Surveillance & Society

Article

Automated Government Benefits and Welfare Surveillance

Mike Zajko

University of British Columbia, Okanagan, Canada
mike.zajko@ubc.ca

Abstract

This article examines the "digital welfare state" historically, presently, and into the future, with a focus on what artificial intelligence means for welfare surveillance. Drawing on scholarship about the development of bureaucracy, the welfare state, and automation, as well as specific examples from the Netherlands, I argue that problems posed by artificial intelligence in public administration are often misplaced or misattributed and that the societal challenges we can expect to encounter in welfare surveillance are more likely to be historically familiar than technologically novel. New technologies do provide some new capabilities, which explains the uptake of algorithmic tools in welfare fraud investigation and the use of chatbots in assisting with welfare applications. Algorithmic systems are also increasingly subject to "audits" and regulations that mandate accountability. However, many of the key issues in the automation of the welfare state are the same as identified in scholarship that long precedes the current hype around artificial intelligence. These issues include a persistent suspicion of welfare recipients to justify surveillance as a form of fraud identification, opaque decision-making, and punitive measures directed against marginalized groups, enacting harm and reproducing inequalities.

Introduction

The "digital welfare state," according to a notable report by UN Special Rapporteur on extreme poverty and human rights, is defined as one in which "systems of social protection and assistance are increasingly driven by digital data and technologies that are used to automate, predict, identify, surveil, detect, target and punish" (Alston 2019: 4). Among these technologies, "there is little doubt that the future of welfare will be integrally linked to digitization and the application of artificial intelligence," which threatens to push us "stumbling, zombie-like, into a digital welfare dystopia" (Alston 2019: 21). More than four years later, there has been increased public sector adoption of AI systems (see Maragno et al. 2021; Neumann, Guirguis, and Steiner 2022), but their uptake for welfare administration remains limited. AI continues to generate attention as the high-tech frontier of governance, but this distracts from the socio-technical configurations in which such developments become embedded. By way of discursive shift, any kind of automation in decision making is increasingly referred to as "AI," and administrative procedures are now being redefined as "algorithms." This being said, there are some distinct features of new algorithmic technologies that are worth considering and that explain their appeal and adoption in certain domains of government.

Contemporary forms of AI are most commonly defined by their ability to "learn" or be "trained" from vast datasets—what earlier scholarship dubbed "Big Data" and predictive analytics (see Andrejevic and Gates 2014). Through "machine learning" (ML), an algorithm changes as it receives feedback about its performance and can be updated through the inclusion of new data. Surveillance is therefore initially implicated in the collection of "training data" for ML and in the assemblage of data sources to categorize people and their characteristics. The use of large datasets with many variables "to generate new information

Zajko, Mike. 2023. Automated Government Benefits and Welfare Surveillance. *Surveillance & Society* 21 (3): 246-258.
https://ojs.library.queensu.ca/index.php/surveillance-and-society/index | ISSN: 1477-7487

about the state of the world and otherwise unobservable correlations compels policymakers and administrators to increase existing and deploy new data generating systems" (Young, Bullock, and Lecy 2019: 310). While simpler, rule-based algorithms are well suited for automating many kinds of eligibility screening, ML-based algorithms excel in finding patterns in data. Therefore, welfare administration tasks for which AI has been implemented include flagging hidden indicators of fraud (Savage 2021), as well as risk indicators in child welfare (Eaton 2019). Governments have also been adopting ML-based chatbots to recommend services or benefits and assist with applications for services (Makasi et al. 2022), but a predominant use has been welfare surveillance.

In this article, I will focus on the use of AI by state agencies to distribute benefits and welfare and to govern problems associated with poverty. Since 2018, I have been studying the automation of federal government services in Canada—a country that has been considered a "world leader" in AI development (Lepage-Richer and McKelvey 2022), although much of the use of AI in Canada's government has remained limited to internal processes (Morrison 2023).[1] While I draw some examples from Canada and summarize broad international trends towards digital government, most of my examples of recent welfare surveillance systems are drawn from the Netherlands, which has emerged as one of the best documented examples of a "digital welfare state." This is largely as a consequence of reporting following a major recent scandal (SyRI) that provided the "clearest warning against technological overreach" in welfare surveillance (Geiger 2023).

Given that surveillance studies is largely interested in problems of governance and information and that governments have become very interested in artificial intelligence (AI) technologies, it is important to relate AI to earlier scholarship in surveillance studies, in addition to other work on state governance. In doing so, we need to counter the reification of AI as a new and distinct phenomenon by locating it within established sets of problems and solutions, dating back to early projects to manage the poor through statistical techniques and discussions of automated welfare surveillance from previous decades. My argument is that we should be more concerned about repeating well-documented issues in the administration of welfare as new information-processing technologies enter the picture, rather than anything that is fundamentally new about AI.

To limit the scope of this article, I will avoid discussing the use of AI in policing, even though police forces have been major adopters of AI technologies, and police are heavily implicated in the governance and reproduction of inequality. Others have analyzed the use of predictive policing to target poor and racialized populations, particularly through "feedback loops" where data collected about individuals and groups leads to further surveillance (Benbouzid 2019; Brayne 2020). While the criminalization of poverty is a historically recurring phenomenon, and police remain deeply involved in various social services, my focus is on the use of technologies by administrative agencies tasked with making decisions about social welfare. Police agencies become relevant primarily to the extent that they are involved in welfare investigations and when they contribute or draw from the assemblage of data used in welfare surveillance (i.e., Davidson and Adriaens 2022).

Artificial Intelligence as a Political Technology

The current period of hype and attention to AI technologies follows in the wake of previous cycles of AI "boom-and-bust" (Strickland 2021) or "springs and winters" (Mitchell 2021) that stretch back to the 1950s. Today's AI applications are built on the basis of fundamental advances in ML techniques, computational power, and large datasets over the past two decades, but our expectations of their power and abilities are often inflated and dangerously anthropomorphized (Weil 2023). While claims and exaggerations about AI have often been made through comparisons between AI and human minds, it is more helpful to ask how AI relates to human organizations. Instead of defining AI through its resemblance to human action or cognition,

[1] One notable public-facing use of AI in Canada's government has been for the processing of immigration applications (including fraud investigation; see Reevely 2021, 2023).

my focus is on how these technologies are developed to make decisions about people in the present and future on the basis of vast datasets collected, extracted, and translated from the past.

This article considers AI as a technology of governance, defined primarily by its use of "machine learning" (ML) algorithms based on statistical reasoning that are then used to make social distinctions and predictions. Understanding key features of machine learning is often less important than an algorithm's "objective function" (Peeters 2020: 512)—what it is "optimized" for or "trained" to maximize/minimize in its decisions (Croll 2018). In this regard, algorithmic systems resemble narrowly goal-focused actors, such as corporations and their goal of maximizing profit or government bureaucracies with their clearly defined sets of inputs and outputs (Penn 2018). Governments have long been automating mundane bureaucratic tasks, and ML techniques provide just another means of constructing administrative algorithms.

AI is a "political technology" in that it has been historically supported and promoted by governments, is defined and understood in relation to the problems of governance, and is "centred on the meta-conditions of social, economic, and political life" (Lepage-Richer and McKelvey 2022: 3). AI technologies are intimately bound in meaning and intent with power and control as enabled through information processing. The use of AI for governance has long been imagined (Natale and Ballatore 2020), but algorithmic technologies labelled as AI have in recent years been adopted by government agencies for a range of information processing and classification tasks (Maragno et al. 2021). In many cases, however, the label of "AI" is applied to relatively simple forms of automation and algorithmic classification rather than the specific (ML) techniques that some associate with "true AI" (Henman 2020: 210).[2] A more general and inclusive term is "automated decision-making" or "ADM systems" (Cobbe 2019)—sometimes used interchangeably with "AI" by scholars of government (Kuziemski and Misuraca 2020) but connoting a deeper history of automation.

As with any "new" technology implicated in surveillance, we need to approach claims of novelty with skepticism, examining their technological antecedents and longstanding governmental goals (Lyon 2014). While it is true that government agencies now have access to vastly more data than they ever have, and are employing statistically-based ML techniques for decision-making using this information, the greatest concerns around such efforts echo those made about earlier periods of automation in government. It is therefore far more helpful to understand these developments through scholarship on administrative state governance, rather than anything that might be new about contemporary AI technologies.

Several authors have traced continuities between supposedly new technologies and established forms of government. Cellard (2022) shows how existing bureaucratic and administrative procedures are now being "rebranded" as "algorithms," while Lepage-Richer and McKelvey (2022) compare attributions of intelligence to machines and state agencies. They argue that national governments might "appear so receptive to the latest innovations in information processing… because they themselves have been previously re-configured as information processing systems," requiring us to consider "government itself as a form of artificial intelligence" (Lepage-Richer and McKelvey 2022: 2). Alkhatib (2021) draws on Scott's (1998) *Seeing Like a State* to "[frame] the massive algorithmic systems that harm marginalized groups as functionally similar to massive, sprawling administrative states" (Alkhatib 2021: 2), and Hoffman (2021) builds on the concept of "administrative violence" carried out by state agencies to develop the notion of "data violence" carried out through information systems. These conceptual analogies between algorithms and administrative states are effective, given all that the two have in common. Administrative agencies are

[2] Others associate "true AI" with the creation of an artificial "general intelligence," which like a human being, would be able to perform a wide range of tasks and reason about the world (Mitchell 2021). However, existing systems remain far from this ideal and remain tied to statistical techniques. Discussions of ML-based systems make generous use of human analogies and metaphors to mischaracterize how these systems operate, further contributing to "AI hype" (Weil 2023).

increasingly adopting algorithmic systems, and as one consequence, we are appreciating just how "algorithmic" the work of government agencies has been for a long time before the digital present.

ADM systems are integrated into public sector organizations that often already function as impersonal decision-making machines, but they do alter the work done by public sector employees, and change the way information flows in government decision-making. A common way of theorizing these developments has been to build on Bovens and Zouridis's (2002) distinction between "street-level bureaucracy" (as conceptualized by Lipsky 2010) and "screen-level bureaucracy," "system-level bureaucracy" (Busch and Henriksen 2018; Elyounes 2021), or the emergence of "street-level algorithms" (Alkhatib 2021). Often, such scholarship examines what these transformations mean for the discretion exercised by those public sector employees (including social workers and police officers) who are empowered to make governmental decisions about people's lives (Bullock, Young, and Wang 2020).

A particularly important kind of decision made by contemporary governments involves allocating a variety of benefits to citizens or otherwise qualified recipients. This requires some process of identification, and often includes the collection of other personal information in order to decide eligibility, making welfare programs a recurring concern of surveillance studies. Welfare can also be used to exemplify the "ambiguous" nature of surveillance (Lyon 2007) in that government recognition and government benefits provide necessary assistance to those in need and may be used to promote equity, while the social sorting inherent in these processes also function as a mechanism of exclusion, targeting, and oppression. Critical scholarship has repeatedly argued that welfare surveillance is a political tool to manage marginalized populations and reproduce their "conditions of abjection" (Monahan 2017), and while it is true that surveillance-driven social assistance can be used to counteract inequalities, "targeted services always risk reinforcing the very divisions they seek to undo" (Henman and Marston 2008: 202). When marginalized groups are targeted for surveillance, the information that is obtained can create a feedback loop, justifying further surveillance and social control (Eubanks 2018). When algorithms are involved, disproportionate outcomes are reproduced under "a veneer of objectivity" (Keddell 2019: 9).

Although welfare surveillance programs do end up targeting individuals and families, they are based on government concerns with larger, collective phenomena, including class, gender, and race or colonial relations. Population-level surveillance began with the creation of the "population" as an object to be measured through official statistics, as part of new "biopolitical" strategies for managing the population's problems and maximizing its potential (Ruppert 2012). Poverty and early forms of social welfare were among the problems of the state and its population that statistics were created to address. These included debates over the causes of "pauperism" and the administration of the "Poor Law" in England during the late-nineteenth and early-twentieth centuries, and the use of techniques such as correlation and regression by the eugenicist pioneers of statistics to model the biological inferiority of the lower classes (Desrosières 1998). Statistical techniques would eventually become a foundation for today's AI systems, many of which classify or predict optimal outcomes based on patterns found in historical datasets that they are "trained" on. In addition to these statistical foundations, more specific ways of measuring and classifying human characteristics in AI systems, such as gender and race classification or facial identification, are based upon old methods for measuring and categorizing human populations (Scheuerman, Pape, and Hanna 2021; Stark and Hutson 2022; Taylor, Gulson, and McDuie-Ra 2021).

Developing the Welfare Surveillance State

The spread of "welfare state" policies around the world (beginning with Germany in the late eighteenth century) and the characteristics these assumed in different countries has been theorized through literature on "policy diffusion" and "policy translation" (see Béland et al. 2022). Given the historical transformations involved, it is difficult to identify the core ideas, rationalizations, and political goals of *the* "welfare state," but in general the development of post-World War II welfare programs should be interpreted not as "as a response to the demand for socioeconomic equality, but to the demand for socioeconomic security" (Flora and Heideheimer 1981: 23). In Canada, for instance, the mid-twentieth century creation of "social security" was tied to “a broad strategy to stabilize society, contain socialism and communism, and ease class conflict”—one that avoided characterizing inequality as the problem or redistribution as its solution but was “strictly about providing a basic minimum of material security” (Sager 2021: 268). In the Netherlands, social security was the outcome of negotiations between employers, unions, and the state, wherein business leaders sought to maintain “harmonious” relations with labor (Oude Nijhuis 2018; Touwen 2023).

The twentieth-century welfare state was made possible by the growth of government bureaucracy, record keeping, and identification systems, all of which enabled closer surveillance of populations as well as individuals. “Public administration” normalized these surveillance practices by requiring citizens to provide personal information in order to receive services, by maintaining records about these citizens, and increasingly by connecting together records from different sources (Webster 2012). Greater use of computer systems in the late-twentieth century “increased the relevance of personal data for the modern welfare state” (Gantchev 2019: 6) and heightened the asymmetry in power between individuals and bureaucratic institutions, gradually making surveillance “more extensive, more efficient, and less obtrusive than former methods” (Gandy 1989: 62).

While social welfare can be broadly defined to include benefits and services provided by state and non-state entities, including education, health insurance, and housing (Henman and Marston 2008), my focus here is on the targeted provision of welfare by state agencies to those most “in need,” which can often translate to identifying and distinguishing the “deserving” and “undeserving poor” (Romano 2018). Welfare regimes that include a “means test” to determine eligibility are particularly prone to expansive surveillance, as these justify investigations and accounting of resources available to welfare applicants or determinations about their ability to work (Gilliom 2001). Such “welfare conditionality” has arguably become more common in the late-twentieth and early-twenty-first centuries (Gantchev 2019), requiring applicants to submit evidence to establish their worthiness, to undergo “forensic” examinations of their financial and medical circumstances, and for successful applicants to be subject to continuous monitoring to maintain their benefits (Whelan 2021). This coincides with growing concerns over efficiency and welfare “fraud” that accompanied what is widely characterized a neoliberal shift in governance at the same time as countries were increasingly turning to computerized databases and algorithms to administer benefits (see Maki 2021).

In the Netherlands for example, the 1980s were the beginning of a period of welfare retrenchment and concerns about worker “inactivity,” which led to a greater emphasis on fraud detection (Oude Nijhuis 2018). By the mid-2000s, an early ADM system was being used to hunt for welfare fraud by comparing recorded water usage at a residence with the number of people claiming to be living there (Gantchev 2019: 16). In more recent years, much of the responsibility for fostering a “participation society” in the Netherlands (toward greater participation in the workforce) has been decentralized or shifted to local governments (van Kersbergen and Metliaas 2020). This local government responsibility includes fighting welfare fraud, with political pressure to tackle fraud increasing as a result of several incidents “involving individuals with an immigration background” (Ranchordás and Schuurmans 2020: 11–12). By 2019, fraud investigations were found to be a major application of “public sector data analytics” in the Netherlands (van Veenstra, Grommé, and Djafari 2020).

ADM systems, including ones based on relatively simple algorithms and more sophisticated machine learning techniques, have been incorporated into a well-defined “problem space” for the public administration of welfare—how to get benefits into the hands of the “right” people and not the “wrong” kind. What exactly defines the category of the “undeserving” (Romano 2018) has been historically contingent, with many recent surveillance regimes being devoted to identifying and excluding those who are guilty of welfare fraud. Such “fraud” can include not providing up-to-date circumstances of one’s situation to authorities or the myriad forms of “everyday resistance” that welfare recipients use as “survival strategies” (Gilliom 2001). The social construction of welfare fraud has been broadly theorized as part of the criminalization of poverty or the criminalization of welfare (Dobson 2019; Gustafson 2011).

What is algorithmically classified as "fraud" can be modeled from a growing number of variables, but computer-mediated welfare surveillance inherits earlier attitudes and assumptions about the moral regulation of the poor (Maki 2021). These include thousands of years of concern over incentives to be "idle" from work, the prevalence of which has more to do with periods of "moral crisis" rather than neoliberal preferences for a shrinking state (Romano 2018). The wider category of "the poor" has long been subjected to high levels of surveillance, not only due to generalized concerns over danger and moral deviance but also as legitimated by the understanding that a person receiving assistance from the state should open themselves up to ongoing state scrutiny. In the words of Coser (1965: 145), "the protective veil [of privacy] available to other members of society is explicitly denied to them." Welfare regimes vary in the extent to which this scrutiny involves assessing "initial eligibility" up-front or "continuing eligibility" (Fellowes and Rowe 2004: 365) and auditing after-the-fact. However, contemporary welfare regimes fundamentally require a way to identify applicants and sort large volumes of files, and several examples from the Netherlands illustrate the dynamics involved as algorithmic systems are combined with a growing abundance of data sources.

Automated Welfare Surveillance in the Netherlands

According to an article in *IEEE Spectrum*, the first instance in which "AI had a hand in forcing a government to resign" occurred in January 2021, after it had become clear that thousands of people in the Netherlands had been wrongly suspected and punished for child benefit fraud (Rao 2022). The Dutch SyRI scandal was presented as a "warning" about the use of AI in government, and while the algorithm that created the problematic risk profiles was described as "self-learning" (Heikkilä 2022), the actual details of the ADM system used were so opaque (lack of transparency being part of the problem) that it is impossible to know whether techniques such as machine learning were actually involved (van Bekkum and Borgesius 2021). The claim that a government institution was "felled by AI" (Rao 2022) is an example of AI hype that attributes far too much agency to these technologies, but it may also reflect a European regulatory definition of AI that can broadly encompass a variety of automated systems. Indeed, the scandal appears to have much in common with other notable failures of ADMs since the early 2000s, wherein pressures to cut costs and identify cases of fraud caused serious harm to large numbers of people who happened to be flagged by an algorithm (Eubanks 2018; Mann 2020; Peachey 2022).

First, regardless of the nature of the algorithm involved, it is indeed correct to see the Dutch SyRI scandal over child care benefits as a "warning" (Heikkilä 2022), because the issues underpinning it will continue to be relevant even as technologies change. Concerns over human involvement and discretion, as well as algorithmic transparency, can be identified across the wide range of ADM systems. Rather than heading toward a future where autonomous machines replace human bureaucrats, we live in a present where bureaucratic decisions are made in socio-technical systems that "imbricate" (Leonardi 2011) human and non-human agency. In the Netherlands (Burgess, Schot, and Geiger 2023; Elyounes 2021: 504) and elsewhere (Burke and Ho 2022; Geiger 2023; Reevely 2021), there has been a tendency for governments to highlight how a human being acts as the final decision-maker in order to counter concerns about automation. This can be used to deflect attention from the ways that technologies shape or direct these human decisions. In many examples of algorithmic governance, human discretion has shifted to the design and configuration of decision-making systems or to decide on cases that have been algorithmically flagged for scrutiny (Elyounes 2021), and such hybrid forms of decision-making will continue to characterize the automation of public services in the future.

One place that human agency remains central for ADM systems is in the decisions around their procurement and deployment. In the Dutch SyRI scandal, a key way in which inequality was reproduced was through the human decision to target the algorithm at low-income or "problem" neighborhoods (van Bekkum and Borgesius 2021; Wieringa 2023), similar to how predictive policing systems are more often directed at disadvantaged populations (Brayne 2020). Dutch welfare surveillance continues to target particular neighborhoods (Davidson and Adriaens 2022), and the choice to deploy these systems against populations that are predominantly poor, racialized, or of immigrant background is just as important for the reproduction of inequality as how these algorithms function.

Inequality is reproduced through a variety of channels that complement one another in reinforcing existing hierarchies and targeting social control against marginalized populations, but ML algorithms can statistically reproduce social inequalities in their training data, even along lines (i.e., race or ethnicity) that are not explicitly being measured. While governments are typically prohibited from discriminating on certain grounds, these characteristics become encoded indirectly or through proxies in data, the number of which increases with the size of the datasets used. In the Netherlands, the government is prohibited from discriminating on the basis of ethnicity, and as a consequence, datasets used in welfare surveillance do not list ethnicity as a variable. However, there are numerous other correlates with ethnicity (such as language fluency), and these can result in the disproportionate targeting of ethnic groups for fraud investigation (Constantaras et al. 2023; see also Belleman, Heilbron, and Kootstra 2023).

More recent examples of welfare surveillance in the Netherlands highlight additional problems of opacity and explainability in complex ADM systems. In 2020, officials in the Dutch region of Walcheren were concerned that their own fraud detection algorithm raised similar issues to SyRI, but they struggled to understand how the algorithm worked because of a lack of transparency from the private-sector developers of the algorithm as well as their own lack of technical expertise. The system "analyzed details of local people claiming welfare benefits and then sent human investigators a list of those it classified as most likely to be fraudsters," but a subsequent audit found that "The risks indicated by the AI algorithm are largely randomly determined" (Meaker 2023). While in the SyRI scandal, decisions about benefits fraud could draw on seventeen broad categories of data from various sources (van Bekkum and Borgesius 2021), a complex ML-based fraud detection algorithm deployed in Rotterdam between 2017 and 2021 used 315 variables to calculate a person's risk score. Even so, Rotterdam's algorithm was described as "alarmingly inaccurate" (Constantaras et al. 2023), and it is worth remembering that the ability of ML-based algorithms to predict individual actions or experience has often been hyped or exaggerated and that models trained on thousands of variables may perform no better than simple algorithms based on a few (Salganik et al. 2020).

The power of these systems lies in their ability to statistically model complex patterns rather than applying clear criteria for decisions, which raises particular problems for government accountability and administrative justice (Henman 2020) and perhaps even undermines core justifications for the existence of administrative agencies (Calo and Citron 2021). While it is true that ML-based algorithms pose additional challenges to transparency and the explainability of government decisions due to their complexity, ambiguity, or incommensurability with human reasoning (de Bruijn, Warnier, and Janssen 2022; Fazi 2020), there are often more fundamental problems in how these systems are used, including forms of opacity, secrecy, or barriers to government accountability that can make it difficult to determine how decisions are being made (Burrell 2016). Although it was complex, Rotterdam's algorithm could still be subjected to an external audit (resulting in its suspension in 2021), and it was disclosed to journalists "when faced with the prospect of potential court action under freedom-of-information laws," allowing for some public accountability for how risk scores were calculated (Constantaras et al. 2023). Details of SyRI also had to be fought for through freedom-of-information, with the Dutch government keeping its risk calculation methods secret to avoid giving "criminals an advantage" (Gantchev 2019: 18). Other governments have similarly opposed disclosing details about their "cutting-edge" algorithmic technologies because they claim this would make it easier to commit fraud (Savage 2021; see also Solano 2022). In other cases, the secrecy results from the private companies that develop these systems protecting what they claim as intellectual property (Meaker 2023; Ranchordás and Schuurmans 2020: 28–29).

Concerns about enabling fraud or violating intellectual property should not prevent the decisions made by ADMs from being systematically monitored. Indeed, just as ML-based techniques have become widely used in auditing (Fedyk et al. 2022) to identify "fraud," "anomalies," and "improper payments" (Bullock, Young, and Wang 2020), algorithmic systems are themselves increasingly being audited by internal and external actors. There are different approaches to what such an audit should encompass (Vecchione, Levy, and

Barocas 2021)—whether to include training data, the resulting algorithms, or their consequences. Some aspects of how algorithmic systems operate may be inscrutable to auditors, particularly if these are external actors who may only be able to track the "impacts" of the system through some of its decisions (Raji et al. 2020). But internally, algorithmic systems can be characterized by a high level of traceability that can enable "self-auditing," even if such audits themselves rely on the use of algorithmic techniques, so that "evaluation and auditing become inter-algorithmic in form" (Power 2022: 6).

The Future of Automated Welfare Surveillance

There is nothing inevitable about the use of new algorithmic technologies for the automation of welfare, and many experimental and pilot projects have not proceeded beyond initial stages, with some rolled back following scrutiny (i.e., Ho and Burke 2022). However, it is likely that, in addition to the continued use of algorithms to identify fraud or other forms of deviance, ADMs will eventually become more integrated into the delivery of services and the allocation of benefits. As one example, Statistics Canada imagines a future where the agency collects "microdata" from many new sources and uses AI or predictive analytics in order to help guide the decisions of both individual "clients" and government providers of social services (Statistics Canada 2022).[3] While one goal is to develop predictive models to identify who will require what kinds of services in the future or to target services based on risk, few governments have implemented such systems, and in the short term, we are more likely to see continued automation in screening for eligibility and fraud.

A predictable consequence of this kind of automation is the production of harm "at scale," whether through algorithmic errors or as the technological reproduction of inequalities. Growing awareness of these harms is contributing to regulatory pressure towards greater accountability for algorithmic systems (Burke and Ho 2023; OECD 2023). In the Netherlands, the SyRi scandal was specifically cited as a reason behind the creation of a new national algorithmic regulator (Out-Law News 2023), with government bodies now required to post information about their algorithmic systems in an online register. The Netherlands has also pushed for greater transparency requirements in the EU's *AI Act* (Bertuzzi 2022), which is being closely watched by countries outside the EU.[4] While opacity and inscrutability remain major concerns for ADMs, particularly for "black box" ML-based models, algorithms are increasingly the targets of regulatory attention.

Conclusion

Welfare surveillance focuses, by definition, on the most vulnerable and marginalized. Driven by a hunt for efficiencies and the reduction of fraudulent claims, governments have increasingly turned to "AI" technologies that effectively automate population-level statistics. We know that when such systems do wrong, denying benefits and punishing those most in need, the consequences can be a matter of life and death. Even when such systems work as intended, they can contribute to structural inequalities and the disproportionate surveillance that has historically been directed against the poor.

There are ML-based techniques for welfare administration that take a more proactive approach in predicting risk and matching people and services. Forms of surveillance that make the conditions of poverty more visible can be justified as ensuring that people do not fall into "gaps"—that they are not excluded from

---

[3] This data collection is to be guided by the agency's "Necessity and Proportionality Framework," which was developed after the federal Privacy Commissioner found that an effort to collect household banking information raised "significant privacy concerns" in 2019. As it currently stands, the Office of the Privacy Commissioner (2021) has found Statistics Canada's Framework to be unsatisfactory in addressing these concerns.

[4] For example, the EU's *AI Act* has significantly influenced Canada's recently proposed AI regulations, which have been designed for "inter-operability" with the European approach (ISED 2023).

services or "miss out on resources that help them overcome their marginalisation" (Clarke, Parsell, and Lata 2021: 410). However, welfare reforms in recent decades have been justified more on the basis of efficiency and fighting fraud. In the SyRI scandal, this resulted in thousands of people being wrongly accused, but it is important to remember that welfare surveillance causes harm even when decisions are not made in error. Tying extensive scrutiny and testing to social benefits is experienced as privacy-invasive and controlling even in the absence of algorithms (Gustafson 2011; Whelan 2021). Using surveillance to detect rule violations and impose punitive sanctions (loss of services, major fines, police involvement) results in a range of harms against people who may find the rules difficult to understand or impossible to comply with while surviving on the margins (Gustafson 2011). Scholarship on welfare surveillance shows how these harms disproportionately affect specific (often racialized) groups, further compounding existing structural inequalities and reproducing structural violence (Monahan 2017).

A more critical argument against welfare surveillance is that "while surveillance 'could' technically be designed to help, considering the history and treatment of the poor, surveillance in the welfare system is highly unlikely to actually benefit them" (Maki 2021: 71). Instead, welfare surveillance has primarily been carried out to meet government's organizational goals, such as increased efficiency and limiting payments to the "undeserving." Even if we accept that welfare surveillance can have benefits, focusing on making "better" surveillance systems can distract us from more fundamental structural and political issues that underlie the production of poverty (Clarke, Parsell, and Lata 2021). Marginalization cannot be meaningfully addressed through better data collection and tinkering with algorithms to make them more fair or inclusive.

In considering the role of AI in the present and future of the digital welfare state, there are a few other things to keep in mind. First, there are the strong continuities with histories of state statistics, bureaucracy, the administrative state, and earlier forms of automation and welfare surveillance, all of which remain highly relevant. The struggles of people against the administrative categories used to make them legible, the associations between social assistance and moral worth, and the use of "neutral" technologies to enforce political preferences will be repeated even as technologies change. What we call "AI" does not deserve credit for transforming these dynamics, and elevating it as a transformative force works to reify AI as an object and contribute to the hype fueling its proliferation.

ML-based systems do have unique abilities in modeling patterns found in large datasets of many variables, which can encourage governments to engage in further data collection and the incorporation of additional databases. However, this approach does not align with the goal of making eligibility decisions based on explicit criteria in a way that people can understand. Instead, ML is used largely to identify cases that do not fit a given pattern (as anomalies), to identify cases that fit a problematic or deviant pattern (as with fraud), or for producing textual outputs based on similarities between an individual case and others it is comparable to (as with chatbots and recommendation systems). Present uses of these technologies owe far more to the contexts of the institutions in which they are employed than anything distinguishing them from previous information-processing systems. As things stand, we have more to fear from repeating historical harms of automation and the administrative state than any dangers that are specific to AI.

Acknowledgments

Funding for this research was provided by the Social Sciences and Humanities Research Council of Canada, Grant/Award Number: 430-2021-00810.

Bibliography

Alkhatib, Ali. 2021. To Live in Their Utopia: Why Algorithmic Systems Create Absurd Outcomes. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, Yokohama, Japan, May 8–13*, 1–9. New York: Association for Computing Machinery.

Alston, Philip. 2019. Digital Technology, Social Protection and Human Rights. *UN – OHCHR*, October 1. https://www.ohchr.org/en/calls-for-input/2019/digital-technology-social-protection-and-human-rights-report [accessed July 3, 2023].

Andrejevic, Mark, and Kelly Gates. 2014. Big Data Surveillance: Introduction. *Surveillance & Society* 12 (2): 185–196.

Béland, Daniel, Gregory P. Marchildon, Michele Mioni, and Klaus Petersen. 2022. Translating Social Policy Ideas: The Beveridge Report, Transnational Diffusion, and Post-War Welfare State Development in Canada, Denmark, and France. *Social Policy & Administration* 56 (2): 315–328.

Belleman, Bas, Belia Heilbron, and Anouk Kootstra. 2023. De Fraudejacht van Duo Treft Bijna Alleen Studenten Met Een Migratieachtergrond. *De Groene Amsterdammer*, June 21. https://www.groene.nl/artikel/ik-dacht-gewoon-ik-pak-je [accessed July 3, 2023].

Benbouzid, Bilel. 2019. To Predict and to Manage. Predictive Policing in the United States. *Big Data & Society* 6 (1): https://doi.org/10.1177/2053951719861703.

Bertuzzi, Luca. 2022. Once Bitten, Netherlands Wants to Move Early on Algorithm Supervision. *EURACTIV*, November 18. https://www.euractiv.com/section/digital/news/once-bitten-netherlands-wants-to-move-early-on-algorithm-supervision/ [accessed July 3, 2023].

Bovens, Mark, and Stavros Zouridis. 2002. From Street-Level to System-Level Bureaucracies: How Information and Communication Technology Is Transforming Administrative Discretion and Constitutional Control. *Public Administration Review* 62 (2): 174–184.

Brayne, Sarah. 2020. *Predict and Surveil: Data, Discretion, and the Future of Policing*. New York: Oxford University Press.

Bruijn, Hans de, Martijn Warnier, and Marijn Janssen. 2022. The Perils and Pitfalls of Explainable AI: Strategies for Explaining Algorithmic Decision-Making. *Government Information Quarterly* 39 (2): https://doi.org/10.1016/j.giq.2021.101666.

Bullock, Justin, Matthew M. Young, and Yi-Fan Wang. 2020. Artificial Intelligence, Bureaucratic Form, and Discretion in Public Service. *Information Polity* 25 (4): 491–506.

Burgess, Matt, Evaline Schot, and Gabriel Geiger. 2023. This Algorithm Could Ruin Your Life. *Wired*, March 6. https://www.wired.com/story/welfare-algorithms-discrimination/ [accessed July 3, 2023].

Burke, Garance, and Sally Ho. 2022. An Algorithm That Screens for Child Neglect Raises Concerns. *Associated Press*, April 29. https://apnews.com/article/child-welfare-algorithm-investigation-9497ee937e0053ad4144a86c68241ef1 [accessed July 3, 2023].

———. 2023. Child Welfare Algorithm Faces Justice Department Scrutiny. Associated Press, January 31. https://apnews.com/article/justice-scrutinizes-pittsburgh-child-welfare-ai-tool-4f61f45bfc3245fd2556e886c2da988b [accessed August 24, 2023].

Burrell, Jenna. 2016. How the Machine "Thinks": Understanding Opacity in Machine Learning Algorithms. *Big Data & Society* 3 (1): https://doi.org/10.1177/2053951715622512.

Busch, Peter André, and Helle Zinner Henriksen. 2018. Digital Discretion: A Systematic Literature Review of ICT and Street-Level Discretion. *Information Polity: The International Journal of Government & Democracy in the Information Age* 23 (1): 3–28.

Calo, Ryan, and Danielle Keats Citron. 2021. The Automated Administrative State: A Crisis of Legitimacy. *Emory Law Journal* 70: 797–845.

Cellard, Loup. 2022. Algorithms as Figures: Towards a Post-Digital Ethnography of Algorithmic Contexts. *New Media & Society* 24 (4): 982–1000.

Clarke, Andrew, Cameron Parsell, and Lutfun Nahar Lata. 2021. Surveilling the Marginalised: How Manual, Embodied and Territorialised Surveillance Persists in the Age of "Dataveillance." *The Sociological Review* 69 (2): 396–413.

Cobbe, Jennifer. 2019. Administrative Law and the Machines of Government: Judicial Review of Automated Public-Sector Decision-Making. *Legal Studies* 39 (4): 636–655.

Constantaras, Eva, Gabriel Geiger, Justin-Casimir Braun, Dhruv Mehrotra, and Htet Aung. 2023. Inside the Suspicion Machine. *Wired*, March 6. https://www.wired.com/story/welfare-state-algorithms/ [accessed July 3, 2023].

Coser, Lewis A. 1965. The Sociology of Poverty: To the Memory of Georg Simmel. *Social Problems* 13 (2): 140–148.

Croll, Alistair. 2018. The Year of the Objective Function. *Medium* (blog), March 2. https://acroll.medium.com/the-year-of-the-objective-function-a2353915b75e [accessed July 3, 2023].

Davidson, David, and Saskia Adriaens. 2022. In Arme Wijken Voorspelt de Overheid Nog Altijd Fraude. *VPRO*, December 20. https://www.vpro.nl/argos/lees/onderwerpen/artikelen/2022/in-arme-wijken-voorspelt-de-overheid-nog-altijd-fraude.html [accessed July 3, 2023].

Desrosières, Alain. 1998. *The Politics of Large Numbers: A History of Statistical Reasoning*. Cambridge, MA: Harvard University Press.

Dobson, Kathy. 2019. Welfare Fraud 2.0? Using Big Data to Surveil, Stigmatize, and Criminalize the Poor. *Canadian Journal of Communication* 44 (3): 331–342.

Eaton, Lynn. 2019. Is It Right to Use AI to Identify Children at Risk of Harm? *The Guardian*, November 18. https://www.theguardian.com/society/2019/nov/18/child-protection-ai-predict-prevent-risks [accessed July 3, 2023].

Elyounes, Doa A. 2021. "Computer Says No!": The Impact of Automation on the Discretionary Power of Public Officers. *Vanderbilt Journal of Entertainment and Technology Law* 23 (3): 451–515.

Eubanks, Virginia. 2018. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. New York: St. Martin's Press.

Fazi, M. Beatrice. 2020. Beyond Human: Deep Learning, Explainability and Representation. *Theory, Culture & Society* 38 (7–8): https://doi.org/10.1177/0263276420966386.

Fedyk, Anastassia, James Hodson, Natalya Khimich, and Tatiana Fedyk. 2022. Is Artificial Intelligence Improving the Audit Process? *Review of Accounting Studies* 27 (3): 938–985.

Fellowes, Matthew C., and Gretchen Rowe. 2004. Politics and the New American Welfare States. *American Journal of Political Science* 48 (2): 362–373.

Flora, Peter, and Arnold J. Heideheimer. 1981. The Historical Core and Changing Boundaries of the Welfare State. In *Development of Welfare States in Europe and America*, edited by Peter Flora and Arnold J. Heideheimer, 17–34. New York: Routledge.

Gandy, Oscar H., Jr. 1989. The Surveillance Society: Information Technology and Bureaucratic Social Control. *Journal of Communication* 39 (3): 61–76.

Gantchev, Valery. 2019. Data Protection in the Age of Welfare Conditionality: Respect for Basic Rights or a Race to the Bottom? *European Journal of Social Security* 21 (1): 3–22.

Geiger, Gabriel. 2023. How Denmark's Welfare State Became a Surveillance Nightmare. *Wired*, March 7. https://www.wired.com/story/algorithms-welfare-state-politics/ [accessed July 3, 2023].

Gilliom, John. 2001. *Overseers of the Poor: Surveillance, Resistance, and the Limits of Privacy*. Chicago, IL: University of Chicago Press.

Gustafson, Kaaryn S. 2011. *Cheating Welfare: Public Assistance and the Criminalization of Poverty*. New York: New York University Press.

Heikkilä, Melissa. 2022. Dutch Scandal Serves as a Warning for Europe over Risks of Using Algorithms. *Politico*, March 29. https://www.politico.eu/article/dutch-scandal-serves-as-a-warning-for-europe-over-risks-of-using-algorithms/ [accessed July 3, 2023].

Henman, Paul. 2020. Improving Public Services Using Artificial Intelligence: Possibilities, Pitfalls, Governance. *Asia Pacific Journal of Public Administration* 42 (4): 209–221.

Henman, Paul, and Greg Marston. 2008. The Social Division of Welfare Surveillance. *Journal of Social Policy* 37 (2): 187–205.

Ho, Sally, and Garance Burke. 2022. Oregon Dropping AI Tool Used in Child Abuse Cases. *Associated Press*, June 2. https://apnews.com/article/politics-technology-pennsylvania-child-abuse-1ea160dc5c2c203fdab456e3c2d97930 [accessed July 3, 2023].

Hoffmann, Anna Lauren. 2021. Terms of Inclusion: Data, Discourse, Violence. *New Media & Society* 23 (12): 3539–3556.

ISED. 2023. The Artificial Intelligence and Data Act (AIDA) – Companion Document. March 13. https://ised-isde.canada.ca/site/innovation-better-canada/en/artificial-intelligence-and-data-act-aida-companion-document [accessed July 3, 2023].

Keddell, Emily. 2019. Algorithmic Justice in Child Protection: Statistical Fairness, Social Justice and the Implications for Practice. *Social Sciences* 8 (10): https://doi.org/10.3390/socsci8100281.

Kuziemski, Maciej, and Gianluca Misuraca. 2020. AI Governance in the Public Sector: Three Tales from the Frontiers of Automated Decision-Making in Democratic Settings. *Telecommunications Policy* 44 (6): https://doi.org/10.1016/j.telpol.2020.101976.

Leonardi, Paul. 2011. When Flexible Routines Meet Flexible Technologies: Affordance, Constraint, and the Imbrication of Human and Material Agencies. *Management Information Systems Quarterly* 35 (1): 147–167.

Lepage-Richer, Théo, and Fenwick McKelvey. 2022. States of Computing: On Government Organization and Artificial Intelligence in Canada. *Big Data & Society* 9 (2): https://doi.org/10.1177/20539517221123304.

Lipsky, Michael. 2010. *Street-Level Bureaucracy: Dilemmas of the Individual in Public Service*. 30th anniversary ed. New York: Russell Sage Foundation.

Lyon, David. 2007. *Surveillance Studies: An Overview*. Cambridge, UK: Polity.

———. 2014. Situating Surveillance: History, Technology, Culture. In *Histories of State Surveillance in Europe and Beyond*, edited by Kees Boersma, Rosamunde van Brakel, Chiara Fonio, and Pieter Wagenaar, 32–46. London: Routledge.

Makasi, Tendai, Alireza Nili, Kevin C. Desouza, and Mary Tate. 2022. A Typology of Chatbots in Public Service Delivery. *IEEE Software* 39 (3): 58–66.

Maki, Krys. 2021. *Ineligible: Single Mothers Under Welfare Surveillance*. Black Point, CA: Fernwood Publishing.

Mann, Monique. 2020. Technological Politics of Automated Welfare Surveillance: Social (and Data) Justice through Critical Qualitative Inquiry. *Global Perspectives* 1 (1): https://doi.org/10.1525/gp.2020.12991.

Maragno, Giulia, Luca Tangi, Luca Gastaldi, and Michele Benedetti. 2021. The Spread of Artificial Intelligence in the Public Sector: A Worldwide Overview. In *Proceedings of the 14th International Conference on Theory and Practice of Electronic Governance, Athens, Greece, October 6–8*, 1–9. New York: Association for Computing Machinery.

Meaker, Morgan. 2023. The Fraud-Detection Business Has a Dirty Secret. *Wired*, March 7. https://www.wired.co.uk/article/welfare-fraud-industry [accessed July 3, 2023].

Mitchell, Melanie. 2021. Why AI Is Harder Than We Think. arXiv. https://doi.org/10.48550/arXiv.2104.12871 [accessed July 3, 2023].

Monahan, Torin. 2017. Regulating Belonging: Surveillance, Inequality, and the Cultural Production of Abjection. *Journal of Cultural Economy* 10 (2): 191–206.

Morrison, Catherine. 2023. Artificial Intelligence Becoming More Commonly Used in Public Service; Feds Update Rules. *Ottawa Citizen*, May 25. https://ottawacitizen.com/news/local-news/feds-update-rules-around-artificial-intelligence-as-use-becomes-more-common-in-the-public-service [accessed July 3, 2023].

Natale, Simone, and Andrea Ballatore. 2020. Imagining the Thinking Machine: Technological Myths and the Rise of Artificial Intelligence. *Convergence* 26 (1): 3–18.

Neumann, Oliver, Katharina Guirguis, and Reto Steiner. 2022. Exploring Artificial Intelligence Adoption in Public Organizations: A Comparative Case Study. *Public Management Review*: 1–28.

OECD. 2023. Global Trends in Government Innovation 2023. https://oecd-opsi.org/publications/trends-2023/ [accessed July 6, 2023].

Office of the Privacy Commissioner of Canada. 2021. Office of the Privacy Commissioner Compliance Monitoring of Statistics Canada's Financial Transactions Project and Credit Agency Data Project: Final Report. December 9.

https://www.priv.gc.ca/en/opc-actions-and-decisions/investigations/investigations-into-federal-institutions/2020-21/pa_20210503_sc/ [accessed July 3, 2023].

Oude Nijhuis, Dennie. 2018. *Religion, Class, and the Postwar Development of the Dutch Welfare State*. Amsterdam, NL: Amsterdam University Press.

Out-Law News. 2023. Regulator of Algorithms Established in the Netherlands. Pinsent Masons, June 28. https://www.pinsentmasons.com/out-law/news/regulator-of-algorithms-established-in-the-netherlands [accessed July 3, 2023].

Peachey, Kevin. 2022. Post Office Scandal: What the Horizon Saga Is All About. BBC News, March 22. https://www.bbc.com/news/business-56718036 [accessed July 3, 2023].

Peeters, Rik. 2020. The Agency of Algorithms: Understanding Human-Algorithm Interaction in Administrative Decision-Making. *Information Polity* 25 (4): 507–522.

Penn, Jonnie. 2018. AI Thinks like a Corporation—and That's Worrying. *The Economist*, November 26. https://www.economist.com/open-future/2018/11/26/ai-thinks-like-a-corporation-and-thats-worrying [accessed July 3, 2023].

Power, Michael. 2022. Theorizing the Economy of Traces: From Audit Society to Surveillance Capitalism. *Organization Theory* 3 (3): https://doi.org/10.1177/26317877211052296.

Raji, Inioluwa Deborah, Andrew Smart, Rebecca N. White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. 2020. Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, Barcelona, Spain, January 27–30,* 33–44. New York: Association for Computing Machinery.

Ranchordás, Sofia, and Ymre Schuurmans. 2020. Outsourcing the Welfare State: The Role of Private Actors in Welfare Fraud Investigations. *European Journal of Comparative Law and Governance* 7 (1): 5–42.

Rao, Rahul. 2022. The Dutch Tax Authority Was Felled by AI—What Comes Next? *IEEE Spectrum*, May 9. https://spectrum.ieee.org/artificial-intelligence-in-government [accessed July 3, 2023].

Reevely, David. 2021. Bots in the Bureaucracy: A New Generation of Algorithmic Decision Aids Is Moving into the Federal Public Service. *The Logic*, August 19. https://thelogic.co/news/bots-in-the-bureaucracy-a-new-generation-of-algorithmic-decision-aids-is-moving-into-the-federal-public-service/ [accessed July 7, 2023].

———. 2023. Feds Use Artificial Intelligence to Hunt for Fraud in Immigration Applications. *The Logic*, June 16. https://thelogic.co/news/feds-use-artificial-intelligence-to-hunt-for-fraud-in-immigration-applications/ [accessed July 7, 2023].

Romano, Serena. 2018. *Moralising Poverty: The 'Undeserving' Poor in the Public Gaze*. London: Routledge.

Ruppert, Evelyn. 2012. Seeing Population: Census and Surveillance by Numbers. In *Routledge Handbook of Surveillance Studies*, edited by Kirstie Ball, Kevin Haggerty, and David Lyon, 209–216. London: Routledge.

Sager, Eric W. 2021. *Inequality in Canada: The History and Politics of an Idea*. Montreal, CA: McGill-Queen's University Press.

Salganik, Matthew J., Ian Lundberg, Alexander T. Kindel, Caitlin E. Ahearn, Khaled Al-Ghoneim, Abdullah Almaatouq, Drew M. Altschul, et al. 2020. Measuring the Predictability of Life Outcomes with a Scientific Mass Collaboration. *Proceedings of the National Academy of Sciences* 117 (15): 8398–8403.

Savage, Michael. 2021. DWP Urged to Reveal Algorithm That "Targets" Disabled for Benefit Fraud. *The Observer*, November 21. https://www.theguardian.com/society/2021/nov/21/dwp-urged-to-reveal-algorithm-that-targets-disabled-for-benefit [accessed July 3, 2023].

Scheuerman, Morgan Klaus, Madeleine Pape, and Alex Hanna. 2021. Auto-Essentialization: Gender in Automated Facial Analysis as Extended Colonial Project. *Big Data & Society* 8 (2): 1–15.

Scott, James C. 1998. *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. New Haven, CT: Yale University Press.

Solano, Joan López. 2022. Data for Dignity: Requirements for the Implementation of Data Systems for Social Programs in Colombia. Translated by Carlos Alberto Arenas París. https://research.tilburguniversity.edu/en/publications/data-for-dignity-requirements-for-the-implementation-of-data-syst [accessed July 3, 2023].

Stark, Luke, and Jevan Hutson. 2022. Physiognomic Artificial Intelligence. *Fordham Intellectual Property, Media and Entertainment Law Journal* 32 (4): 922–978.

Statistics Canada. 2022. The System of National Quality-of-Life Statistics: Future Directions. *Statistics Canada*, January 31. https://www150.statcan.gc.ca/n1/pub/11-633-x/11-633-x2021006-eng.htm [accessed July 3, 2023].

Strickland, Eliza. 2021. The Turbulent Past and Uncertain Future of AI: Is There a Way out of AI's Boom-and-Bust Cycle? *IEEE Spectrum* 58 (10): 26–31.

Taylor, Simon Michael, Kalervo N. Gulson, and Duncan McDuie-Ra. 2021. Artificial Intelligence from Colonial India: Race, Statistics, and Facial Recognition in the Global South. *Science, Technology, & Human Values*, December. https://doi.org/10.1177/01622439211060839.

Touwen, Jeroen. 2023. Why Consult, Why Consent? Employers in Concertation Platforms Facing Welfare State Expansion in the Netherlands, 1920–1960. *Journal of Policy History* 35 (2): 254–280.

van Bekkum, Marvin, and Frederik Zuiderveen Borgesius. 2021. Digital Welfare Fraud Detection and the Dutch SyRI Judgment. *European Journal of Social Security* 23 (4): 323–340.

van Kersbergen, Kees, and Kjersti Metliaas. 2020. Radical Alternative Conceptualizations of the Classical Welfare State? Contrasting the United Kingdom and the Netherlands with Norway. *Social Policy & Administration* 54 (5): 813–826.

van Veenstra, Anne Fleur, Francisca Grommé, and Somayeh Djafari. 2020. The Use of Public Sector Data Analytics in the Netherlands. *Transforming Government: People, Process and Policy* 15 (4): 396–419.

Vecchione, Briana, Karen Levy, and Solon Barocas. 2021. Algorithmic Auditing and Social Justice: Lessons from the History of Audit Studies. In *EAAMO '21: Equity and Access in Algorithms, Mechanisms, and Optimization, New York, October 5–9*, 1–9. New York: Association for Computing Machinery.

Webster, C. William R. 2012. Public Administration as Surveillance. In *Routledge Handbook of Surveillance Studies*, edited by Kirstie Ball, Kevin Haggerty, and David Lyon, 313–320. London: Routledge.

Weil, Elizabeth. 2023. You Are Not a Parrot. *New York Magazine*, March 1. https://nymag.com/intelligencer/article/ai-artificial-intelligence-chatbots-emily-m-bender.html [accessed July 3, 2023].

Whelan, Joe. 2021. We Have Our Dignity, Yeah? Scrutiny under Suspicion: Experiences of Welfare Conditionality in the Irish Social Protection System. *Social Policy & Administration* 55 (1): 34–50.

Wieringa, Maranke. 2023. "Hey SyRI, Tell Me about Algorithmic Accountability": Lessons from a Landmark Case. *Data & Policy* 5 (January): https://doi.org/10.1017/dap.2022.39.

Young, Matthew M., Justin B. Bullock, and Jesse D. Lecy. 2019. Artificial Discretion as a Tool of Governance: A Framework for Understanding the Impact of Artificial Intelligence on Public Administration. *Perspectives on Public Management and Governance* 2 (4): 301–313.