

---

This is an electronic reprint of the original article.  
This reprint may differ from the original in pagination and typographic detail.

Pääkkönen, Juho

## Bureaucratic routines and error management in algorithmic systems

*Published in:*  
Hallinnon tutkimus

*DOI:*  
[10.37450/ht.107880](https://doi.org/10.37450/ht.107880)

Published: 18/11/2020

*Document Version*  
Publisher's PDF, also known as Version of record

*Published under the following license:*  
CC BY-NC-ND

*Please cite the original version:*  
Pääkkönen, J. (2020). Bureaucratic routines and error management in algorithmic systems. *Hallinnon tutkimus*, 40(4), 243-253. <https://doi.org/10.37450/ht.107880>

---

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

# Bureaucratic routines and error management in algorithmic systems



Juho Pääkkönen

## ABSTRACT

This article discusses how an analogy between algorithms and bureaucratic decision-making could help conceptualize error management in algorithmic systems. It argues that a view of algorithms as irreflexive bureaucratic processes is insufficient as an account of errors in complex public sector contexts, where algorithms operate jointly with other organizational work practices. To conceptualize such contexts, the article proposes that algorithms could be viewed as analogous to work routines in bureaucratic organizations. Doing so helps clarify that algorithmic irreflexivity becomes problematic when the coordination of routine work around automation fails. Thus, also the challenges of error management come to concern the wider context of organized work. This argument is illustrated using known examples from the critical literature on algorithms. Finally, drawing on recent studies in routine dynamics, the article formulates empirical research directions on error management in algorithmic systems.

Keywords: Automated decision-making; decision-making errors; irreflexivity; bureaucracy; routine dynamics

## INTRODUCTION

Recent discussions of automated decision-making have drawn on literature in organizational and administrative studies to analyze the role and shortcomings of algorithms. Rule-based algorithms in public administration have been likened to the strictly circumscribed decision processes of classical bureaucracies, which cannot deal with heterogeneous, ambiguous, or frequently shifting decision-making circumstances (Autioniemi 2020). Similarly to bureaucratic processes, algorithms have been argued to

constitute *irreflexive* mechanisms for decision-making, in that at the time of their execution, they cannot accommodate unanticipated or situation-dependent information about the processed cases (see Alkhatib & Bernstein 2019). While new techniques such as machine learning promise efficient data-driven recommendations in increasingly complex contexts, scholars have argued that even sophisticated modeling approaches can be blind to important situational nuances that bear upon the decisions (e.g. Salovaara et al. 2019). In the critical literature, the underlying worry is that the automation of decision-making through irreflexive algorithms can lead to unfair or otherwise erroneous results in complex settings – such as criminal justice (Angwin et al. 2016), child welfare protection (Eubanks 2018), or medical care (Obermeyer et al. 2019) – where decisions have traditionally relied on the discretionary evaluation of human officials.

In this article, I discuss this proposed analogy between algorithms and bureaucratic decision processes to articulate how it could help us understand the errors of automated decision-making. More specifically, I will address the following research question: *How could the view of algorithms as irreflexive processes help us conceptualize and empirically investigate decision-making errors in settings where algorithms operate jointly with other organizational work practices?* I use the term “algorithm” to encompass both the relatively simple rule-governed mechanisms used to automate tasks in highly regulated contexts such as tax administration, as well as the more sophisticated statistical approaches to data modeling that typically underlie artificial intelligence applications (cf. Autioniemi 2020). By “errors”, I mean any results of decision processes, which are unintentionally unfair, systematically biased, or inconsistent with the aims of decision-making. Examples of such adverse outcomes abound in the literature on

algorithms, but a systematic understanding of how they emerge in relation to wider contexts of organizational action has yet to be articulated. The aim of this article is to formulate an initial conceptualization for such an account.

To do so, I will argue that while the view of algorithms as irreflexive decision processes akin to classical bureaucracy provides an intuitive explanation for the errors of automation, it does not suffice to shed light on cases where algorithms interoperate with other organizational work practices. In such cases, a study of how the potential errors of automation are *managed* – that is, understood, recognized, and dealt with – as part of organizational action is crucial for understanding how they might come about. It is this process of *error management*, and in particular its challenges, that I seek to conceptualize through further discussing the analogy between bureaucracies and algorithmic systems.

To answer my research question, I propose that algorithms could be considered as analogous to work routines in bureaucratic organizations (e.g. Crozier 1964; Lipsky 1980). I intend this analogy to have both a positive and a negative side. Positively, the analogy retains that algorithms are similar to strictly circumscribed work patterns in bureaucratic organizations – *bureaucratic routines* – in that they constitute extremely formalized and irreflexive mechanisms for decision-making. However, on the negative side, algorithms bear important differences in relation to other work routines, as they have been characterized in recent organizational studies literature. Most importantly, routines have been argued to be dynamic in nature, in the sense that they can evolve to accommodate unanticipated or shifting circumstances of action (Feldman & Pentland 2003). By contrast, and similarly to bureaucratic routines, algorithms can be difficult to adapt to varying organizational practices, which can lead to inconsistent outputs and failures to identify unfair decisions.

Building on this perspective, I will argue that algorithmic irreflexivity becomes problematic in organizational decision-making when the coordination of routine performances around automation fails. Furthermore, I will propose that the literature on routine dynamics provides

a fruitful starting point for investigating these challenges (e.g. Berente et al. 2016; Glaser et al. 2021; Kremser & Schreyögg 2016). As such, the view of algorithms as bureaucratic routines helps reformulate criticisms of algorithmic irreflexivity into interesting empirical questions about work processes in technological systems. The delineation of these research questions is the primary contribution of this article.

My discussion will proceed as follows. In the next section, I will introduce the view of algorithms as irreflexive bureaucratic processes and argue that it does not suffice to explain errors in cases where algorithms operate jointly with other organizational practices. Thus, error management as an organizational process needs to be considered. To illustrate this point, I discuss more closely a case documented by Eubanks (2018), where decision-making errors result from a host of routine practices operating jointly with automated processes. I will then argue that the view of algorithms as bureaucratic routines can help us conceptualize such interrelations, and to ground an approach to empirically studying error management. Finally, I conclude with some thoughts on the relevance of understanding error management for the ethical and societal debates about automated decision-making.

### **ALGORITHMIC IRREFLEXIVITY AND THE ERRORS OF AUTOMATED DECISION-MAKING**

Critical scholars of algorithms have demonstrated case after another how the automation of decision-making can lead to adverse consequences. For instance, risk prediction systems that draw on data about past cases have been shown to exhibit biases across a number of contexts. As a recent example, Obermeyer et al. (2019) showed that a widely used algorithm for assessing patient need for medical care was racially biased, because its predictions used earlier treatment expenses as a proxy measure for the risk of future illness, which led to prioritizing wealthy clients over the poor. On the other hand, algorithms that rely on predesigned classifications or standardized rules for sorting out people have been shown to oversimplify heterogeneous real-world situations. A recent

case is the dysfunctional COVID-19 vaccine distribution algorithm at the Stanford Medical Centre – a simple rule-based decision-process, which was found to prioritize high ranking doctors over patient facing workers, because the algorithm was not designed to factor in the time that workers actually spend facing patients who carry the disease (Guo & Hao 2020).

Technical literature has variously traced the roots of such algorithmic errors to biased patterns in training data, or to poorly understood decision functions with unintended consequences (e.g. Christian 2020). Around these issues, a vast body of work has developed that analyzes notions such as fairness and accountability to map design directions for algorithms in the public sector and beyond (e.g. Diakopoulos 2016; Friedler et al. 2021). However, in addition to these studies, scholars have recently started to explore how the longstanding literature on earlier decision-making mechanisms could be applied to inform the design and evaluation of algorithmic systems. These works argue that important insights could be learned by viewing algorithms as analogous to more traditional bureaucratic decision processes (Alkhatib & Bernstein 2019; Autioniemi 2020; Alkhatib 2021; Pääkkönen et al. 2020).

The starting point of this approach is that algorithms are rigorously prespecified sets of rules for information processing, which at the time of execution cannot take into account unanticipated or situation-dependent information that might be relevant for decisions (e.g. Alkhatib & Bernstein 2019; Autioniemi 2020). This *irreflexivity* of algorithms becomes especially problematic in complex public sector contexts, where the discretionary evaluation of particular cases has traditionally formed the foundation for policy application (Alkhatib & Bernstein 2019). Decision-making tasks in the public sector are complex in the sense that abstract policy principles and regulations typically cannot be straightforwardly applied to particular cases, due to case heterogeneity and ambiguities in rule application with respect to concrete situations. Past studies of public sector decision-making have shown that this “gap” between abstract principles and their application is bridged in practice by human “street-level bureaucrats” (Lipsky 1980) – such as

social workers, police officers, or judges – who are capable of reflexive assessment of particular cases and practices such as responsive listening (Stivers 1994). The argument from irreflexivity is that, unlike human officials, algorithms must operate according to fixed procedures and prespecified classifications that might only poorly capture the nuanced aspects of people’s lives. At best, the operation of algorithms can be corrected *after decisions*, if information about their errors becomes available. Due to irreflexivity, Autioniemi (2020) has argued that rule-based artificial intelligence applications in public administration should be primarily limited to routine tasks, where decisions follow from unambiguous criteria (see also Mehr 2017). For the same reason, Alkhatib (2021) argues that machine learning applications with implications for large numbers of people should be abandoned, because their potential for adverse consequences is high.

This view of algorithms as *irreflexive decision processes* provides an intuitive explanation for the errors of automation. For instance, the racial biases of risk prediction in medical care can be understood by appeal to the fact that algorithms trained to infer illness risks solely on the basis of patients’ past medical expenses cannot take into account information about inequalities in access to care. Accommodating this further information would require retraining the algorithms – which in fact is what happened in this case, after researchers had uncovered systematic bias in the model’s predictions (Obermeyer et al. 2019). Nevertheless, the point is that capturing the nuances of real-world situations in relatively simple sets of decision rules or predictive models is difficult, and at the time of execution algorithms are not capable of reflecting on the shortcomings of their prespecified decision criteria. Consequently, Alkhatib and Bernstein (2019, 7–8) argue that considerations of irreflexivity should impel designers to focus their efforts on building mechanisms for post hoc error rectification – such as tolerably accessible channels for recourse – instead of technical fixes that seek to avoid errors before they happen.

However, although the idea of irreflexivity helps pinpoint how even sophisticated algorithms might introduce a source of error

into decision-making, the view says little of how these errors become implicated in wider contexts of organizational action. Past research has shown that algorithms rarely operate in isolation, often being used as systems for supporting rather than automating decisions (e.g. Saunders-Newton & Scott 2001; Christin 2017). Even in cases where important organizational processes are delegated to algorithms, research has shown that the implications of automation cannot be adequately understood without analyzing wider organizational circumstances (Shetakofsky 2017). Indeed, recent work on digital organizing has emphasized the importance of developing reliable practices for recognizing and correcting the errors of irreflexive automation (Salovaara et al. 2019; Asatiani et al. 2021). What these studies indicate is that in settings where decision-making involves organized action over and above automated information processing, the irreflexivity of algorithms alone is an insufficient analysis of errors. Rather, in such cases, an account of errors should be able to explain how attempts to recognize and correct the shortcomings of automation might fail, and what challenges are involved in developing mechanisms for error management.

In what follows, I will propose that viewing algorithms as analogous to bureaucratic routines offers a promising approach to conceptualizing such challenges in algorithmic systems. The next section illustrates these challenges by examining a case, where erroneous decisions were a result of diverse organizational practices operating in association with algorithms.

### **ILLUSTRATIVE CASE: AUTOMATED HOUSING ALLOCATION**

In her book *Automating Inequality* (2018), Virginia Eubanks investigates a system for automating homeless services deployed in Los Angeles in 2011. The stated motives behind automation in this case were to improve service efficiency and to ground the housing admission process on assessment of real need. Previously, homeless services in Los Angeles had relied on a complicated system of street-level agencies with long waiting queues, insufficient resources, and inconsistent practices for evaluating applicants. The new system was eventually deployed in

response to escalating numbers of people living on the street, with a focus on providing permanent housing for the chronically homeless – a group identified as having more complex social support needs and higher risk of health problems than the other unhoused. (Eubanks 2018, 91–93; see also <https://www.lahsa.org/ces/about>.)

As documented by Eubanks, the automated system builds on two algorithms, which draw on data from survey interviews with housing applicants to pair them with appropriate housing. The first algorithm is a ranking system, which calculates a risk score for the applicants based on their survey answers. Depending on the score, the applicant is allocated into one of three “triage” classes, which correspond to different levels of housing need (see Karusala et al. 2019 for extensive discussion of triage tools in homeless services). Low scores indicate no need for housing intervention, while higher scores indicate either the need for some limited term supportive services and rental subsidies, or permanent housing. The risk scores are then used as input for a second algorithm, a matching procedure, which pairs high-scoring applicants with available free apartments that best suit their needs in terms of prespecified eligibility criteria. (Eubanks 2018, 93–95.)

The system aspires for extensive standardization of the housing admission process, on the basis of prespecified criteria for evaluating the applicants’ situation. These criteria are operationalized by the survey questionnaire, which includes questions about, for example, the applicants’ use of emergency health care services during the past months, and the frequency of activities such as drug abuse or violent behavior (Eubanks 2018, 93–94). Each of these items corresponds to a pre-assigned score in automated evaluation, which produces the risk score for the applicant.

However, Eubanks’s investigation revealed that housing need evaluation could not be as straightforwardly standardized as presupposed in the system’s design. For instance, the various agencies responsible for organizing the actual surveys had different practices for interviewing applicants. Some interviews were administered under formal circumstances, while others were conducted in overnight shelters, where

the applicants already knew the staff and were more willing to disclose sensitive information (Eubanks 2018, 95–103, 105–107). As a result, the same applicant could receive low scores in some agencies and high scores in others. In addition to these inconsistent data collection practices, Eubanks (2018, 103–109) describes how the interpretation of risk scores varied across city areas. In some areas, a high score more or less meant that the applicant was handed keys to an apartment. In others, high scores were interpreted as implying that the applicant is incapable of maintaining an independent life, and thus cannot be granted housing. (Eubanks 2018, 107.) The prevalent practice in certain areas was to grant high-scoring applicants vouchers for renting an apartment on the private real estate market. However, as many house owners were reluctant to rent their apartments to homeless applicants, vouchers were no guarantee that applicants could actually secure housing, despite scoring high in the system. (Eubanks 2018, 108.)

These observations are corroborated by the study of data collection practices in Los Angeles homeless services by Karusala et al. (2019). The authors interviewed a number of case workers and found them to be acutely aware of the “subjectivities behind the triage tools” that deal with sensitive information about people in vulnerable positions (Karusala et al. 2019, 8). The case workers noted that the accuracy of survey scores is “actually dependent on building rapport” between the interviewer and the interviewee, and that this trust-building process can only happen gradually over time (Karusala et al. 2019, 8). Indeed, as Eubanks (2018, 94) notes, the data gathered in the system is shared with 168 different organizations, including police and government authorities. This makes disclosing scoring-relevant sensitive information difficult for some applicants, who might worry about potential complications with authorities, and highlights the importance of trust for accurate data collection (Eubanks 2018, 121). Thus, the risk scores were sensitive to variance in data collection practices, which was also reflected in the case workers’ perceptions of scoring validity. Karusala et al. (2019, 16) describe how some workers took the scores as ground-truth evidence of housing need, while

others questioned their validity on the basis of intimate knowledge of the applicants’ situations. Overcoming such tensions involved persuasive work and negotiations over case details, to determine how difficult cases should be decided, and how the validity of the automated scorings should be assessed (Karusala et al. 2019, 17–18).

These examples show that the correctness of housing admission decisions in this case depended on diverse work practices and decision-making routines in addition to the automated processes for scoring and matching applicants. The automated scorings were based on prespecified criteria for determining housing need. These criteria can of course themselves be criticized, and Eubanks (2018) in fact argues that the system’s focus on the chronically homeless complicates access to housing for large groups of unhoused people. Thus, the automated system can be argued to be irreflexive in the sense that its criteria fail to accommodate cases where there would be real need for housing. This is an example of case complexity in public services complicating the design of algorithmic decision-making systems. However, even if the prespecified scoring criteria were accepted as appropriately reflecting housing need, the system still produced inconsistent results and negative admission decisions for applicants who by its own criteria should have been eligible for support. These decision-making failings cannot be accounted for solely by appeal to the irreflexivity of the automated processes. Rather, the errors resulted from associations between varying organizational work practices and irreflexive automated processes, which were jointly responsible for the decisions.

## ALGORITHMS AS BUREAUCRATIC ROUTINES

Literature on organizational decision-making has long emphasized the importance of routinized work. Routines are commonly defined as repetitive and recognizable patterns of interdependent actions carried out by multiple actors (Feldman & Pentland 2003). Such patterns of action have been discussed as the standard mechanism for managing work in organizations. Routines have been argued to improve the stability, accountability

and efficiency of decision-making (Feldman & Pentland 2003, 94). However, they can also introduce simplifications into decision processes and be an important source of bias (Lipsky 1980). At worst, overly strict routinization has been argued to lead to mindless adherence to rules (Ashforth & Fried 1988) and inflexible practices which prevent bureaucratic organizations from recognizing and correcting dysfunctional procedures (Crozier 1964; Weiss & Ilgen 1985). For this reason, Lipsky (1980, 84) argued that an analysis of routine work is crucial for understanding the shortcomings of public sector decision-making, where the “policies that result from routine treatment are often biased in ways unintended by the agencies whose policies are being implemented”.

I will next argue that an analogy between strict routinization and algorithmic information processing provides a useful starting point for analyzing decision-making errors in algorithmic systems. On the positive side of this analogy, both algorithms and routines constitute mechanisms for organizing work tasks into repetitive and recognizable patterns. Ideally, both algorithms and routines improve the efficiency of decision-making and facilitate impartiality, through minimizing the role of worker discretion in task completion (see Bovens & Zouridis 2002). However, as Glaser et al. (2021) argue, algorithms also formalize task sequences in a manner that is much more rigid than many earlier mechanisms for managing work, such as written task instructions. As such, algorithms share particular likeness with routinized work in bureaucratic organizations, which seek to control worker behavior through extensively standardized rules and role descriptions (see Kellogg et al. 2020). Accordingly, both algorithms and bureaucratic routines have been argued to constitute “mindless” processes (see Dreyfus & Dreyfus 1984 for discussion of the mindlessness metaphor in artificial intelligence research), which can lead to errors in complex circumstances.

It is in this sense that algorithms – conceived as analogous to strict bureaucratic routines – can be argued to be importantly different from other kinds of routinized work. Although the traditional understanding of routines conceptualized them as a source of stability

and organizational inertia, more recent studies have emphasized that routines are dynamically evolving patterns of action which necessarily involve a *performative* aspect (D’Adderio 2008). Routines must be performed through “specific actions, by specific people, at specific times and places” (Feldman & Pentland 2003, 94), and have been found to change in interaction with each other and in response to shifting circumstances of action (Kremser & Schreyögg 2016; Berente et al. 2016). This means that the performance of routines can involve considerable variety, which over time can also induce changes in the abstract understanding or description of the task sequence itself (Feldman & Pentland 2003, 112–114). In other words, routinized work crucially involves an element of *discretionary assessment*, which makes routines to a certain extent capable of adapting to idiosyncratic circumstances not captured in their abstract description. This performative aspect enables routines to simultaneously improve work efficiency, while allowing low-level workers to creatively manage exceptional situations and contingencies in task execution (Berente et al. 2016; Crozier 1964; Perrow 1967).

By contrast – as the criticism from irreflexivity has it – algorithms cannot at the time of execution adapt to unexpected circumstances or changes in surrounding work practices (Alkhatib & Bernstein 2019). As Autioniemi (2020) argues, this lack of situational adaptivity might not be problematic when algorithms are used to automate simple and rigorously specified task sequences. However, as we saw in the previous section, in public sector decision-making algorithms are often deployed in conjunction with other routine processes, which might vary in performance and be subject to change over time. Consequently, the automated processes that draw on varying routine performances will yield inconsistent outputs, which moreover are used in conflicting ways in other parts of the system.

The case of homeless service automation illustrates these difficulties well. In this case, the standardized survey can be thought of as a routine process, where multiple actors perform prespecified tasks sequences to produce data for automated risk scoring. The standardized questionnaire embodies an abstract prescription

for how data collection should be accomplished, which however was variously performed in different agencies. Due to this variance, the automated scorings could only superficially be treated as comparable to each other, as in reality they more closely reflected differences across routine performances. Furthermore, the outputs of the automated system were used as grounds for decision-making downstream, in housing admission routines which also differed from each other at crucial points. Importantly, while the data collection routine could have been rigidly enough standardized, the system's overall decisions would still have remained inconsistent and in some cases contradictory to the prespecified triage classification, given that scorings were differently interpreted and deployed across city areas. The underlying problem is that the design of the automated system could not take into account variance in the routine performances associated with its operation.

Thus, by explicitly considering the similarities and differences between algorithms and routinized work, we can see that in organizational contexts algorithmic irreflexivity is not only problematic because automation might simplify nuanced real-world situations. Rather, the issue of irreflexivity also comes to concern the challenge of coordinating varying routine performances in relation to the automated processes. To understand how erroneous decisions come about, we need to investigate the difficulties involved in achieving this coordination. Algorithmic irreflexivity certainly is a source of error in organizational contexts, but an account of how these errors become implicated in decision-making should also shed light on how attempts to deal with them fail. This is what the literature on routine dynamics can help conceptualize and study in a systematic fashion.

Building on the performative view of routines, Kremser and Schreyögg (2016) argue that organizations tend to develop clusters of interrelated but semi-autonomous routines, which adapt over time to each other and to the specific circumstances of their respective tasks. The advantage of developing such "well oiled" systems of routines is that they enable efficient coordination of even very complicated tasks.

Within a cluster, single routines are autonomous in the sense that their performers need not know how work elsewhere is carried out, as long as they are able to adapt their own performances in relation to the circumstances of their focal domain. (Kremser & Schreyögg 2016, 700–702.) However, this also means that information exchange between routines becomes limited, and clusters become incapable of adapting to changes outside their particular focus (Kremser & Schreyögg 2016, 715). This limited possibility of coordinating interrelated but autonomously operating routines can also become a hindrance for error management in algorithmic systems.

Consider again the case of homeless service automation, where such locking of routines seems to be at play. As Karusala et al. (2019, 17) describe, case workers in different positions experienced difficulties in explaining their understandings of the scorings to each other. For instance, the data collectors' intimate knowledge about applicant situations could not be easily transferred to workers in housing admission. However, addressing inconsistencies would require that the divergent routine performances be coordinated with respect to each other. To some extent, this can happen through street-level practices such as case conferences, where workers negotiate difficult-to-decide cases and the appropriate interpretation of risk scores (Karusala et al. 2019, 3). However, as Eubanks (2018, 107–112) argues, more thorough efforts at alignment – for instance across city areas – are likely to require top-down intervention and extensive resource investments. Given that the diverging routines at play might be adapted to handle quite specific circumstances, achieving effective coordination might be difficult without also bringing about unexpected consequences. Indeed, a central characteristic of established routine clusters is that they resist change (Kremser & Schreyögg 2016, 715–716), and this can also complicate error rectification in algorithmic systems.

A related challenge due to routinized work is that identifying erroneous decisions might be difficult, because it might not be clear what constitutes result correctness in the first place. As the studies of Eubanks (2018) and Karusala et al. (2019) show, the criteria for decision-making correctness can vary across routines. This is



again evidenced by the diverging status ascribed to the triage scorings. Some data collectors challenged the scorings on the basis of intimate familiarity with applicant situations, while others regarded them as ground truth (Karusala et al. 2019, 16). Arguably, this is because workers in other positions could not “see” into the contingencies involved in case evaluation. From the perspective of error management, the problem is to decide how the correctness of decisions should be assessed, given such conflicting street-level judgments of the system’s functioning. Potential reconfigurations of algorithms must resolve tensions between such conflicting approaches, which make the objective measurement of errors difficult.

### **ERROR MANAGEMENT IN ALGORITHMIC SYSTEMS: DIRECTIONS FOR EMPIRICAL RESEARCH**

In the previous sections I have argued that the view of algorithms as analogous to bureaucratic routines could help us conceptualize the challenges of error management in algorithmic systems. The primary import of this conceptualization is that it offers promising starting points for empirical research on error management. One interesting issue concerns how algorithms are used to automate action in organizations, and the implications of this for error management. Glaser et al. (2021) have argued that algorithms can relate to other work routines through 1) automating parts of their action sequences, 2) establishing entirely novel sequences of action that other routines may or may not be associated with, or 3) through influencing how routines relate to each other. In each case, irreflexivity might have different implications for how decision-making errors can be detected and addressed. This calls for empirical work on the different uses of algorithms, with potential for comparisons between error management challenges and strategies across the uses. For instance, are there systematic differences in how decision-making correctness is understood and measured across the various uses of algorithms? Are errors handled as part of routine operations, or through discretionary adaptation? Previous studies have argued that automation might shift the

location of discretionary work in organizations (Pääkkönen et al. 2020). The focus on work routines could help investigate these changes in a systematic fashion.

A related issue concerns how performative variance in routines becomes problematic for irreflexive algorithms. Berente et al. (2016) have argued that dynamically adapting routines provide an important source of stability in organizations, in that they can adjust to accommodate for the potential dysfunctions of technological reforms. Somewhat paradoxically, this stability is not achieved by rigidly aligning routines to match information system design, but rather by allowing routines to develop their own performative adaptations (Berente et al. 2016, 564–567). However, in the case of homeless service automation we saw that too little standardization can lead to errors with irreflexive algorithms. This points to an interesting question: what determines when routine performances around algorithms should be rigidly coordinated, versus allowed to vary? The sociological theory of technological accidents (Perrow 1999) predicts that when technological processes are fairly simple but tightly connected to each other – in the sense that human intervention in their execution is limited – work around those processes should be highly standardized. A promising avenue for research would be to investigate whether this holds also for decision-making automation in public sector organizations.

The routine dynamics literature also provides further hypotheses for investigating the conditions under which irreflexivity can become challenging. Kremser and Schreyögg (2016, 718) note that the extent to which routines tend to form stable clusters depends on their relative centrality to the organizations’ core tasks, as well as on the age of the routines in question. Thus, a reasonable expectation would be that in such locations of tight clustering, the error potential of automation is increased, given that established clusters might not adapt well to top-down interventions and changes. On the other hand, recognizing the errors of algorithms might be easier in such situations, if the established clusters involve highly aligned criteria for decision correctness. Investigation of these issues could be combined with the

notion that algorithms can relate differently to extant routines, to form a typology of error management challenges. For instance, is replacing parts of individual routines easier than establishing novel automated sequences in tightly adapted routine clusters, and what are the associated challenges?

Finally, the routine dynamics literature provides methodological recommendations for how to study work patterns around algorithms. Kremser and Schreyögg (2016, 716–717) argue that studies of evolving routines should preferably be longitudinal, and focus on tracing how work patterns react to changes in the technological environment. According to them, routine patterns can be distinguished from each other by investigating information flows between action sequences. Two sequences of action can be analyzed as separate routines, if there is no “exchange and processing of real-time information” between the actors who perform them (Kremser & Schreyögg 2016, 716). A study of the interrelations between routines involves a reconstruction of routine clusters through investigating such information flows, and examining which routines use outputs from the others. This same approach could be applied to study work patterns that relate to automated processes, and how these evolve over time in association with algorithmic information processing.

## CONCLUSION

In this article, I have argued that the analogy between algorithms and bureaucratic routines provides a useful starting point for conceptualizing error management in algorithmic systems. Thoroughly demonstrating the usefulness of this view would require empirical studies on algorithms in organizational decision-making. Scholars of routine dynamics have begun to investigate these issues (Glaser et al. 2021), but the field’s focus on algorithmic systems is only now emerging. In particular, the focus has not been on error management in public sector automation. On the other hand, work practices that support reliable decision-making in technological systems have been discussed in

the contexts of business and industry (Salovaara et al. 2019; Shestakofsky 2017), and in the extensive literature in Computer Supported Collaborative Work. Error management in public sector automation merits similar thorough investigation, and in this respect the account I propose signposts promising directions for empirical research.

In this regard, however, it should also be recognized that there are important ethical issues related to automated decision-making that my discussion has glossed over. For instance, as Mittelstadt et al. (2016, 12) note, the problem of how responsibility should be distributed in networks of human and algorithmic actors is a pressing one for automated decision-making. The worry is that the delegation of decision-making to algorithms can lead to the “de-responsibilisation of human actors”, and an associated tendency to assume that the results of sophisticated computational models are correct (Mittelstadt et al. 2016, 12–13). Recognizing that the debate on automated decision-making spans such ethical issues is doubly important for my argument, as these issues underscore the societal relevance of the topic of error management. Algorithmic technologies are becoming a standard tool in attempts to reorganize and streamline organizational processes, both in private industries and public services (see Kellogg et al. 2020). Previous literature on public sector decision-making and e-governance has argued that the increasing automation of public services and administration can compromise decision-making legitimacy and lead to the detachment of decision processes from the heterogeneous particularities of real-life situations (Bovens & Zouridis 2002; Jansson & Erlingsson 2014; Smith et al. 2010). With the current hype about artificial intelligence in the public sector (e.g. Wirtz et al. 2019), issues such as accountability and the allocation of responsibilities in decision-making are only gaining traction. Systematic research on how algorithmic errors depend on the wider context of organized action might provide empirical grounds for informing these debates.

”This work was supported by the Kone Foundation project *Algorithmic Systems, Power and Interaction*. I would like to thank Airi Lampinen, Matti Nelimarkka and Jesse Haapoja, as well as two anonymous referees, for their instructive comments on the earlier version of this paper.”

## REFERENCES

- Pääkkönen, Juho, Nelimarkka, Matti, Haapoja, Jesse & Lampinen, Airi (2020). Bureaucracy as a lens for analyzing and designing algorithmic systems. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, Article 651, 1–14. <https://doi.org/10.1145/3313831.3376780>
- Alkhatib, Ali. (2021). To live in their utopia: Why algorithmic systems create absurd outcomes. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–14. <https://doi.org/10.1145/3411764.3445740>
- Alkhatib, Ali. & Bernstein, Michael. (2019). Street-level algorithms: A theory at the gaps between the policy and decisions. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, Article 530, 1–13. <https://doi.org/10.1145/3290605.3300760>
- Angwin, Julia, Larson, Jeff, Mattu, Surya & Kirchner, Lauren (2016). Machine bias. *ProPublica*. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- Asatiani, Aleksandre, Malo, Pekka, Nagbøl, Per Rådberg, Penttinen, Esko, Rinta-Kahila, Tapani & Salovaara, Antti (2021). Sociotechnical envelopment of artificial intelligence: An approach to organizational deployment of inscrutable artificial intelligence systems. *Journal of the Association for Information Systems*, 22(2), 325–352. <https://doi.org/10.17705/1jais.00664>
- Ashforth, Blake & Fried, Yitzhak (1988). The mindlessness of organizational behaviors. *Human Relations*, 41(4), 305–329. <https://doi.org/10.1177/001872678804100403>
- Autioniemi, Jari. (2020). Tekoälyn yhteiskehittäminen julkisella sektorilla. *Hallinnon Tutkimus*, 39(1), 5–20. <https://doi.org/10.37450/ht.98075>
- Berente, Nicholas, Lyytinen, Kalle, Yoo, Youngjin & King, John Leslie (2016). Routines as shock absorbers during organizational transformation: Integration, control, and NASA’s enterprise information system. *Organization Science*, 27(3), 551–572. <https://doi.org/10.1287/orsc.2016.1046>
- Bovens, Mark & Zouridis, Stavros (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. <https://doi.org/10.1111/0033-3352.00168>
- Christian, Brian (2020). *The Alignment Problem: Machine Learning and Human Values*. New York: W. W. Norton & Company.
- Christin, Angèle (2017). Algorithms in practice: Comparing web journalism and criminal justice. *Big Data & Society*, 4(2), 1–14. <https://doi.org/10.1177/2053951717718855>
- Crozier, Michel (1964). *The Bureaucratic Phenomenon*. Chicago: University of Chicago Press.
- D’Adderio, Luciana (2008). The performativity of routines: Theorising the influence of artefacts and distributed agencies on routine dynamics. *Research Policy*, 37(5), 769–789. <https://doi.org/10.1016/j.respol.2007.12.012>
- Diakopoulos, Nicholas (2016). Accountability in algorithmic decision making. *Communications of the ACM* 59(2), 56–62. <http://doi.acm.org/10.1145/2844110>
- Dreyfus, Hubert & Dreyfus, Stuart (1984). Mindless machines. *Sciences*, 24(6), 18–22. <https://doi.org/10.1002/j.2326-1951.1984.tb02749.x>
- Eubanks, Virginia (2018). *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. New York: St. Martin’s Press.
- Feldman, Martha & Pentland, Brian (2003). Reconceptualizing organizational routines as a source of flexibility and change. *Administrative Science Quarterly*, 48(1), 94–118. <https://doi.org/10.2307/3556620>
- Friedler, Sorelle, Scheidegger, Carlos & Venkatasubramanian, Suresh (2020). The (im) possibility of fairness: Different value systems require different mechanisms for fair decision making. *Communications of the ACM*, 64(4), 136–143. <https://doi.org/10.1145/3433949>
- Glaser, Vern, Valadao, Rodrigo & Hannigan, Timothy (2021). Algorithms and routine dynamics. In Feldman M. S., Pentland B. T., D’Adderio L., Dittrich K., Rerup C., & Seidl D. (Eds.). *Handbook of Routine Dynamics*. Cambridge University Press.
- Guo, Eileen & Hao, Karen (2020). This is the Stanford vaccine algorithm that left out frontline doctors. *MIT Technology Review*. <https://www.technologyreview.com/2020/12/21/1015303/stanford-vaccine-algorithm/>
- Jansson, Gabriella & Erlingsson, Gissur (2014). More e-government, less street-level bureau-

- cracy? On legitimacy and the human side of public administration. *Journal of Information Technology & Politics*, 11(3), 291–308. <https://doi.org/10.1080/19331681.2014.908155>
- Karusala, Naveena, Wilson, Jennifer, Vayanos, Phebe & Rice, Eric (2019). Street-level realities of data practices in homeless services provision. *Proceedings of the ACM on Human-Computer Interaction*, Vol 3, CSCW, Article 184, 1–23. <http://doi.acm.org/10.1145/3359286>
- Kellogg, Katherine, Valentine, Melissa & Christin, Angèle (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals*, 14(1), 366–410. <https://doi.org/10.5465/annals.2018.0174>
- Kremsler, Waldemar & Schreyögg, Georg (2016). The dynamics of interrelated routines: Introducing the cluster level. *Organization Science*, 27(3), 698–721. <https://doi.org/10.1287/orsc.2015.1042>
- Lipsky, Michael (1980). *Street-Level Bureaucracy: The Dilemmas of the Individual in the Public Services*. New York: Russell Sage Foundation.
- Mehr, Hila (2017). *Artificial Intelligence for Citizen Services and Government*. Harvard Kennedy School, Ash Center for Democratic Governance and Innovation. <https://ash.harvard.edu/publications/artificial-intelligence-citizen-services-and-government>
- Mittelstadt, Brent, Allo, Patrick, Taddeo, Maria-Rosaria, Wachter, Sandra & Floridi, Luciano (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2), 1–21. <https://doi.org/10.1177/2053951716679679>
- Obermeyer, Ziad, Powers, Brian, Vogeli, Christine & Mullainathan, Sendhil (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447–453. <https://doi.org/10.1126/science.aax2342>
- Perrow, Charles (1967). A framework for the comparative analysis of organizations. *American Sociological Review*, 32(2), 194–208. <https://doi.org/10.2307/2091811>
- Perrow, Charles (1999). *Normal Accidents: Living With High-Risk Technologies*. Princeton: Princeton University Press.
- Salovaara, Antti, Lyytinen, Kalle & Penttinen, Esko (2019). High reliability in digital organizing: Mindlessness, the frame problem, and digital operations. *MIS Quarterly*, 43(2), 555–578. <https://doi.org/10.25300/MISQ/2019/14577>
- Saunders-Newton, Desmond & Scott, Harold (2001). “But the computer said!”: Credible uses of computational modeling in public sector decision making. *Social Science Computer Review*, 19(1), 47–65. <https://doi.org/10.1177/089443930101900105>
- Shestakofsky, Benjamin (2017). Working algorithms: Software automation and the future of work. *Work and Occupations*, 44(4), 376–423. <https://doi.org/10.1177/0730888417726119>
- Smith, Matthew, Noorman, Merel & Martin, Aaron (2010). Automating the public sector and organizing accountabilities. *Communications of the Association for Information Systems*, 26(1), 1–16. <https://doi.org/10.17705/1CAIS.02601>
- Stivers, Camilla (1994). The listening bureaucrat: Responsiveness in public administration. *Public Administration Review*, 54(4), 364–369. <https://doi.org/10.2307/977384>
- Weiss, Howard & Ilgen, Daniel (1985). Routinized behavior in organizations. *Journal of Behavioral Economics*, 14(1), 57–67. [https://doi.org/10.1016/0090-5720\(85\)90005-1](https://doi.org/10.1016/0090-5720(85)90005-1)
- Wirtz, Bernd, Weyerer, Jan & Geyer, Carolin (2019). Artificial Intelligence and the public sector —Applications and challenges. *International Journal of Public Administration*, 42(7), 596–615. <https://doi.org/10.1080/01900692.2018.1498103>