the visualisation skills above.

[13] https://www.gov.uk/government/publications/data-science-accelerator-programme/introduction-to-the-data-science-accelerator

[14] A quick google search reveals over 8 million references – not bad for a concept that's about 5 years old.

[15] http://www.opendatamanchester.org.uk/

While the engineer, analyst and designer roles quite clearly relate to the data itself and are in some ways inseparable from the skills they exercise, the other two roles are much more concerned with relationships. These include the relationship between the data and the problem it seeks to solve, the analytic team and its policy "client", and the data owners and its broader audience. And these roles are different and need very different skills.

Given the rarity of naturally occurring "unicorns", how best can we make sure that all the required roles are being filled? Beyond bromides about "good communication" (which as advice goes manages to be entirely correct and entirely useless simultaneously) some approaches that organisations have adopted include the following.

1 Work closely with the partner at all stages of the project. In particular, design projects so that they are built around constant communication between the analyst and client so that the process itself creates the role. One corollary of this is that more intelligent and appropriate forms of project management such as Agile are used to ensure delivery[16].

2 Building the capacity to undertake the translator roles especially into the data scientist's skill set. The Data Science degree developed by the Data Science Campus, for example, identifies that the first skill of data scientist is to "identify and clarify problems and organisation aces, and reformulate them into Data Science problems", which might work as a definition of the translator role.

3 Recognising the different strengths already in data teams and bring bringing together teams based upon complementary strengths. One of the first tasks that the Greater Manchester analysts network undertook was a skills audit for precisely this reason. In a slightly, light-hearted vein, Open Data Manchester provide the following self-assessment tool to help identify comfort in playing the different roles.

Figure 2: Data roles self-assessment (courtesy of Open Data Manchester)

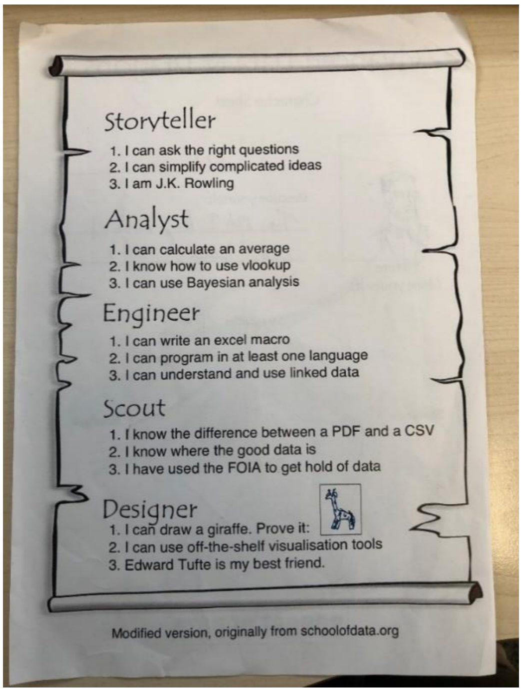

16 This is also widely recognised as a key defence against data projects becoming self-indulgent and unwieldy

Where we belong - social positioning of the organisation

This fourth prerequisite links to both higher purpose and the ability to influence the system. It is somewhat tricky to define tightly, but encompasses a range of related issues:

- The data team's conceptualisation either of its role within its organisation or its relationship with outside organisations with which it partners.
- The team or organisation's understanding of its position insider the broader system, what it can uniquely do, and the context which creates opportunities or barriers for it to further its mission.
- A range of ethical issues which can be aggregated as an understanding of the risk of doing harm (even the most careful and ethically conscious of organisations may do so inadvertently)

1 interactions and relationship

Understanding how one fits into the system is a prerequisite *how* one goes about achieving one's mission, in much the same way that higher purpose determines what one's mission is. This is highlighted very clearly in the contrasts between a strategic body such as the DSC compared with the "coal face" agencies in Greater Manchester.

DSC emphasise partnership with others as their way of working. Not only is this essential to gain access to both the problems that need to be understood and the data to analyse, it is a prerequisite for the work to influence policy and its implementation. In fact, a clear risk for central strategic agencies is that this partnership is not well enough managed and that, as a consequence, the commissioning agency has insufficient commitment to an end result.

For the local government departments, the issue is much more about who does this type of work within analytic teams and how they then interrelate with the broader team and the organisation more widely.

In both instances there is a sub issue about how to avoid a type of "golden child" syndrome, where the development is seen as a small group of privileged analysts getting to do "cool stuff" with data outside of the mainstream of work, breeding resentment and resistance. The importance of robust oversight and project management while allowing the freedom to experiment was emphasised among both groups.

The freedom to experiment creates an interesting paradox of failure. If all projects are a success then the likelihood is that not enough risks are being taken, and opportunities for insight are being missed. Yet a reputation for using techniques that can fail can be disastrous, especially in risk averse cultures. (Conversely never taking a risk and missing opportunities for improvement can be reputation enhancing.)

2 Context

Context can be wide ranging. It can relate to the policy, financial and governance in which an organisation operates. It can relate to the political culture, organisational mandate, media or even physical environment.

Such is the breadth of what can constitute context for an organisation working with data it is probably more useful to cite examples from the two organisations I visited than attempt a comprehensive analysis of these.

For example, local government analysts highlighted to complexities of bringing together data from the two large and arguably competing bureaucracies of health and social care. They also reflected on the way in which normative power structures can kill off data innovation; who 'owns' the data and who makes decisions about its use can be major barriers. Similarly, the data itself can become corrupted by local circumstances, especially when funding rests on getting the right results.

On a more positive note, the Data Science Campus reflected on the value of having their own physical space, separate and differently designed to the rest of the Office for National Statistics. They also emphasised the need to establish their reputation and brand as having an exciting, credible and important mission in order to attract the right people to work for them.

3 Ethical issues

The dangers of constructing or implementing policies from data which do harm to vulnerable populations is starting to be considered more widely. There are obvious risks of disclosure in using micro level data through careless presentation of results. However, perhaps more pressing and harder to solve are analyses that entrench inequity (one would hope unintentionally, but some instances are egregious enough to make one wonder…). A wide range of examples, especially from the US have been chronicled by Propublica[17].

This subject is complex enough to merit papers all of its own. Some of the analysis of the problem and potential solutions identified by the Turing Institute[18] for example are both extremely interesting and potentially of great value. However, in this context the issue for an organisation using data is to consider the potential risks and implications. There are some approaches to avoid some of the worst risks of algorithmic *un*fairness, for example ACC have made the decision never to reject a claim based on algorithm alone (anything that does not meet criteria for immediate approval goes to claims assessors)[19]. However, whether such a low risk approach could be used by all agencies in all circumstances is not clear.

On the importance of driveshafts – influencing beyond the data bubble

The worry about investing heavily in data and the people to use it is that we create Rolls Royce engine which is not connected to the "wheels" of delivery. This seems an odd thing to say as government departments are increasingly concerned to be (or at least *say* that they are) data-driven. Certainly, agencies that provide services to government increasingly complain about the burden of data collection (which creates an opportunity cost which limits their capacity to provide the service that they are being paid to provide).

Yet research by the Behavioural Insights team (BIT) [20] argues that in fact data is frequently nothing more than a support to the decisions predetermined by heuristics and political pressures, rather than a basis for decision making. Policy implementation is driven by biases in what policy makers notice, how they make decisions about policy and how it is executed (and one of these failures is precisely that inappropriate measurement creates an illusion of control that isn't real).

This returns us to the issue of data for insight not data for control purposes. When data is used as a mechanism for control, metrics are chosen to fit a predetermined frame, and results assume

[17] https://www.propublica.org/series/machine-bias
[18] https://www.turing.ac.uk/research/research-projects/counterfactual-fairness
[19] https://www.acc.co.nz/about-us/news-media/latest-news/claims-approval-process-documents-released/
[20] https://www.behaviouralinsights.co.uk/publications/behavioural-government/

causation. In the development of policy this leads to confirmation bias. In delivery, the data derived leads to both group reinforcement and feeds optimism bias.

This is made more complex because technocratically rational data driven positions can fail to understand either political context or system complexity. Thus, policies which are rational on their own terms can fail when they come into contact with political realities (BIT argue that the poll tax introduced in the 1990s in the UK is one example of this), while there are other policies of such symbolic importance in signalling political priority and “direction of travel” that their implementation is a success even if there is little unambiguous evidence that they achieved their stated aims (BIT cite reducing class sizes in schools). This sense that “data isn’t everything” represents a considerable barrier to successfully implementing data driven policy.

BIT propose two approaches for avoiding this problem. For the wider public service they recommend training in “de-biasing” helping public servants to identify and avoid where biases enter into their thinking. This has been demonstrated to reduced biased judgements and increase accuracy of predictions made by policy analysts.

More proactively, they have identified a set of preconditions likely to help behavioural science influence policy.[21] These apply equally well to data science teams. These are as follows:

Administration: to be a success, a central team needs to have people on it who understand the machinery of government. This helps to gain traction inside government

Politics: in the early days especially, it is very important to have senior political support, both from politicians and very senior public servants.

People: BIT think this to be the most important lesson of all. You are nothing without your people. So being able to, for example, control your own recruitment is essential for a team that requires specialist skills.

Location: Physical location is surprisingly important. Being physically close to your political sponsors is both symbolically important and helps cement ties.

Experimentation: embedding a culture of testing and trialling as a means of being able to demonstrative efficacy is essential to success.

Scholarship: alongside the administrative ties, strong links with academia are essential.

So, what does this all look like?

When I started on this journey of discovery, I had a hope that there would be some clear prescription for what this might imply in terms of team structure, leadership and incentivisation; while my expectation, based upon experience, was that I would end up saying something like “well, it all rather depends”.

I find myself with a triumph of experience over hope.

It does all rather depend. Different people that I have spoken to have different team structures, different leadership and communication models, different staffing structures, different routes in. And this is OK because these reflect both the realities in which organisations work and the mission that they are tasked with accomplishing.

[21] https://www.behaviouralinsights.co.uk/uncategorized/the-global-spread-of-behavioural-insights-conditions-for-success-of-a-central-unit/

But there are some commonalities that can be drawn out – this paper sets some of these out in detail – their application requires specific intentional thought. Appendix 1 sets out a checklist of how to think about these things.

Perhaps the possibilities of Process Automation are themselves the best analogy to draw here. Just as these create space for additional thought rather than replace the need for humans to do it; the outputs of this study may in the end provide support for ways of thinking about how we exploit the new possibilities of data rather than a blueprint for success.