

Chapter 1

Computational Social Science for Public Policy



Helen Margetts and Cosmina Dorobantu

Abstract Computational Social Science (CSS), which brings together the power of computational methods and the analytical rigour of the social sciences, has the potential to revolutionise policymaking. This growing field of research can help governments take advantage of large-scale data on human behaviour and provide policymakers with insights into where policy interventions are needed, which interventions are most likely to be effective, and how to avoid unintended consequences. In this chapter, we show how Computational Social Science can improve policymaking by detecting, measuring, predicting, explaining, and simulating human behaviour. We argue that the improvements that CSS can bring to government are conditional on making ethical considerations an integral part of the process of scientific discovery. CSS has an opportunity to reveal bias and inequalities in public administration and a responsibility to tackle them by taking advantage of research advancements in ethics and responsible innovation. Finally, we identify the primary factors that prevented Computational Social Science from realising its full potential during the Covid-19 pandemic and posit that overcoming challenges linked to limited data flows, siloed models, and rigid organisational structures within government can usher in a new era of policymaking.

1.1 Introduction

These are exciting times for social science. Large-scale data was formerly the province of the physical and life sciences, while social science relied mostly on qualitative data or survey data to understand human behaviour. The data revolution from the 2010s onwards, where huge quantities of transactional data are

H. Margetts (✉) · C. Dorobantu
The Alan Turing Institute, London, UK

Oxford Internet Institute, University of Oxford, Oxford, UK
e-mail: helen.margetts@oii.ox.ac.uk; cdorobantu@turing.ac.uk

generated by people's online actions and interactions, means that for the first time, social scientists have access to large-scale, real-time transactional data on human behaviour. With this influx of data, social scientists can and need to develop and adapt computational methods for analysis of large-scale social data. Computational Social Science—the marriage of computational methods and social sciences—can transform how we detect, measure, predict, explain, and simulate human behaviour. Given that public policy is about understanding and potentially changing the world outside—society and the economy—Computational Social Science is well placed to help policymakers with a wide range of tasks, combining as it does computational methods with social scientific lines of enquiry and theoretical frameworks. Given the struggle that social science often has to demonstrate or receive recognition for policy impact (Bastow et al., 2014), CSS might act as a channel for social science to be appreciated in a policy context. This chapter examines how CSS might assume an increasingly central role in policymaking, bringing social science insight and modes of exploration to the heart of it.

The seminal article on CSS (Lazer et al., 2009) laid out how the capacity to analyse massive amounts of data would transform social science into Computational Social Science, just as data-driven models and technologies had transformed biology and physics. 'We define CSS as the development and application of computational methods to complex, typically large-scale, human (sometimes simulated) behavioural data [. . .] Whereas traditional quantitative social science has focused on rows of cases and columns of variables, typically with assumptions of independence among observations, CSS encompasses language, location and movement, networks, images, and video, with the application of statistical models that capture multifarious dependencies within data' (Lazer et al., 2009, p. 1060). Although there is no definitive list of the methodologies that would fall into the category of CSS, it is clear that agent computing, microsimulation, machine learning (ML), complex network analysis, and statistical modelling would all fall into the field. We might also add large-scale online experimental methods and some of the ethical thinking that should accompany the handling of large-scale data about human behaviour.

Lazer et al.'s early article did not discuss how policymaking or government might also be transformed, although the second article 12 years on (Lazer et al., 2021) emphasised the need to articulate how CSS could tackle societal problems. Given CSS's emphasis on data and data analysis, the transformative potential of Computational Social Science for policymaking is huge. Traditionally, governments have made little use of transactional data for policymaking (Margetts & Dorobantu, 2019). That is not surprising, given that bureaucratic organisation from the earliest forms of the state relied on 'the files' for information (Muellerleile & Robertson, 2018). Paper-based files offer the capability to find individual pieces of data but generate no usable data for analysis. Likewise, the large-scale computer systems which gradually replaced these files from the 1950s onwards in the largest governments also had no capacity to generate usable data (Margetts, 1999). For decades, governments' transactional data resulting from their interactions with citizens languished in 'legacy systems', unavailable to policymakers. During this period, data and modelling existed in government but relied on custom built 'official

statistics’ or performance indicators, or long-running annual surveys, such as the UK ‘British Crime Survey’. Only with the internet and the latest generation of data-driven models and technologies has there been the possibility for policymakers to use large-scale transactional data to inform decision-making.

This chapter outlines key policymaking tasks for which CSS can be used: detection, measurement, prediction, etiology,¹ and simulation. It discusses how CSS needs to be ‘ethics-driven,’ revealing bias and inequalities and tackling them by taking advantage of research advancements in ethics and responsible innovation. Then, the chapter examines how the potential of CSS tools has been highlighted in the pandemic crisis but also how CSS failed to realise this potential due to weaknesses in data flows, models, and organisational structures. Finally, the chapter considers how CSS might be used to tackle future crisis situations, renewing the policy toolkit for more resilient policymaking.

1.2 Detection

Detection is one of the ‘essential capabilities that any system of control must possess at the point where it comes into contact with the world outside’ (Hood & Margetts, 2008). Government is no exception, needing to understand societal and economic behaviour, trends, and patterns and to calibrate policy accordingly. That includes detection of unwanted (or less often, wanted) behaviour of citizens and firms to inform policy responses.

Data-intensive technologies, such as machine learning, lend themselves very well to the performance of detection tasks. Advances in machine learning over the past decade make it a powerful tool in the analysis of both structured and unstructured data. Structured data refers to data points that are stored in a machine-readable format. ML is well suited to the performance of detection tasks that rely on structured data, such as pinpointing fraudulent transactions in large-scale financial data. The progress made by researchers and practitioners in the fields of natural language processing (NLP) and computer vision now also makes ML well suited to the analysis of unstructured data, such as human language and visual data (Ostmann & Dorobantu, 2021). ML can perform detection tasks that were outside the realm of possibilities in earlier decades due to our inability, in the past, to process large quantities of structured and unstructured data.

A good illustration of where Computational Social Science and policymakers can work side-by-side to detect unwanted behaviour relates to online harm. Online harm is a growing problem in most countries, including (but not limited to) the gen-

¹ Etiology is a term often used in the medical sciences, meaning ‘the cause, set of causes, or manner of causation of a disease or condition’. (Oxford Languages from Oxford University Press). Here we use it more broadly, to refer to the cause, set of causes, or manner of causation of phenomena of interest within the social sciences.

eration, organisation, and dissemination of hate speech, misinformation, misleading advertising, financial scams, radicalisation, extremism, terrorist networks, sexual exploitation, and sexual abuse. Nearly all governments are tackling at least some of these harms via a range of public agencies. Criminal justice agencies need to track and monitor the perpetrators of harm; intelligence agencies need to scrutinise security threats, while regulators need to detect and monitor the behaviour of a huge array of data-powered platforms, particularly social media firms.

How can Computational Social Science help policymakers? A growing number of computational social scientists are focusing on the detection of harmful behaviour online, seeking to understand the dissemination and impact of such behaviour, which is a social as well as a computational task. Machine learning classifiers need to be built, and this is a highly technical task, requiring cutting edge computer science expertise and facing huge challenges (see Röttger et al., 2021). But it is those with social science training that are comfortable dealing with the normative questions of defining terms such as ‘hate’. And it is social scientists who are able to explore the motivations behind harmful online behaviour; to understand the differential impacts of different kinds of harm (e.g., misinformation has different dynamics from hate speech, see Taylor et al., 2021); and to explore how we can build distinct classifiers for different kinds of online harm or different targets of harm, such as misogyny (Guest et al., 2021) or sinophobia (Vidgen et al., 2020). By bringing together the development of technical tools and the rigour and normative stance of the social sciences, Computational Social Science offers a holistic and methodologically sound solution to policymakers interested in tackling online harm.

Regulation for online safety is a key area where CSS is uniquely qualified to help. Regulators need to develop methodological expertise but often struggle to keep ahead of the perpetrators of unwanted online behaviour and the massive platforms where these harms play out. While CSS expertise is growing in this area, the platforms themselves have incubated parallel streams of in-house research, with different motivations, confidentiality, secrecy, and lack of data sharing preventing knowledge transfer between the two. This leaves an important role for academic researchers, working directly with regulators to help them understand the ‘state-of-the-art’ research in promoting online safety.

1.3 Measurement

Another key capability of government is measurement. Policymakers need to be able to monitor and track societal and economic trends and patterns in order to understand when interventions are needed.

The technologies that were available to us prior to the data revolution limited our ability to collect, store, and analyse data. These technological limitations meant that in the past, policymakers and academic researchers alike were at best able to measure socio-economic phenomena imprecisely and at worst unable to measure them at all. For example, policymakers and researchers have been trying

for decades to understand visitation rates at public parks (see, e.g., Cheung, 1972). This understanding is needed for a range of policy interventions, from protecting green spaces and increasing investment in parks to driving up community usage. But what seems like a simple metric, the number of visitors to a park, has been difficult to produce in practice. The solution preferred by many local authorities has been to hire contractors and ask them to stand at the entrance of a park and count the number of people going in. This solution has obvious limitations: it is costly, it can only measure park attendance for limited periods of time, it is prone to measurement error, and it fails to capture characteristics of the people visiting the park—to name only a few.

Complex socio-economic phenomena are even more difficult to measure. Firms, consumers, and policymakers are increasingly worried about inflation, a phenomenon that threatens the post-pandemic economic recovery. Yet despite the fact that so many eyes and newspaper headlines focus on the consumer price index, few know the difficulties of collecting and generating it. In the UK, for example, the Office for National Statistics calculates the Consumer Prices Index. The index largely rests on the physical collection of data in stores across 141 locations in the UK. At a time when we needed precise inflation measures the most, during the Covid-19 crisis, the data collection efforts for the Consumer Prices Index were severely affected by store closures and social distancing measures. Furthermore, the labour-intensive nature of collecting and generating the Consumer Prices Index means that it cannot be, with its current design, a real-time measure. National statistical offices usually publish it once a month with the understanding that it reflects the reality of a few weeks back.

Computational Social Science allows new opportunities to measure and monitor socio-economic phenomena—from park usage to inflation. Recent research uncovered the value of using social media data and mobile phone app data to measure park visitation (see, e.g., Donahue et al., 2018; Hamstead et al., 2018; Sinclair et al., 2021; Suse et al., 2021). Attempts to create real-time measures of inflation go back more than a decade. In 2010, Google’s chief economist, Hal Varian, revealed that the company was working on a Google Price Index—a real-time measure of price changes calculated by monitoring prices online. Although Google never published this measure, it hints at the possibilities of using computational methods and economic expertise to move beyond the inflation measures that we have today.

More generally, Computational Social Science could facilitate a wholesale rethinking of how we measure key socio-economic indicators. As Lazer et al. (2021) reflected in their study of ‘Meaningful measures of human society in the twenty-first century’:

Existing measures of key concepts such as gross domestic product and geographical mobility are shaped by the strengths and weaknesses of twentieth century data. If we only evaluate new measures against the old, we simply replicate their shortcomings, mistaking the gold standard of the twentieth century for objective truth.

Traditional social science methods of data analysis tend to perpetuate themselves. Survey researchers, for example, are reluctant to relinquish either long-running

surveys or questions within them. This means that over time, surveys become longer and longer and increasingly unsuited to measuring behavioural trends in digital environments (e.g., asking people what they did online is a highly inaccurate way of determining digital behaviour compared with transactional data). Computational Social Science gives us the ability to improve our measurements so that everything—from basic summary stats to the most sophisticated measures—can move away from having to rely on old measurements that are limited by the technologies and data that were available decades ago.

1.4 Prediction

Another tool that Computational Social Science has to offer to policymakers is predictive capability. Machine learning is increasingly used within the private sector to perform prediction and forecasting tasks, as it is well suited to the performance of these tasks. Governments and public sector organisations in general do not have a good record on forecasting and prediction, so this is another area where CSS can add to policymakers' toolkit. Policymakers can use machine learning to spot problematic trends and relationships of concern before they have a detrimental impact and to predict points of failure within a system. One of the most common uses of machine learning by local and central governments is to predict where problems are most likely to arise with the aim of identifying 'objects' (from restaurants and schools to customs forms) for inspection and scrutiny. The largest study on the use of machine learning in US federal government provides the example of the US Food and Drug Administration, which uses machine learning techniques to model relationships between drugs and hepatic liver failure (Engstrom et al., 2020, p. 55), with decision trees and simple neural networks used to predict serious drug-related adverse outcomes. The same agency also uses regularised regression models, random forest, and support vector techniques to construct a rank ordering of reports based on their probability of containing policy-relevant information about safety concerns. This allows the agency to prioritise for attention those that are most likely to reveal problems.

Machine learning can also be used to predict demand, helping policymakers plan for the future. When used in this way, it can be a good way to optimise resources, allowing government agencies to be prescient in terms of service provision and to direct human attention or financial resources where they are most required. For example, some police forces use machine learning to predict where crime hotspots will arise and to anticipate when and where greater police presence will be needed. Recent studies on the use of data science in UK local government (Bright et al., 2019; Vogl et al., 2020) estimate that 15% of UK local authorities were using data science to build some kind of predictive capability in 2018, when the research was carried out.

The use of machine learning for prediction in policymaking is controversial, however. Some have argued that the predictive capacity of Computational Social

Science brings tension to the field, sitting happily with the epistemological aims of computer scientists, but going against the tradition of social science research, which prioritises explanations of individual and collective behaviour, ideally via causal mechanisms (Hofman et al., 2021, p. 181). Kleinberg et al. (2015) argue that some important policy problems do benefit from prediction alone and that machine learning can generate high policy impact as well as theoretical insights (Kleinberg et al., 2015, p. 495). But this use of machine learning generates important ethical questions of fairness and bias (discussed below), as the use of the COMPAS (Correctional Offender Management Profiling for Alternative Sanctions) system for predictive sentencing in the USA has shown (see Hartmann & Wenzelburger, 2021). Furthermore, as Athey (2017) explains, many of the prediction solutions described (e.g., in health care and criminal justice) require some kind of causal inference to achieve payoffs, even where prediction is most commonly cited as beneficial, such as the identification of building sites or other entities for inspection and scrutiny. Overall, she concludes, multidisciplinary approaches are needed that build on the development of machine learning algorithms but also ‘bring in the methods and practical learning from decades of multidisciplinary research using empirical evidence to inform policy’. In a similar vein, Hofman et al. (2021) make the case for integrative modelling, developing models that ‘explicitly integrate explanatory and predictive thinking’, arguing that such an approach is likely to add value over and above what can be achieved with either technique alone and deserves more attention than it has received so far.

1.5 Etiology

The possibilities of detection, measurement, and prediction that CSS methods afford to tackle policy problems do not obviate the need for understanding the underlying causes of observed behaviours, as discussed in the preceding section. Etiology is particularly important when policymakers try to understand human behaviour in digital settings, where they need also to understand how the digital context, including the design of platforms and the algorithms they use, drive behaviour. Wagner et al. (2021) observe that in the ‘algorithmically infused society’ in which we now live, algorithms shape our behaviour in many contexts: shopping, travelling, socialising, entertainment, and so on. In such a world, the data that we derive from platforms like Twitter gives us useful clues about our interactions, but the social sciences is the only lens through which we can learn to separate what is ‘natural’ human behaviour and what is algorithm-induced human behaviour. The social sciences are also the domain that gives us the theoretical starting point for re-examining frameworks, models, and theories that were developed when algorithms were not a prevalent part of our lives. We need to understand both how algorithmic amplification (e.g., via recommender systems or other forms of social information) influences relationship formation, while also understanding how social adaptation causes algorithms to change. This understanding is particularly important

for regulators, who need to know how digital platforms are influencing consumer preferences and behaviour (e.g., through targeted advertising) and which elements of the behaviours we notice online are attributable to the algorithms themselves. Scientific researchers need to develop this kind of expertise. Although streams of research are being developed within, for example, social media companies, around issues of content moderation and algorithm design, the primary aim of this work is to limit reputational damage. The companies themselves have little motivation to invest in programmes of research that uncover the organisational dynamics of online harms or the impact of such harms on different groups of citizens. They also have limited incentives to share the findings of such research, even if they decide to carry it out.

CSS can also help with etiology via experimental methods. Early social science experiments used survey data or laboratory-based experiments, which were expensive and labour-intensive and quickly resulted in small numbers problems. In contrast, online randomised controlled trials based on large-scale datasets can operate at huge scale and in real time. Such behavioural insights have been used by governments, for example, testing out the effects of redesigning letters and texts urging people to pay tax on time (Hallsworth et al., 2017). Large-scale digital data also offers the possibility of identifying ‘natural experiments’ (Dunning, 2012) in policy settings, where some disruption of normal activity at a point in time or in a particular location occurs, and the data is analysed after the disruption, as an ‘as if random’ treatment group. An example is provided by Transport for London’s analysis of their Oyster card data to understand the effects of a 2014 industrial dispute which led to a strike of many of the system’s train drivers (described in Dunleavy, 2016). During the strike, millions of passengers switched their journey patterns to avoid their normal lines and stations hit by the strike. Larcom et al. (2017) examined Oyster card data for periods before and after the strike period, linking journeys to cardholders. They found that 1 in 20 passengers changed their journey, and a high proportion of these stayed with their new journey pattern when normal service resumed, suggesting their new route was better for them. The findings suggested that Tube travellers only ‘satisfice’ and had originally gone with the first acceptable travel solution that they found, later settling on the new route because it saved them time. The analysts also showed that the travel time gains made by the small share of commuters switching routes as a result of the Tube strike more than offset the economic costs to the vast majority (95%), who simply got disrupted on this one occasion. So the strike led to net gains, suggesting that possible side benefits of disruptions might be factored in by policymakers when making future decisions (like whether to close a Tube line wholly in order to accomplish urgent improvements (Dunleavy, 2016)).

Natural experiments like this can be hard to systematise or find. But large-scale observational data can be used to identify causal inference even where there is no identifiable ‘as if random’ treatment group or no counterfactual control group. Large-scale data analysis offers ‘New tricks for Econometrics’ (Varian, 2014), for example, where datasets are split into small worlds, creating artificial ‘control groups’ via a predictive model based on a function of past history and possible

predictors of success. CSS methods have developed hugely in this area, especially in economics. Athey and Imbens (2017) discuss a range of such strategies, including regression discontinuity designs, synthetic control and differences-in-differences methods, methods that deal with network effects, and methods that combine experimental and observational data—as well as supplementary analyses (such as sensitivity and robustness analysis)—where the results are intended to convince the reader of the credibility of the primary analysis. They argue that machine learning methods hold great promise for improving the credibility of policy evaluation, particularly through these supplementary strategies.

1.6 Simulation

Another way in which CSS can tackle policy issues is through the development of simulation methods, allowing policymakers to try out interventions before implementing the measures in the real-world and having them give rise to unintended and unanticipated consequences. As noted above, policy choices need to be informed by counterfactuals: if we implemented this measure—or didn't implement it—what would happen?

An increasing range of modelling approaches can now be used for simulation, including complex network analysis and microsimulation, involving highly detailed analysis of, for example, traffic flows, labour mobility, urban industrial agglomeration patterns, or disease spread. One modelling approach that is gaining popularity with the growing availability of large-scale data is agent computing. Agent-based models (ABMs) have been used to study socio-economic phenomena for decades. Thomas Schelling was among the first to use agent-based modelling techniques within the social sciences. In the early 1970s, he published a seminal paper that showed how a simple dynamic model sheds light on how segregation can arise from the interplay of individual choices (see Schelling, 1971). But models like Schelling's—and many others that followed—were 'toy models': formal models without any real-world data to ground them in the socio-economic reality that they were meant to study. In contrast, the agent computing models used now are based on large-scale data, which transforms them into powerful tools for researchers and policymakers alike. Rob Axtell, one of the pioneers of Computational Social Science, recently developed a model of the US private sector, in which 120 million agents self-organise into 6 million firms (Axtell, 2018). Models like Axtell's are extremely powerful tools for studying the dynamics of socio-economic phenomena and carrying out simulations of complex systems, from economies to transport networks. Today's agent computing models can also be used in combination with machine learning methods, where the models provide a practical framework to combine data and theory without constraining oneself with too many unrealistic *a priori* assumptions about how socio-economic systems behave, such as 'fully rational agents' or 'complete information'.

An agent computing model consists of individual software agents, with states and rules of behaviour and large corpuses of data pertaining to the agents' behaviour and relationships. Running such a model could theoretically amount to instantiating an agent population, letting the agents interact, and monitoring what happens; 'Indeed, in their most extreme form, agent-based computational models will not make any use whatsoever of explicit equations' (Axtell, 2000, p. 3). But models usually involve some combination of data and formulae. Researchers have started to explore the possibilities of 'societal digital twins' (Birks et al., 2020), a combination of spatial computing, agent-based models, and 'digital twins'—virtual data-driven replicas of real-world systems that have become popular for modelling physical systems, in engineering or infrastructure planning, for example. Such 'societal' twins would use agent computing to model the socio-economic world, although the proponents warn that the complexity of socio-economic systems and the slower development of real-time updating means that the societal equivalent of digital twins is 'a long way from being able to simulate real human systems' (Birks et al., 2020, p. 2884).

Agent computing has gained popularity as a tool for transport planning or providing insight for decision-makers in disaster scenarios such as nuclear attacks or pandemics (Waldrop, 2018). UNDP are also trialling the use of an agent computing model to help developing countries work out which policy areas—health, education, transport, and so on—should be prioritised in order to meet the sustainable development goals (Guerrero & Castañeda, 2020). Mainstream economics modelling has struggled to keep pace with the new possibilities brought about by the growing availability of large-scale data, meaning that computational social scientists can and should play a key role in developing collaborations with policymakers and forging a new field of research aimed at enabling governments to design evidence-based policy interventions.

1.7 An Ethics-Driven Computational Social Science

CSS methods are data-driven. Machine learning models in this field are trained on data from human systems. For example, a model to support judicial decision-making will be trained on large datasets generated by earlier judicial decisions. That means that if decision-making in the past or present is biased—clearly the case in some areas, such as policing—then the machine learning algorithms trained on this data will be biased also. The use of the resulting machine learning tools in decision-making processes will reinforce and amplify existing biases. In part for this reason, extensive controversy has accompanied the use of machine learning for decision support, particularly in sensitive areas such as criminal justice (Hartmann & Wenzelburger, 2021; Završnik, 2021) or child welfare (Leslie et al., 2020).

The CSS methods discussed in this chapter raise numerous ethical concerns, from replicating biases to invading people's privacy, limiting individual autonomy, eroding public trust, and introducing unnecessary opaqueness into decision-making

processes—to name only a few. To tackle these issues, CSS should take advantage of the work that has been done on the ethical use of AI technologies in government. Guidance on the responsible design, development, and implementation of AI systems in the public sector (Leslie, 2019) and a framework for explaining decisions made with AI (Information Commissioner’s Office & The Alan Turing Institute, 2020) are used across UK departments and agencies. These publications focus on how the principles of fairness, sustainability, safety, accountability, and transparency can—and should—guide the responsible design, development, and deployment of AI systems. In contrast, Computational Social Science research has focused far more on the technical details of these data-intensive technologies rather than the ethical concerns, which tend to be underplayed. A recent special issue of *Nature* on CSS,² for example, mentioned ethics and responsible innovation only once in the editorial, and none of the articles focused on the topic. So in this case, CSS could have something to learn from recent work on trustworthy and responsible AI innovation for the public sector.

There are significant gains to be had if computational social science makes ethics an integral part of the process of scientific discovery. CSS methods are data-driven, using data generated by existing administrative systems. Rather than replicating biases, CSS can play an important role in shedding light, sometimes for the first time, on the bias endemic in human decision-making. As large-scale data sources become available, CSS could be used to reveal and tackle bias in modern digital public administration and policymaking. Identifying bias and understanding its origins can be a first step towards tackling long-running failings of administration.

1.8 Building Resilience: CSS at the Heart of a Reinvented Policy Toolkit

Nowhere are the possibilities of CSS for public policy—and the importance of realising them—illustrated more starkly than in the coronavirus pandemic of 2020 onwards. Computational Social Science seemed, to these authors at least, to have huge potential for the design of policy interventions and informing decision-making during the pandemic, for example, through undertaking the key tasks of detection, measurement, prediction, etiology, and simulation laid out above. But somehow, the use of CSS in this setting was disappointing. While it was good to see data, modelling, and science in such high relief throughout the pandemic, the use of CSS was limited and many interventions were introduced with no real evidence of their expected payoffs.

The difficulties seemed to be threefold. First, many countries discovered that they did not collect the kind of real-time, fine-grained data that was needed to inform policy design. In the UK, for example, it turned out that there was no availability

² *Nature* volume 595, issue 7866, 2021

of data on the number of people dying of Covid-19 until weeks after the deaths had taken place, making it impossible to calibrate the use of interventions. Economic policymakers had to design financial support mechanisms such as furlough schemes and stimulus packages without fine grained data about the areas of the economy that would be most affected by social distancing measures and supply chain disruptions. This meant that blanket schemes were applied, helping sectors that benefited from the pandemic (such as delivery companies and many technology firms) along with those that had been devastated (such as travel and hospitality). Policymakers and computational social scientists need to work together to identify the data streams that are likely to be needed in a crisis and ‘develop dynamic capabilities’ (Mazzucato & Kattel, 2020).

Second, there seemed to be a universal lack of integrated modelling. The focus tended to be on modelling one policy area at a time. There were models that tracked the spread of the virus and separate models that examined the economic effects. These two issues, however, were inextricably intertwined. The absence of integrated models to capture these interdependencies meant that policymakers often pointed to the trade-off between ‘public health’ and ‘economic recovery’ but were never able to pinpoint optimal interventions. There is a need for CSS to develop more integrated, generalised models that policymakers could turn to in an emergency. Besides their inability to capture interdependencies between policy areas, many economic models proved to be incapable of dealing with surprises. Models of commodity prices, for example, were based on the assumption that negative oil prices were impossible. During the pandemic, it became clear that not enough attention is given to quantifying uncertainty, which can have a cascading effect in complex multi-level systems. To help policymakers equip themselves for future crises, we need to develop CSS models that are based on robust assumptions and are able to quantify uncertainty. Integrated modelling, data-centric policymaking, causal inference, and uncertainty quantification are all ways in which CSS might build resilience into policymaking processes (MacArthur et al., 2022).

Third, it became clear that the organisational structures involved in policymaking to some extent worked against the kind of computational and modelling expertise that was required during the pandemic. Big departments of state have few incentives to share data, and very little tradition of sharing technical solutions to policy problems. This is unfortunate, because the vertical nature of data-intensive methods means that they lend themselves to being transferred across organisational boundaries. Yet policymakers seeking to meet a generic modelling challenge—such as how to identify vulnerable groups, quantify uncertainty or use machine learning to derive causal explanations as laid out above—are much more likely to seek help in their own department than to turn to departments or agencies in other parts of government. This siloed approach works against building up of expertise.

Overcoming these issues could allow CSS to usher in a new era of policymaking. As we begin to emerge from the pandemic, the word ‘resilience’ has become widespread in policy circles. Resilience is an organisational value that underpins how a government designs its policymaking systems and processes (Hood, 1991). Governments that value resilience prioritise stability, robustness, and adaptability.

Developing the CSS tools and models we have discussed here, with the focus on detecting and measuring trends and patterns, predicting and understanding human behaviour, and developing integrative modelling techniques that can simulate policy interventions all point in this direction. A resilient approach of this kind could equip policymakers to tackle the aftermath of the pandemic and face future crises (MacArthur et al., 2022).

1.9 Conclusion

This chapter has shown some of the transformational potential of Computational Social Science, bringing analysis of large-scale social and economic data into policymaking. CSS can renew the toolbox of contemporary government, refreshing and sharpening the essential tasks of detection, measurement, prediction, simulation, and etiology. None of these tasks can, alone, transform the policy toolkit. They need to be used in concert and require large-scale, real-time, fine-grained data sources. Measurement, for example, requires detection to be able to observe trends in the variable under scrutiny. Both are needed for prediction, which on its own is of questionable value in policy settings that lack the ability to pinpoint causality. Many researchers are making the case for integrative modelling that incorporates prediction and causal inference. Simulation requires large-scale data and is often used in conjunction with more predictive techniques.

New possibilities for the use of large-scale data about human behaviour bring new responsibilities, in terms of implementing and developing guidelines and frameworks for responsible innovation. Substantial progress has already been made in building ethical frameworks for the growing use of artificial intelligence in government. Guided by these frameworks, CSS researchers have a real opportunity to make explicit long-running biases and entrenched inequalities in public policy and administration. Their scholarship and methodologies have the potential to usher in a new era of policymaking, where interventions and administrative systems are more fair than ever before, as well as more efficient, effective, responsive, and prescient (Margetts & Dorobantu, 2019).

The need to respond to the coronavirus pandemic has raised the profile of data and modelling but has also illustrated missed opportunities in terms of data flows, integrative modelling, and the development of expertise. To face future crises, we need to overcome these challenges, bringing CSS methods to the heart of policymaking and developing models to inform the design of resilient policy interventions.

References

- Athey, S. (2017). Beyond prediction: Using big data for policy problems. *Science*, 355(6324), 483–485. <https://doi.org/10.1126/science.aal4321>
- Athey, S., & Imbens, G. W. (2017). The state of applied econometrics: Causality and policy evaluation. *Journal of Economic Perspectives*, 31(2), 3–32. <https://doi.org/10.1257/jep.31.2.3>
- Axtell, R. L. (2000). *Why agents? On the varied motivations for agent computing in the social sciences*. <http://www2.econ.iastate.edu/tesfatsi/WhyAgents.RAxtell2000.pdf>.
- Axtell, R. (2018). Endogenous firm dynamics and labor flows via heterogeneous agents. In C. H. Hommes & B. D. LeBaron (Eds.), *Handbook of computational economics* (Vol. 4, pp. 157–213). Elsevier. <https://doi.org/10.1016/bs.hescom.2018.05.001>
- Bastow, S., Dunleavy, P., & Tinkler, J. (2014). *The impact of the social sciences: How academics and their research make a difference*. SAGE Publications. <https://doi.org/10.4135/9781473921511>
- Birks, D., Heppenstall, A., & Malleson, N. (2020). Towards the development of societal twins. In: *Frontiers in artificial intelligence and applications. 24th European Conference on Artificial Intelligence (ECAI 2020)* (pp. 2883–2884).
- Bright, J., Ganesh, B., Vogl, T., & Cathrine, S. (2019). *Data science for local government*. Oxford Internet Institute. https://pure.rug.nl/ws/portalfiles/portal/117364515/SSRN_id3370217.pdf
- Cheung, H. K. (1972). A day-use park visitation model. *Journal of Leisure Research*, 4(2), 139–156. <https://doi.org/10.1080/00222216.1972.11970070>
- Donahue, M. L., Keeler, B. L., Wood, S. A., Fisher, D. M., Hamstead, Z. A., & McPhearson, T. (2018). Using social media to understand drivers of urban park visitation in the Twin Cities, MN. *Landscape and Urban Planning*, 175, 1–10. <https://doi.org/10.1016/j.landurbplan.2018.02.006>
- Dunleavy, P. (2016). ‘Big data’ and policy learning. In G. Stoker & M. Evans (Eds.), *Evidence-based policy making in the social sciences: Methods that matter* (Vol. 143). Policy Press.
- Dunning, T. (2012). *Natural experiments in the social sciences: A design-based approach*. Cambridge University Press. <https://doi.org/10.1017/CBO9781139084444>
- Engstrom, D. F., Ho, D. E., Sharkey, C. M., & Cuéllar, M.-F. (2020). Government by algorithm: Artificial intelligence in Federal Administrative Agencies. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3551505>
- Guerrero, O. A., & Castañeda, G. (2020). Policy priority inference: A computational framework to analyze the allocation of resources for the sustainable development goals. *Data & Policy*, 2, e17. <https://doi.org/10.1017/dap.2020.18>
- Guest, E., Vidgen, B., Mittos, A., Sastry, N., Tyson, G., & Margetts, H. (2021). An expert annotated dataset for the detection of online misogyny. In: *Proceedings of the 16th conference of the European chapter of the Association for Computational Linguistics: Main Volume* (pp. 1336–1350). <https://doi.org/10.18653/v1/2021.eacl-main.114>
- Hallsworth, M., List, J. A., Metcalfe, R. D., & Vlaev, I. (2017). The behavioralist as tax collector: Using natural field experiments to enhance tax compliance. *Journal of Public Economics*, 148, 14–31. <https://doi.org/10.1016/j.jpubeco.2017.02.003>
- Hamstead, Z. A., Fisher, D., Ilieva, R. T., Wood, S. A., McPhearson, T., & Kremer, P. (2018). Geolocated social media as a rapid indicator of park visitation and equitable park access. *Computers, Environment and Urban Systems*, 72, 38–50. <https://doi.org/10.1016/j.compenvurbysys.2018.01.007>
- Hartmann, K., & Wenzelburger, G. (2021). Uncertainty, risk and the use of algorithms in policy decisions: A case study on criminal justice in the USA. *Policy Sciences*, 54(2), 269–287. <https://doi.org/10.1007/s11077-020-09414-y>
- Hofman, J. M., Watts, D. J., Athey, S., Garip, F., Griffiths, T. L., Kleinberg, J., Margetts, H., Mullainathan, S., Salganik, M. J., Vazire, S., Vespignani, A., & Yarkoni, T. (2021). Integrating explanation and prediction in computational social science. *Nature*, 595(7866), 181–188. <https://doi.org/10.1038/s41586-021-03659-0>

- Hood, C. (1991). A public management for all seasons? *Public Administration*, 69(1), 3–19. <https://doi.org/10.1111/j.1467-9299.1991.tb00779.x>
- Hood, C. C., & Margetts, H. Z. (2008). The tools of government in the digital age. *Public Administration*, 86(4), 1137–1138. https://doi.org/10.1111/j.1467-9299.2008.00756_4.x
- Information Commissioner's Office, & The Alan Turing Institute. (2020). *Explaining decisions made with AI*. <https://ico.org.uk/for-organisations/guide-to-data-protection/key-dp-themes/explaining-decisions-made-with-artificial-intelligence/>
- Kleinberg, J., Ludwig, J., Mullainathan, S., & Obermeyer, Z. (2015). Prediction policy problems. *American Economic Review*, 105(5), 491–495. <https://doi.org/10.1257/aer.p20151023>
- Larcom, S., Rauch, F., & Willems, T. (2017). The benefits of forced experimentation: Striking evidence from the London underground network. *The Quarterly Journal of Economics*, 132(4), 2019–2055. <https://doi.org/10.1093/qje/qjx020>
- Lazer, D., Hargittai, E., Freelon, D., Gonzalez-Bailon, S., Munger, K., Ognyanova, K., & Radford, J. (2021). Meaningful measures of human society in the twenty-first century. *Nature*, 595(7866), 189–196. <https://doi.org/10.1038/s41586-021-03660-7>
- Lazer, D., Pentland, A., Adamic, L., Aral, S., Barabasi, A. L., Brewer, D., Christakis, N., Contractor, N., Fowler, J., Gutmann, M., Jebara, T., King, G., Macy, M., Roy, D., & Van Alstyne, M. (2009). Life in the network: The coming age of computational social science. *Science (New York, N.Y.)*, 323(5915), 721–723. <https://doi.org/10.1126/science.1167742>
- Leslie, D. (2019). *Understanding artificial intelligence ethics and safety: A guide for the responsible design and implementation of AI systems in the public sector*. The Alan Turing Institute. <https://doi.org/10.5281/zenodo.3240529>
- Leslie, D., Holmes, L., Hitrova, C., & Ott, E. (2020). *Ethics review of machine learning in children's social care*. What Works for Children's Social Care. https://whatworks-csc.org.uk/wp%20content/uploads/WWCSC_Ethics_of_Machine_Learning_in_CSC_Jan2020_Accessible.pdf
- MacArthur, B. D., Dorobantu, C. L., & Margetts, H. Z. (2022). Resilient government requires data science reform. *Nature Human Behaviour*, 6(8), 1035–1037. <https://doi.org/10.1038/s41562-022-01423-6>
- Margetts, H. (1999). *Information technology in government: Britain and America*. Routledge.
- Margetts, H., & Dorobantu, C. (2019). Rethink government with AI. *Nature*, 568(7751), 163–165. <https://doi.org/10.1038/d41586-019-01099-5>
- Mazzucato, M., & Kattel, R. (2020). COVID-19 and public-sector capacity. *Oxford Review of Economic Policy*, 36(Supplement_1), S256–S269. <https://doi.org/10.1093/oxrep/graa031>
- Muellerleile, C., & Robertson, S. L. (2018). Digital Weberianism: Bureaucracy, information, and the techno-rationality of neoliberal capitalism. *Indiana Journal of Global Legal Studies*, 25(1), 187. <https://doi.org/10.2979/indjglolegstu.25.1.0187>
- Ostmann, F., & Dorobantu, C. (2021). *AI in financial services*. The Alan Turing Institute. <https://doi.org/10.5281/zenodo.4916041>
- Röttger, P., Vidgen, B., Nguyen, D., Waseem, Z., Margetts, H., & Pierrehumbert, J. (2021). HateCheck: Functional tests for hate speech detection models. In: *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 41–58). <https://doi.org/10.5281/zenodo.4916041>
- Schelling, T. C. (1971). Dynamic models of segregation. *The Journal of Mathematical Sociology*, 1(2), 143–186. <https://doi.org/10.1080/0022250X.1971.9989794>
- Sinclair, M., Zhao, Q., Bailey, N., Maadi, S., & Hong, J. (2021). Understanding the use of greenspace before and during the COVID-19 pandemic by using mobile phone app data. In: *GIScience 2021 Short Paper Proceedings. 11th International Conference on Geographic Information Science. September 27–30, 2021. Poznań, Poland (Online)*. <https://doi.org/10.25436/E2D59P>

- Suse, S., Mashhadi, A., & Wood, S. A. (2021). Effects of the COVID-19 Pandemic on Park Visitation Measured by Social Media. In: *Companion Publication of the 2021 Conference on Computer Supported Cooperative Work and Social Computing* (pp. 179–182). <https://doi.org/10.1145/3462204.3481754>
- Taylor, H., Vidgen, B., Anastasiou, Z., Pantazi, M., Inkster, B., & Margetts, H. (2021, June). Investigating vulnerability to online health-related misinformation during COVID-19. In: *MISDOOM 2021: 3rd Multidisciplinary International Symposium on Disinformation in Open Online Media*. https://www.researchgate.net/publication/355436534_Investigating_vulnerability_to_online_health-related_misinformation_during_COVID-19
- Varian, H. R. (2014). Big data: New tricks for econometrics. *Journal of Economic Perspectives*, 28(2), 3–28. <https://doi.org/10.1257/jep.28.2.3>
- Vidgen, B., Botelho, A., Broniatowski, D., Guest, E., Hall, M., Margetts, H., Tromble, R., Waseem, Z., & Hale, S. (2020). Detecting East Asian Prejudice on social media. *ArXiv:2005.03909 [Cs]*. <http://arxiv.org/abs/2005.03909>
- Vogl, T. M., Seidelin, C., Ganesh, B., & Bright, J. (2020). Smart technology and the emergence of algorithmic bureaucracy: Artificial intelligence in UK local authorities. *Public Administration Review*, 80(6), 946–961. <https://doi.org/10.1111/puar.13286>
- Wagner, C., Strohmaier, M., Olteanu, A., Kıcıman, E., Contractor, N., & Eliassi-Rad, T. (2021). Measuring algorithmically infused societies. *Nature*, 595(7866), 197–204. <https://doi.org/10.1038/s41586-021-03666-1>
- Waldrop, M. M. (2018). Free agents. *Science*, 360(6385), 144–147. <https://doi.org/10.1126/science.360.6385.144>
- Završnik, A. (2021). Algorithmic justice: Algorithms and big data in criminal justice settings. *European Journal of Criminology*, 18(5), 623–642. <https://doi.org/10.1177/1477370819876762>

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

