Check for updates

Information Polity 25 (2020) 491–506
DOI 10.3233/IP-200223
IOS Press

Artificial intelligence, bureaucratic form, and discretion in public service

Justin Bullock[a,*], Matthew M. Young[b] and Yi-Fan Wang[c]
[a]*Public Service and Administration Department, The Bush School of Government and Public Service, Texas A&M University, Texas, TX, USA*
[b]*Department of Public Administration and International Affairs, Maxwell School of Citizenship and Public Affairs, Syracuse University, Syracuse, NY, USA*
[c]*School of Public Administration, University of Nebraska-Omaha, Omaha, NE, USA*

Abstract. This article examines the relationship between Artificial Intelligence (AI), discretion, and bureaucratic form in public organizations. We ask: How is the use of AI both changing and changed by the bureaucratic form of public organizations, and what effect does this have on the use of discretion? The diffusion of information and communication technologies (ICTs) has changed administrative behavior in public organizations. Recent advances in AI have led to its increasing use, but too little is known about the relationship between this distinct form of ICT and to both the exercise of discretion and bureaucratic form along the continuum from street- to system-levels. We articulate a theoretical framework that integrates work on the unique effects of AI on discretion and its relationship to task and organizational context with the theory of system-level bureaucracy. We use this framework to examine two strongly differing cases of public sector AI use: health insurance auditing, and policing. We find AI's effect on discretion is nonlinear and nonmonotonic as a function of bureaucratic form. At the same time, the use of AI may act as an accelerant in transitioning organizations from street- and screen-level to system-level bureaucracies, even if these organizations previously resisted such changes.

Keywords: Discretion, artificial intelligence, bureaucracy, policing, improper payments

Key Points for Practitioners

- Organizations considering implementing AI need to account for its goodness of fit with respect to both organizational form and the degree of discretion required for a given task.
- The decision to use AI may have consequences for the organization's broader form and function over time.
- AI's impact on bureaucratic discretion is more pronounced in organizations where discretion is already systematically limited, and more dynamic in organizations where discretion is central to agent performance.
- AI may act as an accelerant in transitioning organizations from street- and screen-level to system-level bureaucracies, even if these organizations previously resisted such changes.

1. Introduction

The diffusion of information and communication technologies (ICTs) over the past several decades has profoundly changed administrative behavior in public organizations (Fountain, 2001). As their capacity for data generation, storage, and processing continued to grow exponentially over the late 20th and early

*Correspoding author: Justin Bullock, Public Service and Administration Department, The Bush School of Government and Public Service, Texas A&M University, Texas, TX, USA. E-mail: jbull14@tamu.edu.

1570-1255/20/$35.00 © 2020 – IOS Press and the authors. All rights reserved

21st centuries, ICTs became increasingly central to organizational decision-making processes (Gil-Garcia et al., 2018). This increased emphasis on technologically-mediated decision-making directly affects the quantity and quality of discretion – the latitude afforded individuals with delegated responsibilities to use their judgement when making a decision – afforded public administrators (Bullock, 2019).

From a top-down, managerial perspective, one of the principle benefits of ICT-mediated discretion is the standardization of bureaucratic behavior. Standardization reduces the risk of moral hazard endemic to the delegation of responsibilities and the corresponding degree of discretion afforded both street-level and mid-level agents. This standardization by ICT may also improve decision making and outcomes (Greer & Bullock, 2018). A bottom-up perspective, by contrast, views these changes as the erosion of professionalization, particularly among street-level bureaucrats who interface directly with the public and have considerable discretion to determine how to provide services and enforce policies (Lipsky, 2010). As street-level bureaucrats' work processes transition from person-to-person interactions to interfacing with ICTs, service-providing public organizations transition from street- to screen-level, and ultimately system-level, bureaucracies (Bovens & Zouridis, 2002). More recent theoretical and empirical work adds further nuance to this theory of change, calling attention to how task, technology, and organizational characteristics contextualize the effect of "digital discretion" in practice (Hannah-Moffat, 2019; Peeters & Schuilenburg, 2018; van Eijk, 2020).

With respect to technological context, recent advances in Artificial Intelligence (AI) have led to its increasing use in both private and public organizations (Rahwan et al., 2019). In particular, these uses involve a particular form of AI: machine learning-based applications that adjust their decision criteria dynamically using stochastic optimization processes. The exponential growth in AI's use, capabilities, and risks has not gone unnoticed (Bostrom, 2014; Dietvorst et al., 2018; Zou & Schiebinger, 2018). Much of this work has focused on AI's potential impact on the economy, the labor market, and national defense (Frank et al., 2019; Frey & Osborne, 2017; Korinek & Stiglitz, 2017; McClure, 2018). Another vein of related work focuses on the use of machine learning-based decision support systems in criminal justice for bail setting and sentencing (Binns, 2019; Hannah-Moffat, 2013, 2019; Hannah-Moffat et al., 2009). Less attention, however, has been paid to AI's implications for public organizations and the governance of the public sector more broadly. Moreover, much of the related work focusing on public organizations is framed around the concepts of "big data" (Brayne, 2017; Ferguson, 2017) or "smart/algorithmic governance" (Eubanks, 2018; Gil-Garcia et al., 2014). Missing from this discussion is explicit attention to the implications for the public sector that arise both from allocating discretion away from humans to AI, and the corresponding changes to the form of bureaucratic organizations, often towards systems level bureaucracies.

This is now beginning to change. Recent scholarship has begun to develop frameworks for describing, understanding, and evaluating the use of AI to perform tasks in public sector and organizational contexts (Bullock, 2019; Drexler, 2019; Young et al., 2019). These scholars argue that, unlike other forms of digital discretion, AI is best understood as an agent embedded in an organizational context that executes tasks using stochastic, or non-deterministic, approaches. These approaches require learning, sensing, and probabilistic reasoning, and allow AI to perform tasks that are ill-suited for deterministic expert systems algorithms, which are built using logic-based rules. These complex and contingent tasks, and the institutional framework necessary to enable (and constrain) their execution, were previously the exclusive purview of human agents. However, for a growing set of increasingly complex tasks involving decision-making, this is no longer the case – the man in the Mechanical Turk has been made redundant. Exactly how the use of AI affects the nature of discretion in public organizations, and how this use is conditioned in turn by organizational context, however, is largely unknown.

We address this lacuna by identifying uses of AI in different organizational contexts in the public sector and analyzing them through the theoretical lenses of systems bureaucracies, digital discretion, and artificial discretion. Two research questions motivate this work: How is the use of AI both changing and changed by the bureaucratic form of public organizations, and what effect does this have on the use of discretion within these organizations? In answering these questions, we examine both the use of AI and the form of bureaucracy for predictive policing, facial recognition, and criminal pattern recognition, and how the United States Centers for Medicare & Medicaid Services (CMS) implements AI in its antifraud and improper payment reduction efforts.

These two domains are chosen as initial, illustrative explorations into the use of artificial intelligence and form of bureaucracy for several reasons. Perhaps most importantly, both involve the current, active use of AI in core organizational functions, as well as past and potential future expansion of its use over time. Both cases are also highly salient in terms of the cost of service provision and the impact on individual citizens and society. Finally, the two domains provide reasonable variation across where the organizations are located along a continuum from street-level form to system-level form. These factors, taken together, make our cases important early explorations into the advance of both artificial discretion and the system-level bureaucracies in the public sector.

We proceed by first summarizing our theoretical framework, which draws upon Bovens & Zouridis' (2002) theory of system-level bureaucracy, its extension in Busch and Henriksen's (2018) digital discretion, and its relationship to the AI-specific focus of Young et al.'s (2019) theory of artificial discretion. We then turn to our cases, first providing descriptive analysis of how AI is currently being used *in situ*, and then applying our framework to test its explicative power. We conclude with a distillation of our findings coupled with suggestions for future research.

2. Theoretical framework

Scholarship on how ICT implementation affects discretion in public organizations is multifaceted, with theoretical and empirical work alternatively focusing on the national level (Fountain, 2001), local level (Busch & Eikebrokk, 2019; Tummers & Bekkers, 2014), and different policy domains including welfare (Hetling et al., 2012; Houston, 2015; Pors, 2015), criminal justice (Hannah-Moffat, 2019; Marks et al., 2015; van Eijk, 2017; Završnik, 2019), and others. A unifying feature across empirical contexts is the recognition that as public organizations increase their use of ICT, and as the tools themselves increase in their scope and capacities, the result is a reshaping of administrative work itself, leading to new forms of bureaucracy (Bovens & Zouridis, 2002; Danaher, 2016; Dunleavy et al., 2006; Peeters & Schuilenburg, 2018; Peeters & Widlak, 2018).

One early and highly influential work is Bovens and Zouridis' (2002) theory of ICT-driven organizational change from street- to system-level bureaucracies. The theory explains that as ICT become increasingly central to the execution of routine tasks, the nature of this work changes from person-to-person interactions to person-to-computer interactions, a stage termed "screen-level bureaucracy." The theory calls explicit attention to the impact of this shift to screen-level bureaucracies for street-level bureaucrats – those whose responsibilities required direct interaction with the public – and the relatively high degree of discretion they are afforded absent these technologies. Bovens and Zouridis go on to argue that ICT may eventually become not only the primary but also the *sole* structure for communication, flow of information, and organization of work tasks; at this stage of development the organization is said to be a system-level bureaucracy with little to no place for street-level bureaucrats as traditionally defined (Bovens & Zouridis, 2002).

Importantly, the transition from traditional to screen- and ultimately system-level bureaucracy is not conceived of as an irresistible force. Rather, the propensity for this change is explicitly contingent on the nature of the work tasks required, and by extension the initial degree of discretion afforded to its agents and other organizational characteristics. The examples provided in the original theory are education and policing; the authors acknowledge that these domains of public service provision are extremely complex and contingent, and require empowering street-level bureaucrats – teachers and officers, respectively – with a high degree of discretion (Bovens & Zouridis, 2002). This conditioning of expectations was later enriched in Busch and Henriksen's (2018) review and extension of the literature on ICT and discretion that followed the publication of the theory of system-level bureaucracies. Their analysis calls attention to both the technological characteristic of the ICT and the importance of the unit of analysis – micro, meso, or macro – within the organization where the ICT-mediated task takes place. It also reveals a public sector still staffed with street-level bureaucrats, but with virtually all of them embedded in varying forms of screen-level bureaucracy, with some tasks mediated by ICT and some automated entirely.

Other studies of the relationship between ICT and discretion also call attention to the importance of the technology's characteristics. For example, technology does not just curtail discretion; it can enable it as well by affording new capabilities – whether intentional or not (Buffat, 2015; Hupe & Buffat, 2014; Thunman et al., 2020). In describing and explaining what they term "digital discretion," Busch and Henriksen note that technological advances in ICT – most notably AI – pose unique questions for the structure of public organizations and the use of discretion within them vis-a-vis automation.

The question of AI's impact is addressed directly in Young et al.'s (2019) framework for understanding and evaluating the use of AI in public organizations, which they term "artificial discretion." Central to their argument is the distinction between machine learning-based AI and traditional, expert systems-based algorithmic automation. Unlike expert systems, most modern AI applications – and all of the more powerful, complex forms, e.g. neural networks – are stochastic. As a consequence, it is not possible to reverse engineer a decision reached by AI with perfect certainty despite having complete, perfect knowledge of all available input variables/features. This is reflected in how computer scientists evaluate AI decision performance: AI is evaluated probabilistically (e.g., a 93% success rate at making the correct classification decision), rather than in terms of absolute fidelity. Moreover, the final architecture of any AI system after it has been trained is fundamentally unique – even when the trainer uses identical initial architectures and training data (Russell & Norvig, 2009).

This machine learning characteristic of AI is what makes it capable of performing tasks that were previously considered solely the domain of human agents. Furthermore, AI's stochastic behavior introduces the same information asymmetries between manager-or-organization-as-principal and AI-as-agent endogenous to administrative discretion as traditionally understood. For example, consider the use of AI for facial recognition. The outcome of a discrete task in this context is to determine whether an image of a face can be reliably paired to a separate image with associated metadata about the individual (e.g., name, age, known residence, etc.). The AI-as-agent ultimately produces a decision: it reports either a match/a set of possible matches, or that it failed to find a credible match. On a sub-task level, this process involves comparing n pre-known images of faces against the new image; at a level below this the AI agent is performing k subroutines, such as evaluating the geometries, shade gradients, and other constituent elements of the images against each other. Thus, from an organizational perspective, AI agents possess the same afforded discretion for this task as any human agent: they can be trained, and held to performance standards, but the precise outcome of any discrete task is not knowable *a priori*.

Young et al. (2019) also integrate the contextual effects of task type and unit of analysis previously identified in both the theory of system-level bureaucracy and digital discretion. They then derive propositions for how AI use is likely to vary based on these contexts. For tasks with low levels of required

Table 1
Theoretical expectations: AI, discretion, and bureaucratic form

Required discretion	Bureaucratic form	Ratio of human to artificial discretion	Common forms of artificial discretion	Magnitude of change in bureaucratic form	Magnitude of change in share of discretion
High	Street- and screen-level	High and waning	Data collection, decision support systems	High; accelerates transition from street- to system-level	Low
Low	System-level	Low and waning	Data collection, decision support systems, automation	Low; already possesses system-level traits	High

discretion, automation by AI may be possible removing much of the need for human discretion. For tasks with medium levels of required discretion, AI is mostly likely to be used as a tool of prediction, finding patterns in multidimensional data, and generating new insights into variables that affect task execution. Finally, neither automation nor prediction may be an appropriate use of AI for tasks that require a high level of discretion. Instead, AI may be used to generate better data from unstructured inputs such as images, sensors, and text that can aid human decision makers in executing a task more effectively (Young et al. 2019).

Along with the degree of discretion, the unit of analysis at which a task is embedded within a bureaucracy is also helpful for identifying what types of tasks are likely candidates for AI. The organizational context refers in part to the scope of influence or effect a given task has – localized (micro), organizational (meso), or institutional (macro). The micro level of analysis consists of tasks that are implemented at the individual level, in general at lower levels in the organization such as street-level bureaucrats. Meso-level tasks "shape and affect the organizational environment in which individual agents are embedded" (Young et al. 2019; 5). Macro-level tasks are institutional or enterprise level tasks that include the contextual factor of formulation of rules and general policies for the organization (Busch & Henriksen, 2018).

The organizational context includes not only the level (micro, meso, and macro) within the organization at which the task is embedded, but also the location of the organization along the continuum of bureaucracy form from street to system-level, which also shapes the organizational context for decision making and task execution. The location of the organization along this continuum is a function of the type of tasks that dominate that organization and its policy domain. If the tasks, in general, require less discretion then it is more likely that ICT tools have begun playing a decisive role for the decision making of that organization (it is *easier* to adapt these simpler tasks to automation by machine), while tasks that require more discretion have been more likely to rely on human discretion and keep the street-level bureaucracy form.

To summarize, the theory of system-level bureaucracy is further extended and developed by both digital and artificial discretion. The theoretical expectations are summarized in Table 1. Viewed synthetically, the extension of these theories of ICT-driven organizational change in the public sector to the specific case of AI lead to empirically testable expectations for how AI's implementation will both be shaped by organizational and task contexts and, in turn, affect the way discretion is allocated and used within these organizations and, by extension, the organization itself. Accordingly, the characteristics of the tasks – in this case, the required level of discretion – that dominate a bureaucracy influence both the location of the bureaucracy on the street to system-level continuum and the allocation of uses of artificial discretion. In street- and screen-level bureaucracies we expect to see AI used collaborative with human agents as a decision support tool and for gathering new information. In system-level bureaucracies we expect to see AI used more predominantly to automate tasks, reducing or eliminating human agents in decision processes.

This theoretical framework integrates insights from traditional ICT literature, advances in our understanding of digital discretion, and recent work highlighting the unique characteristics and impacts of AI and its use by public organizations. Building from previous literature we have argued that AI presents unique challenges to public administration. For policy domains that are characterized predominantly by high discretion tasks we expect more traditional street-level and screen-level bureaucracies to dominate and for a less reliance on artificial discretion and when artificial discretion is employed the uses are to gather more information and to provide decision support. However, for policy domains that are characterized by low discretion tasks we expect system-level bureaucracies to emerge and for a stronger reliance of artificial discretion to develop and for artificial discretion to take the form of automation.

3. Case analyses: Policing and public health insurance administration

To test the explanatory power of this framework, we selected two important policy domains: policing and public health insurance. These domains were chosen to provide cases with strongly differing organizational contexts and associated requirements for discretion (Seawright & Gerring, 2008). Both policy domains have long histories to observe any shifts along the continuum from street-level bureaucracy to system-level bureaucracy, and the impacts of artificial intelligence. We collected government reports and reviewed related prior research for our analysis.

3.1. The use of AI in policing

Policing is one of the largest public sector domains where AI is being implemented, capturing the attention of the public, elected officials, and academics. Broadly speaking, there are three distinct ways in which police organizations (hereafter, departments) are using AI. The first is to mine administrative and other data to forecast future criminal activity, or what is commonly referred to as "predictive policing." The second identifies underlying relationships between multiple past crimes to establish when there is likely a "criminal pattern" – crimes committed by the same individual(s). The third use consists of facial recognition-based systems to generate new information at speeds and scales that are impossible when using human agents. This section discusses each of these uses in turn and contextualizes them in the broader political and institutional history of police policy and practice.

Predictive policing is an umbrella term for the use of data analytics and operations management to first forecast future criminal activity, and then deploy police resources proactively to prevent occurrences, facilitate arrests (Meijer & Wessels, 2019) and predict recidivism (Dressel & Farid, 2018). The basic premise of proactive policework is as old as professional police forces, if not older. What distinguishes predictive policing from more traditional strategies is both its automaticity and scalability. Whereas the knowledge and insight necessary to know where to focus preventative attention was previously the domain of high-performing and experienced police officers, predictive policing leverages large and high-dimensional datasets to effectively automate the application of that knowledge across an entire jurisdiction.

Predictive policing can employ either places in space and time, or individuals as the unit of analysis for its models (Bennett et al., 2018; Norton, 2013; Perry, 2013). Place-based predictive policing is concerned with forecasting the likelihood of future criminal activity across subunits of geographic space, while individual-based approaches are used to forecast the likelihood that a given person will be involved in a criminal act in the future. Irrespective of the different units of analysis, all predictive policing approaches are predicated on both the analysis of large volumes of complex and multifaceted data, and the normative

argument that public safety organizations should be proactive rather than reactive with respect to criminal behavior (Meijer & Wessels, 2019).

Two examples of private vendors offering place-based predictive policing systems are PredPol and Azavea. PredPol was founded by researchers from the University of California Los Angeles, and its eponymous software, first used in 2012, is based on their work with the Los Angeles Police Department in 2008 (Brayne, 2017). Azavea was founded by a former crime analyst with the Philadelphia Police Department. Its software, HunchLab, is an extension of the Crime Spike Detector system developed as a geospatial add-on to the department's COMPSTAT system (Benbouzid, 2019). Place-based predictive policing rapidly diffused from these initial cases to jurisdictions throughout the United States (Ferguson, 2012, 2016). Known implementations of individual-based predictive policing include the City of Chicago's Strategic Subject List and the use of Palantir's security/military intelligence platform by the cities of New Orleans, New York, and Los Angeles (Ferguson, 2017; Saunders et al., 2016). Frey and colleagues (2018) note that the NYPD is known to engage in online surveillance of individuals, and that law enforcement professional associations estimate that approximately 95% of member departments use social media (though it is unknown how much of that use is specifically for predictive policing systems vs. other uses of data collection and integration).

A related but ontologically distinct use of AI in policing is to generate new information by identifying criminal patterns in administrative data and individuals from photographic and video data. Both of these applications leverage AI's capacity for pattern recognition. Identifying criminal patterns from discrete crime report data is perhaps the *sine qua non* of pattern recognition in law enforcement, whether assisted by technology or not. But the technology behind facial recognition also relies on pattern recognition. Faces are identified by decomposing images into constituent vectors, whose characteristics and relationships within a given face can be used to find matching patterns of vector characteristics in preexisting databases that link images of individuals' faces to administrative data.

One example of using AI to identify criminal patterns is the New York Police Department's (NYPD) deployment of Patternizer. Patternizer is a program designed to identify patterns in reported property crimes (burglary, robbery, grand larceny) and make recommendations to officers when multiple crimes are likely 'patterns' for further investigation and follow-up. The program is integrated into NYPD's 'Domain Awareness System' (DAS), which is "a citywide network of sensors, databases, devices, software, and infrastructure" (Levine et al., 2017). Patternizer is available to NYPD officers through the DAS desktop application. Its implementation included hiring 100 new civilian crime analysts who were trained on how to use Patternizer at the start of their job duties (Chohlas-Wood & Levine, 2019).

Brayne's (2017) ethnographic analysis of the Los Angeles Police Department's (LAPD) use of Palantir's platform provides another example of using AI to identify patterns. The Palantir platform automatically generates notifications when a user either runs a query or enters new data if the content of either use is similar to other users' – even when that other user is in a different police jurisdiction across the country. The LAPD also uses facial recognition programs to automatically scan all individuals who pass within some 600 feet of one of the department's networked camera installed throughout the city (Ferguson, 2017; Garvie, 2016). While the LAPD's system is real-time, other local, county, and state law enforcement agencies throughout the United States are also known to use *ex post* facial recognition systems, where prerecorded video is analyzed and compared against existing data to identify individuals of interest (Garvie, 2016). As with predictive policing, there is already a private sector market for facial recognition software. Several vendors, including NEC and 3M, also offer systems that claim to allow for real-time recognition.

Taken individually, these examples of AI in policing illustrate how new technologies enable, constrain, and otherwise alter preexisting institutions around the exercise of discretion by agents within these

organizations. Viewed systemically, the use of AI in policing can be seen as an imbricated or fractal set of instances of applied pattern recognition. Image data are analyzed to match individuals to particular places in space and time. AI systems mine large-n, multidimensional data to identify patterns of criminal activity across space and time. Other AI systems plumb these and other data to forecast likely criminal activity forward in time across space and individuals. Before ICT were ubiquitous, pattern recognition in policing was largely the domain of veteran foot patrolmen (the platonic street-level bureaucrat) and domain experts in the form of detectives and analysts. As the scope of AI use increases, and as disparate AI systems within policing organizations become more integrated, these organizations shift further towards system-level bureaucracies.

3.2. The use of AI in public health insurance administration

One of the more expensive public services provided by government is public insurance. Publicly funded insurance spans areas such as health, education, income, and old age. One financially large and societally important public insurance in the United States is public health insurance. The two largest programs are Medicare and Medicaid and they are administered by the Centers for Medicare and Medicaid Services (CMS). These programs are also systematically beset with large dollar amounts of improperly paid claims. As with all insurance, public health insurance suffers from moral hazard, adverse selection, and other incentive misalignments that may lead to improper payments, either due to accident or fraud. Public health tasks, more generally, often require high levels of discretion to make determinations about who should receive health care, how much they should receive, and how much should be charged. However, the processing of the insurance claims that correspond with the received health care is much more *deterministic* and rule bound. As the organization charged with processing these insurance claims, CMS processes an enormous amount of low discretion, relatively routine tasks.

In fiscal year 2019, the outlays of Medicare and Medicaid were projected to be $800 billion and $623 billion, respectively (Centers for Medicare and Medicaid, 2019). The Department of Health and Human Services (HHS) estimates that nearly 8 percent of insurance spending on Medicare and Medicaid is improperly paid (Department of Health and Human Services, 2018). The nature of insurance data contributes to the difficulty in detecting insurance fraud and improper payments (Dora & Sekharan, 2015). Insurance claims data is vast, and it is difficult for humans to identify patterns that indicate improper or fraudulent payments. Additionally, insurance claims data come in various forms and formats from different stakeholders. These features make insurance data difficult for humans to parse but are precisely the type of inputs for tasks that artificial discretion tools are purpose-built for. Two specific examples of uses of AI within the delivery of public health insurance can be found in the Fraud Prevention System (FPS) utilized by CMS and the Health Care Fraud and Abuse Control Program (HCFAC).

CMS deploys several technological tools to reduce improper payments and fraud in the claims delivery system. They include tools for data collection, a fraud prevention system (FPS), pre-payment claims edits, and post-payment error detection. The ICTs in this case play important roles throughout the FPS process. Each of these systems involves the collection of immense quantities of data in a variety of formats. This is another strength of AI, which can be used to quickly integrate heterogeneous data into a single relational database or other structure (Bauder & Khoshgoftaar, 2018; Bostrom, 2014; Dora & Sekharan, 2015). Additionally, AI can be used as a decision support tool for payment error detection. Another of AI's strengths is finding patterns that might identify a payment error. AI tools in this domain could also be used to predict when and where payment errors are likely to occur, providing overall improvements to the public health insurance system. CMS further leverages this capacity to reduce improper payments by

sharing aggregated data with multiple stakeholders through Program Integrity Centers (PIC) (Centers for Medicare and Medicaid, 2017). In the FPS, CMS uses AI in the form of anomaly-detection and predictive models to prevent and identify prior potential fraudulent claims. Anomaly-detection models identify unusual patterns by detecting fraud in the claims data and comparing these abnormal claims to more typical, non-improper claims. The predictive models analyze the historical data to detect fraud and improperly paid claims (Government Accountability Office, 2012, 2015, 2018c, 2018a, 2019).

Another example program that utilizes artificial discretion to combat improper payments and fraud is the Health Care Fraud and Abuse Control Program (HCFAC). HCFAC is a partnership with the Department of Health and Human Services (the broader department in which CMS is situated) and the Department of Justice that works to detect and prosecute health insurance fraud and abuse (Government Accountability Office, 2018a, 2018b). HCFAC utilizes several automated processes and tools, including anomaly-detection and predictive models, to identify potential fraudulent claims for action. This program recovered $2.6 billion in the fiscal year 2018 and acquired $4 for every $1 spending on this program (Department of Health and Human Services, 2019).

Public health insurance processing continues to suffer from fraud and improper payments. In the US, CMS is responsible for ensuring the accuracy of these claims. Furthermore, health insurance processing tasks are often rule bound and deterministic. Either an insurance claim is paid correctly give the insurance rules, or it is paid incorrectly. This type of claim classification is a strength for AI. Additionally, CMS' extensive history on working with large amounts of numerical data, along with the fact that many tasks require low levels of discretion, highlights a bureaucracy that was primed and ready for the transition to a system-level bureaucracy. The use of the Fraud Protection System and both the Program Integrity Centers and the Health Care Fraud and Abuse Control Program are all examples of a bureaucracy making the transition to a system-level bureaucracy.

4. Discussion

In our two illustrative cases, we have described some of the uses of AI that can be found in both the policing and public health insurance domains. Through these cases we demonstrate how the use of AI affects the exercise of discretion and how both organizational and task characteristics influence the uses of AI in turn. We further show that the use of AI itself may pressure bureaucracies to transition from street- and screen-level to system-level bureaucracies. These effects carry significant potential consequences for organizational transparency and accountability: the complexity and automaticity of system-level bureaucracies makes it difficult to provide satisfactory ex post explanations of their decisions. On the

Table 2
Case study findings related to theoretical expectations

Policy domain	Required discretion	Bureaucratic form	Ratio of human to artificial discretion	Common forms of artificial discretion	Magnitude of change in bureaucratic form	Magnitude of change in share of discretion
Policing	Moderate to high	Street- and Screen-level	High and in flux	Data collection, decision support systems, automation	High; accelerates transition from streeet- to system-level	Low to Moderate
Insurance process-ing	Low	System-level	Low and waning	Data collection, decision support systems, automation	Low; already possesses system-level traits	High

other hand, these effects may yield significant gains in efficiency and effectiveness for many of the outcomes these organizations are held accountable for providing to the public. Table 2 summarizes our findings, which we discuss in more detail above.

The use of AI affects the exercise of discretion. CMS is a system-level bureaucracy that employs AI to expand both the scope of its work efforts and its efficiency, or the rate at which it can perform its required tasks. While humans remain integral to the overall structure, many of CMS tasks no longer require human discretion or with the use of AI and digital discretion have expanded into tasks that would not be feasible for humans to complete. This use of AI, then, constitutes a process improvement that does not require the organization to adapt its system-level behavior by having to create new positions for human agents. Here it may seem that human discretion is often completely absent, as system level bureaucracy theory would suggest. Indeed, one use of automation by CMS is the identification of improperly paid public health insurance claims and automated prepayment edits to help prevent fraud. AI can also be used to predict claimants' risk levels or where patterns of fraud may be occurring by utilizing predictive analytics. Finally, with the widespread digitization of personal and event data, AI is used to gather new types of data on claimants and claims. AI can be used in this way to help gather new types of data to be considered as part of the claims process.

Street-level bureaucrats in law enforcement settings, however, are still entrusted with a great degree of discretion, both explicit and tacit, in performing their duties because of both the high degree of uncertainty associated with many of their tasks and the political importance of their outcomes. This discretion can be problematic. For example, information asymmetries and resulting moral hazard between street-level officers and their managers, and social and political harm when officers make mistakes or abuse their discretion to needlessly punish individuals and socioeconomic groups (Pierson et al., 2020; Selbst, 2017). Additionally, COMPSTAT's effectiveness is as much about its ability to regulate and make for predictable, measurable behavior on the part of street-level officers as it is about reducing crime rates (Benbouzid, 2019). Studies on the applied use of predictive policing have shown that it is recognized both by management and by street-level employees as a way for management to exert more control and have more certainty about how the street-level officers are spending their time (Brayne, 2017).

All three of AI's identified uses in law enforcement encroach on tasks that have historically been the purview of the officer on the street or an entire professional class of law enforcement employees called crime analysts. Historically, crime forecasting would either involve an ad-hoc, experience-based decision by veteran officers, or crime analysts using tools like Geographic Information Systems (GIS) to produce hotspot mapping and other forecast outputs. Similarly, finding criminal patterns across multiple cases was either the job of a detective to notice patterns based on experience, or a professional crime analyst. Facial recognition in policing was classically part of the community policing model of foot patrolmen getting to know the residents in a small area and knowing the individuals on the street by virtue of familiarity, or stopping unknown individuals and asking for identification.

The organizational context shapes the way AI is used. One example for policing is the institutional history around investments in information technology in American law enforcement. This can be traced back to the 1968 Omnibus Crime Control and Safe Streets Act, which created a federal resource stream for local law enforcement to fund investments in information technology. The causal theory behind the allocation of federal funds for local police to invest in ICT was that too little and too disjointed information was an impediment to solving and preventing crime. As investments in ICT continued and technological capacities grew over time, most medium- to large-sized departments found themselves with a surfeit of data stored in multiple, disjointed systems in a myriad of formats (Andrejevic, 2017). The COMPSTAT program developed by William Bratton in New York in the late 1990s/early 2000s

and brought with him to Los Angeles systematically brought these – and yet more – data together in a model of *ex post* performance measurement and management that was rapidly emulated throughout the US and in the United Kingdom (Sherman, 2013). The COMPSTAT system is also the direct technological progenitor to modern predictive policing: as mentioned earlier both PredPol and HunchLab platforms were spun-off from previous work on COMPSTAT-based software. Given the amount of data that law enforcement agencies have, and the political pressure to put those data to use, the diffusion of AI in this context seems almost inevitable.

The current use of predictive policing represents a high-discretion task. As an organizing methodology, predictive policing has broad and pervasive impacts on organizational structure, behavior, operations, and performance. At the same time, the discretion required to forecast likely criminal activity and direct armed police officers to intervene is necessarily broad. Criminal pattern recognition is a medium-discretion, micro-level task. It replicates work done previously by individual detectives and/or crime analysts, but also requires sophisticated pattern recognition and analysis that do not currently lend themselves to complete automation. Finally, facial recognition involves a relatively low level of discretion, but the organizational context is difficult to map due to the scaling effects of AI. Whereas a patrol officer recognizing a suspect vs. an innocent passerby was previously a micro-level task, that is difficult to equate to the city-wide dragnet capabilities of AI-enabled policing.

For public health insurance one important organizational context is the abundance of numerical data. These numerical data have a history of being organized for use by quantitative analysis, and the data have already been regularly analyzed by digital computing systems. This history of data collection, organization, and analysis suggest that public health insurance tasks may be even more particularly suitable to the use of artificial discretion and the use of automation of some tasks. Tasks that involve processing health insurance claims lend themselves to task characteristics that contain clear rules, clear processes, and can be numerically described. The boundedness of many public health insurance claims tasks is also narrower and often more clearly defined than, for example, tasks with the delivery of broader health services where tasks often require dynamic interaction with human clients. This suggests that there are likely significant sets of tasks within the provision of public health insurance that have lower levels of required discretion, which also may lead to more automation and further development of the system-level bureaucracy.

We find that improper payment and fraud detection, tasks that require low levels of discretion and are embedded at the micro level of organization, are beginning to be automated. We also find tasks such as predicting hubs of fraudulent or criminal activity are being augmented and enhanced with predictive analytics tools. These tasks require more discretion and thus predictive analytics working together with human analysts, rather than complete automation of the tasks. Finally, we see that AI can also be used to gather more data for insurance claims and process vast amounts of video for police departments to help provide more detailed data to decision makers across these policy domains.

The use of AI may also affect organizational structure. We observe the role AI itself may play in moving organizations along the continuum towards system-level bureaucracies. Thus, it seems that by developing the organizational infrastructure for the use of AI furthers the organizational along its path towards a more decisive role of ICTs in the decision-making function. Police departments, which are much closer to traditional street-level bureaucracies, but also incorporate significant elements of screen-level bureaucracies, use AI for numerous types of tasks. Policing tasks generally require more discretion, but we still find that AI is used for data collection, decision-support, and automation of some tasks. This spread of AI throughout policing appears to be shifting police departments to more widespread and decisive uses of ICT and thus towards the form of system-level bureaucracy. For CMS, within this task

set of reducing fraud and improper payments in the delivery of public health insurance, we find that CMS is using AI to help gather and organize insurance claims data, analyze that data for patterns, and automate large parts of the search function for identifying problematic claims. These are the types of tasks that are prevalent in a system-level bureaucracy, and where our theory expects to find AI employed in those organizations.

Jointly, this suggests that the use of AI encourages a transition into a system-level bureaucracy and that when system-level bureaucracies use AI it is often to curb and restrict human discretion, in a manner consistent with the original theory of system bureaucracies. However, our findings suggest that these pressures are neither linear nor monotonic. Instead, they are moderated by the preexisting bureaucratic form.

AI as implemented in both policing and insurance administration reduces human discretion in some cases. But this effect is much more substantive at the margin in insurance administration because its system-level bureaucratic form eliminates potential discretion-increasing spillover effects. On the other hand, AI's use for various forms of pattern recognition in policing has a more balanced effect on discretion: opportunities for human agents to exercise discretion are both destroyed and created. The reciprocal effect of AI on bureaucratic form, however, is much stronger in the case of policing than it is for insurance administration. While AI-augmented policing still retains features of both street- and screen-level bureaucracy, the logic of scalability and efficiency provided by AI-focused systems integration signals an acceleration towards system-level bureaucracy in organizations previously identified as being highly resistant to such changes. Meanwhile, AI's effect on bureaucratic form is negligible in the case of insurance administration. We attribute this null effect to the inherent congruence between system-level bureaucracies and the automaticity that AI enables and enhances; they are naturally complementary systems.

5. Conclusion

In the final edition of *Administrative Behavior*, Simon (1997) argued that while ICTs were diffusing throughout organizations in the form of personal computers and the internet, they had not yet *drastically* altered the bureaucratic form or decision-making processes – including discretion – of administrative organizations. By 2002, Bovens & Zouridis argued that, in fact, for many bureaucracies, ICTs *had* drastically altered the bureaucratic form, giving rise to the system-level bureaucracy and reshaping the use of discretion in the process. These system-level bureaucracies take a different shape than their street and screen-level forms. Discretion is much more centralized, tasks more routinized, and few if any street-level bureaucrats are present. Since the dawn of the 21st Century, ICTs have continued to spread throughout bureaucracies, giving rise to more system-level bureaucracies, and the ICTs themselves have continued to increase in their capabilities.

Modern AI has further increased the capabilities of ICTs to conduct the work of public organizations. Furthermore, AI's unique machine learning characteristics allow it to operate in a non-deterministic fashion, making it much more akin to the decision-making processes of human discretion. In this article we have examined how the use of AI alters the exercise of discretion within an organization and its bureaucratic form, as well as how the form also moderates the processes for which AI is used. Drawing from extant theoretical literature on ICTs, the characteristics of AI, and its relationship to discretion and organizational context in the public sector, we argue that when AI is used within system-level bureaucracies like the Centers for Medicare and Medicaid Services it is often used for automation

of tasks and the further reduction of human discretion. This stems in part from the fact that system-level bureaucracies consist of tasks that require lower levels of discretion and have already reshaped the decision-making process to make it more amendable for automation. Street- and screen-level bureaucracies such as those found in policing, on the other hand, consist of more tasks that require higher levels of discretion and thus they have been more resistant to a transition to a system-level bureaucracy. However, we find in our policing case that AI is adopted for many uses across many tasks, and thus we suggest that the use of AI may also represent a sea change in the push to transform these organizations towards a system-level bureaucratic form.

At the intersection of AI, bureaucratic form, and discretion, we find that as AI is deployed within public organizations, discretion is altered. In system-level bureaucracies discretion continues to shift away from street-level bureaucrats to software developers and machines. In bureaucracies that have not transitioned to a system-level bureaucracy, AI does reduce some forms of discretion by automating tasks. But at the same time, it also enables new uses of discretion by providing new capabilities in data gathering, decision support, and predictive analytics. Additionally, the use of AI itself requires decision making processes and bureaucratic forms that better integrate ICT infrastructure into the decision-making process. This infrastructure requirement of AI use, along with the variety of tasks for which AI may then be applied, suggest that AI implementation accelerates organizations' transition towards a system-level bureaucracy, even when those organizations traditionally resist such efforts.

Additional theoretical and empirical work is required to further improve our understanding of how bureaucracies are evolving in response to implementing AI, and how bureaucratic forms shape the nature of AI implementations in turn. While AI can improve the execution of increasingly complex tasks, what is less known is how to carefully integrate these tools into the decision-making process and structure of public organizations. We are beginning to explore the actual uses of AI within public organizations and the consequences of that integration. However, more detailed cases studies, surveys, and interviews are needed to better understand how public servants and public organizations are implementing AI and the resulting consequences. Additionally, more theoretical work needs to be done to understand the general consequences to bureaucratic form as AI becomes more intimately integrated. This work must also revisit classic questions of public governance and administration, such as the relationship between principals and agents, agents' motivations, and the structure of policies as human and artificial agents work together to govern.

References

Andrejevic, M. (2017). To Preempt a Thief. *International Journal of Communication (19328036)*, *11*, 879–896. Communication Source.

Bauder, R., & Khoshgoftaar, T. (2018). A survey of medicare data processing and integration for fraud detection. *2018 IEEE International Conference on Information Reuse and Integration (IRI)*, 9–14.

Benbouzid, B. (2019). To predict and to manage. Predictive policing in the United States. *Big Data & Society*, *6*(1), 2053951719861703.

Bennett Moses, L., & Chan, J. (2018). Algorithmic prediction in policing: Assumptions, evaluation, and accountability. *Policing and Society*, *28*(7), 806–822.

Binns, R. (2019). Human Judgement in Algorithmic Loops; Individual Justice and Automated Decision-Making. *Individual Justice and Automated Decision-Making*, (September 11, 2019).

Bostrom, N. (2014). *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press.

Bovens, M., & Zouridis, S. (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, S. (2017). Big data surveillance: The case of policing. *American Sociological Review*, *82*(5), 977–1008.

Buffat, A. (2015). Street-level bureaucracy and e-government. *Public Management Review*, *17*(1), 149–161.
Bullock, J. B. (2019). Artificial Intelligence, Discretion, and Bureaucracy. *The American Review of Public Administration*, *49*(7), 751–761. doi: 10.1177/0275074019856123.
Busch, P. A., & Eikebrokk, T. R. (2019). *Digitizing Discretionary Practices in Public Service Provision: An Empirical Study of Public Service Workers' Attitudes*.
Busch, P. A., & Henriksen, H. Z. (2018). Digital discretion: A systematic literature review of ICT and street-level discretion. *Information Polity*, *23*(1), 3–28.
Centers for Medicare and Medicaid. (2017). *Medicare and Medicaid Integrity Programs Fiscal Year 2017 Annual Reports*. Centers for Medicare and Medicaid.
Centers for Medicare and Medicaid. (2019). *National Health Expenditure Data*. Centers for Medicare and Medicaid.
Chohlas-Wood, A., & Levine, E. S. (2019). A Recommendation Engine to Aid in Identifying Crime Patterns. *INFORMS Journal on Applied Analytics*, *49*(2), 154–166. doi: 10.1287/inte.2019.0985.
Danaher, J. (2016). The Threat of Algocracy: Reality, Resistance and Accommodation. *Philosophy & Technology*, *29*(3), 245–268. doi: 10.1007/s13347-015-0211-1.
Department of Health and Human Services. (2018). *Agency Annual Report for Fiscal Year 2018*. Department of Health and Human Services. https://www.hhs.gov/sites/default/files/fy-2018-hhs-agency-financial-report.pdf.
Department of Health and Human Services. (2019). *Health Care Fraud and Abuse Control Program Annual Report for Fiscal Year 2018*. Department of Health and Human Services.
Dietvorst, B. J., Simmons, J. P., & Massey, C. (2018). Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management Science*, *64*(3), 1155–1170.
Dora, P., & Sekharan, G. H. (2015). Healthcare insurance fraud detection leveraging big data analytics. *International Journal of Science and Research*, *4*(4), 2073–2076.
Dressel, J., & Farid, H. (2018). The accuracy, fairness, and limits of predicting recidivism. *Science Advances*, *4*(1), eaao5580.
Drexler, K. E. (2019). *Reframing Superintelligence: Comprehensive AI Services as General Intelligence* (No. 2019-1; FHI Technical Report). Future of Humanity Institute, University of Oxford. https://www.fhi.ox.ac.uk/wp-content/uploads/Reframing_Superintelligence_FHI-TR-2019-1.1-1.pdf.
Dunleavy, P., Margetts, H., Bastow, S., & Tinkler, J. (2006). New public management is dead – Long live digital-era governance. *Journal of Public Administration Research and Theory*, *16*(3), 467–494.
Eubanks, V. (2018). *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press.
Ferguson, A. G. (2012). Predictive Policing and Reasonable Suspicion. *Emory Law Journal*, *62*(2), 259.
Ferguson, A. G. (2016). Policing Predictive Policing. *Washington University Law Review*, *94*, 1109.
Ferguson, A. G. (2017). *The Rise of Big Data Policing: Surveillance, Race, and the Future of Law Enforcement*. NYU Press.
Fountain, J. E. (2001). *Building The Virtual State: Information Technology and Institutional Change*. Brookings Institution Press.
Frank, M. R., Autor, D., Bessen, J. E., Brynjolfsson, E., Cebrian, M., Deming, D. J., Feldman, M., Groh, M., Lobo, J., & Moro, E. (2019). Toward understanding the impact of artificial intelligence on labor. *Proceedings of the National Academy of Sciences*, *116*(14), 6531–6539.
Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, *114*, 254–280. doi: 10.1016/j.techfore.2016.08.019.
Frey, W. R., Patton, D. U., Gaskell, M. B., & McGregor, K. A. (2018). Artificial Intelligence and Inclusion: Formerly Gang-Involved Youth as Domain Experts for Analyzing Unstructured Twitter Data. *Social Science Computer Review*, *38*(1), 42–56. doi: 10.1177/0894439318788314.
Garvie, C. (2016). *The perpetual line-up: Unregulated police face recognition in America*. Georgetown Law, Center on Privacy & Technology.
Gil-Garcia, J. R., Dawes, S. S., & Pardo, T. A. (2018). Digital government and public management research: Finding the crossroads. *Public Management Review*, *20*(5), 633–646. doi: 10.1080/14719037.2017.1327181.
Gil-Garcia, J. R., Helbig, N., & Ojo, A. (2014). Being smart: Emerging technologies and innovation in the public sector. *Government Information Quarterly*, *31*, I1–I8.
Government Accountability Office. (2012). *Medicare Fraud Prevention: CMS Has Implemented a Predictive Analytics System, but Needs to Define Measures to Determine Its Effectiveness. (GAO Publication No. 13-104)*. Government Accountability Office.
Government Accountability Office. (2015). *Medicare: Potential Uses of Electronically Readable Cards for Beneficiaries and Providers. (GAO Publication No. 15-319)*. Government Accountability Office.
Government Accountability Office. (2018a). *Artificial Intelligence: Emerging Opportunities, Challenges, and Implications*. (GAO Publication No. 18-142SP). Government Accountability Office.
Government Accountability Office. (2018b). *Improper Payments: Actions and Guidance Could Help Address Issues and Inconsistencies in Estimation Processes*. (GAO Publication No. 18-377). Government Accountability Office.
Government Accountability Office. (2018c). *Medicare: Actions Needed to Better Manage Fraud Risks*. (GAO Publication No. 18-660T). Government Accountability Office.

Government Accountability Office. (2019). *Insurance Markets: Benefits and Challenges Presented by Innovative Uses of Technology*. (GAO Publication No. 19-423). Government Accountability Office.
Greer, R. A., & Bullock, J. B. (2018). Decreasing improper payments in a complex federal program. *Public Administration Review*, *78*(1), 14–23.
Hannah-Moffat, K. (2013). Actuarial sentencing: An "nsettled" proposition. *Justice Quarterly*, *30*(2), 270–296.
Hannah-Moffat, K. (2019). Algorithmic risk governance: Big data analytics, race and information activism in criminal justice debates. *Theoretical Criminology*, *23*(4), 453–470.
Hannah-Moffat, K., Maurutto, P., & Turnbull, S. (2009). Negotiated risk: Actuarial illusions and discretion in probation. *Can. JL & Soc.*, *24*, 391.
Hetling, A., Watson, S., & Horgan, M. (2012). "We Live in a Technological Era, Whether You Like It or Not": Client Perspectives and Online Welfare Applications. *Administration & Society*, *46*(5), 519–547. doi: 10.1177/0095399712465596.
Houston, S. (2015). Reducing Child Protection Error in Social Work: Towards a Holistic-Rational Perspective. *Journal of Social Work Practice*, *29*(4), 379–393. 10.1080/02650533.2015.1013526.
Hupe, P., & Buffat, A. (2014). A Public Service Gap: Capturing contexts in a comparative approach of street-level bureaucracy. *Public Management Review*, *16*(4), 548–569. doi: 10.1080/14719037.2013.854401.
Korinek, A., & Stiglitz, J. E. (2017). *Artificial intelligence and its implications for income distribution and unemployment*. National Bureau of Economic Research.
Levine, E. S., Tisch, J., Tasso, A., & Joy, M. (2017). The New York City Police Department's Domain Awareness System. *INFORMS Journal on Applied Analytics*, *47*(1), 70–84. doi: 10.1287/inte.2016.0860.
Lipsky, M. (2010). *Street-level bureaucracy*. Russell Sage Foundation.
Marks, A., Bowling, B., & Keenan, C. (2015). Automatic justice? Technology, crime and social control. *The Oxford Handbook of the Law and Regulation of Technology, OUP, Forthcoming*.
McClure, P. K. (2018). "You're fired," says the robot: The rise of automation in the workplace, technophobes, and fears of unemployment. *Social Science Computer Review*, *36*(2), 139–156.
Meijer, A., & Wessels, M. (2019). Predictive Policing: Review of Benefits and Drawbacks. *International Journal of Public Administration*, *42*(12), 1031–1039. doi: 10.1080/01900692.2019.1575664.
Norton, A. A. (2013). Predictive policing: The future of law enforcement in the Trinidad and Tobago Police Service (TTPS). *International Journal of Computer Applications*, *62*(4).
Peeters, R., & Schuilenburg, M. (2018). Machine justice: Governing security through the bureaucracy of algorithms. *Information Polity*, *23*(3), 267–280.
Peeters, R., & Widlak, A. (2018). The digital cage: Administrative exclusion through information architecture – The case of the Dutch civil registry's master data management system. *Government Information Quarterly*, *35*(2), 175–183.
Perry, W. L. (2013). *Predictive policing: The role of crime forecasting in law enforcement operations*. Rand Corporation.
Pierson, E., Simoiu, C., Overgoor, J., Corbett-Davies, S., Jenson, D., Shoemaker, A., Ramachandran, V., Barghouty, P., Phillips, C., Shroff, R., & Goel, S. (2020). A large-scale analysis of racial disparities in police stops across the United States. *Nature Human Behaviour*, *4*(7), 736–745. doi: 10.1038/s41562-020-0858-1.
Pors, A. S. (2015). Becoming digital – passages to service in the digitized bureaucracy. *Journal of Organizational Ethnography*, *4*(2), 177–192. doi: 10.1108/JOE-08-2014-0031.
Rahwan, I., Cebrian, M., Obradovich, N., Bongard, J., Bonnefon, J.-F., Breazeal, C., Crandall, J. W., Christakis, N. A., Couzin, I. D., & Jackson, M. O. (2019). Machine behaviour. *Nature*, *568*(7753), 477–486.
Russell, S., & Norvig, P. (2009). *Artificial Intelligence: A Modern Approach* (3 edition). Pearson.
Saunders, J., Hunt, P., & Hollywood, J. S. (2016). Predictions put into practice: A quasi-experimental evaluation of Chicago's predictive policing pilot. *Journal of Experimental Criminology*, *12*(3), 347–371. doi: 10.1007/s11292-016-9272-0.
Seawright, J., & Gerring, J. (2008). Case Selection Techniques in Case Study Research: A Menu of Qualitative and Quantitative Options. *Political Research Quarterly*, *61*(2), 294–308. doi: 10.1177/1065912907313077.
Selbst, A. D. (2017). Disparate impact in big data policing. *Ga. L. Rev.*, *52*, 109.
Sherman, L. W. (2013). The Rise of Evidence-Based Policing: Targeting, Testing, and Tracking. *Crime and Justice*, *42*, 377–451. doi: 10.1086/670819.
Simon, H. A. (1997). *Administrative Behavior* (4th ed.). Simon & Schuster.
Thunman, E., Ekström, M., & Bruhn, A. (2020). Dealing With Questions of Responsiveness in a Low-Discretion Context: Offers of Assistance in Standardized Public Service Encounters. *Administration & Society*, *52*(9), 1333–1361. doi: 10.1177/0095399720907807.
Tummers, L., & Bekkers, V. (2014). Policy implementation, street-level bureaucracy, and the importance of discretion. *Public Management Review*, *16*(4), 527–547.
van Eijk, G. (2017). Socioeconomic marginality in sentencing: The built-in bias in risk assessment tools and the reproduction of social inequality. *Punishment & Society*, *19*(4), 463–481.
van Eijk, G. (2020). Inclusion and exclusion through risk-based justice: Analysing combinations of risk assessment from pretrial detention to release. *The British Journal of Criminology*.

Young, M. M., Bullock, J. B., & Lecy, J. D. (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. doi: 10.1093/ppmgov/gvz014.

Završnik, A. (2019). Algorithmic justice: Algorithms and big data in criminal justice settings. *European Journal of Criminology*, 1477370819876762.

Zou, J., & Schiebinger, L. (2018). *AI can be sexist and racist – It's time to make it fair*. Nature Publishing Group.