

# AI and bureaucratic discretion

Kate Vredenburg

To cite this article: Kate Vredenburg (29 Sep 2023): AI and bureaucratic discretion, *Inquiry*, DOI: [10.1080/0020174X.2023.2261468](https://doi.org/10.1080/0020174X.2023.2261468)

To link to this article: <https://doi.org/10.1080/0020174X.2023.2261468>



© 2023 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group



Published online: 29 Sep 2023.



Submit your article to this journal [↗](#)



Article views: 110



View related articles [↗](#)



View Crossmark data [↗](#)

# AI and bureaucratic discretion

Kate Vredenburg

Department of Philosophy, Logic, and Scientific Methods, The London School of Economics and Political Science, London, UK

## ABSTRACT



Algorithmic decision-making has the potential to radically reshape policy-making and policy implementation. Many of the moral examinations of AI in government take AI to be a neutral epistemic tool or the value-driven analogue of a policymaker. In this paper, I argue that AI systems in public administration are often better analogised to a street-level bureaucrat. Doing so opens up a host of questions about the moral dispositions of such AI systems. I argue that AI systems in public administration often act as indifferent bureaucrats, and that this can introduce a problematic homogeneity in the moral dispositions in administrative agencies.

**ARTICLE HISTORY** Received 27 April 2023; Accepted 11 September 2023

**KEYWORDS** AI; discretion; organizations

## 1. Introduction

Virginia Eubanks (2018, Chapter 4) tells the story of Pat Gordan, an intake screener in the Department of Human Services in Allegheny County, Pennsylvania. The Department deploys a risk assessment tool to aid decisions about when to investigate a report of child mistreatment, called the Allegheny Family Screening Tool (AFST). Eubanks, with a keen eye for the ridiculous, recounts how Pat has become adept at predicting the risk score that the algorithm will output, without understanding why it has output what it has. Furthermore, Pat sometimes finds herself disagreeing with the algorithm's output. Eubanks recounts a story in which she and Pat both rate two cases as low risk, but the AFST rates one case as high risk, as Pat predicted it would. But, because Pat doesn't understand the

**CONTACT** Kate Vredenburg  K.Vredenburg@lse.ac.uk  Department of Philosophy, Logic, and Scientific Methods, the London School of Economics and Political Science, London, UK

© 2023 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group  
This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

reasons behind the risk score, she must choose between ignoring the AFST's risk score or overriding her own judgement, on very little evidence about which one is likely to be correct in this particular case.

AI has already changed how many are governed. It is used by governments to identify those entering and leaving their territory, to detect cases of financial fraud, or to inform decisions about release from prison. And it has the potential to accelerate the trend of so-called algorithmic governance, or the use of information technologies in public administrations and in the provision of governmental services (Garson 2006; Henman 2010). But, there are many reasons to be concerned about algorithmic governance. The case of Pat Gordon contains a number of these reasons, such as opacity, or how AI systems can encode bias and advance unjust political agendas.<sup>1</sup> Such concerns have received attention across philosophy, computer science, and the social sciences, galvanising debates in the philosophy of AI on algorithmic bias and justice (Johnson 2020; Zimmermann and Lee-Stronach 2022) or opacity (Creel 2020; Sullivan 2022; Vredenburg 2022).

But, in this paper, I want to focus on an overlooked moral concern about the use of AI in administrative agencies. In Eubanks' story, Pat doesn't treat the algorithmic system like a policymaker who has come in to implement a new policy or to exercise top-down control over her actions. Instead, she treats the system like a predictable but frustrating peer, whose opaque, firmly-held judgments she has to either ignore or accept. And, Pat is not mistakenly anthropomorphising the system. Instead, her treatment of it is apt, and insightful. Algorithmic decision-making systems sometimes act as street-level bureaucrats like Pat, rather than as policymakers. Accordingly, they raise moral issues that are unique to street-level bureaucrats. But, because AI systems are usually conceptualised as neutral epistemic tools, or analogised to policymakers, these moral issues have been overlooked.

This paper makes three contributions to the literature on the ethics of algorithmic decision tools. The first contribution is to argue that discretion is morally valuable, not merely epistemically valuable (Sections 3 and 4). The second contribution is to argue that many algorithms are better analogised to street-level bureaucrats than they are to policymakers or top-level bureaucrats, and that they act as indifferent bureaucrats (Sections 5 and 6). The third is to argue that an over reliance on algorithms,

---

<sup>1</sup>Eubanks (2018, 94) that the algorithm disproportionately targets poor individuals for further investigation, despite Allegheny County's laudable attempts to de-bias the algorithm.

especially opaque algorithms, leads to a problematic imbalance in the distribution of moral dispositions in organisations (Section 7).

## 2. Street-level bureaucrats and the problem of discretion

In law, economics, information science, and public administration, there is a robust debate about whether algorithmic governance is superior to human governance because it is more accurate and uniform (e.g. Sunstein 2022). One mechanism that can increase the accuracy and uniformity of governmental decision-making is the reduction of decision-maker discretion (Busch and Henriksen 2018). In particular, efforts to reduce discretion in government have targeted *street-level bureaucrats*, or front line workers in public administrations who interact with the public in the course of service provision. To understand the moral nuances of discretion in governmental decision-making, we need to examine street-level bureaucrats, and how their work could be impacted by the introduction of algorithmic governance systems.

At least since Lipsky's (1980) groundbreaking book *Street Level Bureaucracy: The Dilemmas of the Individual in Public Service*, mainstream research in public administration has recognised that street-level bureaucrats have a distinct role in policy implementation. Policymakers, aided by bureaucratic experts, pass statutes, and top-level bureaucrats create rules and regulation to enforce these statutes. Street-level bureaucrats, by contrast, are the decision-makers on the ground who, in the words of Maynard-Moody and Musheno (2012, S19), 'occupy an organisational space in which rules and other abstractions confront specific circumstances and specific people, both citizen-clients and fellow workers'. They are the teachers, police officers, and social workers who are responsible for making decisions about individuals in light of bureaucratic rules and regulations.

The property which is often taken to mark out street-level bureaucrats is the *discretion* they have in decisions about how to allocate resources and to whom. In public administration research, discretion is conceptualised differently across different paradigms, varying according to the background theory of the organisation at play (Maynard-Moody and Musheno 2000; Zacka 2017, Chapter 1). Common to these accounts of discretion, however, is a focus on 'the freedom that street-level bureaucrats have in determining the sort, quantity and quality of sanctions, and rewards during policy implementation' (Tummers and Bekkers 2014, 529, discussing Lipsky's (1980) account of bureaucratic discretion). The definition of discretion that I will utilise in this paper adds two further components

to this account. Discretion, first of all, is a set of powers and prerogatives grounded in the public official's role in the administration. These powers and prerogatives concern both the scope of her decision-making power – which decisions about sanctions and rewards are in her remit? – and how constrained that decision-making power is – is she accountable to other actors, or subject to oversight?. Secondly, discretion involves the *de facto* power to choose permissible and impermissible sanctions and rewards, from both a moral and an institutional perspective. Discretion can be used to contravene institutional rules and policies, or to choose a morally unjustifiable sanction that may or may not be permitted by the rules and policies.

It is often argued that bureaucrats must exercise discretion to make decisions in their day-to-day jobs. That is because their directives are often ambiguous, rules and procedures conflict, and they do not have the resources to satisfy all justified client demands (Zacka 2017, Chapter 1). There are a number of rules and regulations that apply to any particular case, and not all of the rules and regulations can be applied at the same time, as they often conflict (Maynard-Moody and Musheno 2000). Furthermore, following a rule often requires some interpretation by the agent. Bureaucratic rules, for example, often state what actions ought to be taken, but not how. For example, a police officer's role involves interviewing witnesses, but their manner will differ depending on how cooperative the witness is (Zacka 2017, Chapter 1). Furthermore, the need for exercises of discretion also arises because of the organisational constraints in which street-level bureaucrats make decisions: the information that they have to make a decision under-determines what to do, and they often lack the time or resources to gather more information (Tummers and Bekkers 2014). However, as Zacka (2017, 34) stresses, 'discretion is not "doing as one pleases" within the bounds of the law'. Bureaucratic decision-making is constrained by sanctions and rules, as well as standards of reasonableness that inform managerial or peer criticism.

Recently, political philosophy has also begun to consider moral questions raised by street-level bureaucrats' ability to exercise discretion. Bureaucratic discretion is often seen as a threat to fairness, legitimacy, equality, and the effectiveness of policy. Discretion, the worry goes, is an opportunity for bureaucrats to pursue their own private goals, creating problems of corruption (Kolodny 2023) or mistakes. This pursuit of private ends raises concerns of legitimacy, as bureaucrats fail to adhere to the letter or the spirit of the democratic will (Cordelli 2020; Lipsky 1980).

Fairness concerns arise because bureaucratic discretion may be exercised in a way that is partial or discriminatory, leading bureaucrats to treat relevantly alike cases differently and thereby violating norms of procedural fairness. Discretion can also undermine the effectiveness of policy: if bureaucrats pursue their own private ends, resources will not make their way to the intended recipients. Finally, the inappropriate exercise of discretion creates objectionable relations of unequal standing between bureaucrats and the clients they serve, because bureaucrats have 'untempered', asymmetric power over clients (Kolodny 2023, 91). Bureaucratic rules and procedures are important mechanisms to reduce such problematic discretion (Mashaw 1983).

For those who are concerned about street-level exercises of discretion, AI systems hold great promise to improve bureaucratic decision-making. That is not only because these systems are often more accurate than human decision-makers. In addition, their predictions are often less noisy, or display less 'unwanted variability in judgments' (Sunstein 2022, 1178). Noise can be most easily illustrated in contexts where a fixed fact is measured at different points in time: for example, my doctor's judgments of my height would be noisy if I stay the same height, but she measures me to be five foot seven on Monday, five foot eight on Wednesday, and five foot six on Friday. Some, such as Sunstein (2022), have argued that algorithms are less noisy than human decision-makers. Kleinberg et al. (2018), for example, trained an algorithm on judicial decisions about whether to release an incarcerated person on bail. They found that the algorithm was less noisy than judges themselves, hypothesising that judges tend to respond to non-predictive or non-explanatory features in the data as if they were meaningful signals of the true state of the world. Scientists have also argued that using methods from machine learning for scientific model building can reduce noise (Yarkoni and Westfall 2017).

However, many balk at the greater integration of AI into the administrative state. One serious ground for concern is the pervasive and pressing problem of algorithmic bias, or the encoding of racial biases into algorithms and attendant discriminatory decisions that perpetuate race-based subordination (Barocas and Selbst 2016; Noble 2018). Algorithmic bias, however, is not the focus of this paper. Instead, I will examine concerns about AI's reduction of discretion in street-level bureaucratic decision-making. The empirical evidence on whether AI actually reduces bureaucratic discretion is both scant and mixed (de Boer and Raaphorst 2023; see also Buffat 2015 on information technologies generally).

So, rather than assume that algorithmic governance will reduce bureaucratic discretion, I will examine moral concerns about its potential to do so.

When is it morally concerning that AI reduces street-level bureaucratic discretion? To answer this question, we need to know why discretion is morally valuable, if it is. I will consider two accounts of the value of discretion in the next two sections. The first account of the value of discretion I will consider is epistemic (Section 3). I will argue that this epistemic account has serious flaws, and that we ought to move to a moral account of the value of bureaucratic discretion (Section 4). Once we do, we can see that AI poses a threat to street-level bureaucratic decision-making because it reduces the range of moral dispositions exercised by street-level bureaucrats (Section 7).

### 3. The epistemic value of bureaucratic discretion

In this section, I will consider an epistemic defence of discretionary decision-making. An epistemic defence of the value of bureaucratic discretion is, I would hazard, the prevailing view in the literature, although it is often assumed rather than defended (see Alkhatib and Bernstein 2019 for one defence). On this defence, street-level bureaucratic discretion is epistemically valuable because it promotes desirable epistemic states like true belief, knowledge, and understanding.

Why would street-level bureaucratic discretion better promote truth or knowledge, compared to a lack of discretion? Human beings, one might argue, have certain capacities to elicit and respond to information that artificial agents, say, do not. So, there is a connection between human epistemic capacities, which street-level bureaucrat's role-related powers and prerogatives allow them to exercise in decision-making, and valuable epistemic outcomes.

One such set of capacities allow human bureaucrats to access evidence in the decision context that a rule-based algorithm lacking discretion cannot access. Zacka (2017, Chapter 1) discusses a hypothetical case where a child protective services worker must decide whether a child has been abused. The worker has access to evidence that could be input into an algorithm: facts about bruises on the child's hands, or the parent's record. But, Zacka claims, the worker also has access to evidence that cannot be quantified, and so cannot be an algorithmic input. For example, the worker may have evidence about whether the parent's explanation is convincing, such as tone of voice, the logical flow of the

story, and so on. Since, all else equal, access to more evidence improves the accuracy of decision-making, the worker can make a more accurate judgement than an agent without discretion could.

Another such set of capacities is the ability in a new decision context to form new beliefs about how likely some outcome is, given new information. Alkhatib and Bernstein (2019) argue this point in terms of a human bureaucrat's ability to change their decision boundary before making a decision about a new case. A decision-boundary is used in machine learning to separate outputs into different classes. A human analogue to a decision boundary is a rule by which one maps properties of some person or situation into a class. Take a police officer whose job on patrol is to determine whether someone is breaking into a property. The property of entering through the front door after unlocking it, say, makes someone fall on the 'no' side of the decision boundary, whereas the property of entering through a window makes someone fall on the 'yes' side. However, in a new or rare context, she may do better by shifting her decision boundary. Say the police officer observes a barbecue in full swing in the front lawn of a small house. She then sees a very distressed looking person being lifted into an open first-floor window by a friend. In this case, it is epistemically better for her to shift her decision boundary to classify this event as lawful entry, as the house's occupant has locked themselves out in the middle of their party.

One reason that these epistemic capacities promote more epistemic value is the problem of missing data alluded to above (Alkhatib and Bernstein 2019; Zacka 2017). The data may not be quantifiable, or the event may be rare enough that the best-fixed predictor would not sacrifice overall accuracy by misclassifying it. Or, data may be missing because of the institutional features of bureaucracies. Data collection practices are often primarily for purposes other than model building, such as recording bureaucrats' decisions (Veale and Brass 2019). So, the modeler may not have access to the data needed to build a good model for all cases. Discretion is thus valuable because the exercise of epistemic dispositions can compensate for this missing data.

An epistemic account of the value of discretion is attractive because it negates many of the concerns that one might have about street-level bureaucratic discretion. Exercises of epistemic discretion can improve the fairness of decision-making, because bureaucrats more accurately sort like cases according to the rules. The account also assuages concerns about legitimacy or the creation of relations of unequal standing, because bureaucrats do not bring their own moral judgement to bear on decisions



about clients. And, all else equal, more accurate bureaucratic decisions increase the welfare of clients, as bureaucrats can better identify clients in greater need.

The epistemic case for bureaucratic discretion makes two assumptions, both of which are doubtful. The first assumption is that an agent without discretion could not access the relevant evidence. This is sometimes called the problem of soft evidence (Zacka 2017, Chapter 1). One understanding of the problem of soft evidence is that the problem arises because the relevant fact is not an input to the decision system as it is currently designed. But, the epistemic defence goes, there will always be some evidence that is left out. That is because there will be some properties of people that are not good general predictors, even though they are highly predictive for a small sub-group. But, those properties are not utilised by the system, because doing so would not produce a model that is more accurate overall. A decision-maker that uses discretion, by contrast, is not constrained by the types of evidence that she can take into account in making a decision. So, by using discretion, street-level bureaucrats can make more accurate decisions.

One way of understanding this criticism imputes a modeling failure in the design of the decision system. This version of the criticism, though, can be handled by building a better non-discretionary decision system that takes those missing features into account. Furthermore, decision-makers could use different models for different sub-populations, if the properties that are predictive for each population are non-overlapping. So, a decision system without discretion need not exclude predictively useful features.

A more challenging understanding of the problem of soft evidence is that the problem arises because certain facts cannot be an input into an algorithm in principle (Zacka 2017, Chapter 1). Whether a parent's story is trustworthy, for example, is a fact that the child protective services worker could infer on the basis of available evidence, but an algorithm could not, because that evidence – facts such as tone of voice or facial expression – cannot be an input to the algorithm.

This version of the argument, however, is not as strong as it appears. There are some types of evidence, such as facial expression and tone of voice, which could not have been an input to an algorithm ten years ago, but are now used by hiring algorithms to sort qualified from unqualified job candidates. The more general point is that what types of evidence can be input to an algorithm will depend in part on the current state of technology. Furthermore, we can grant the proponent of this

view the claim that some types of evidence can be perceived by a human decision-maker but not, say, an algorithm that does not exercise discretion, but deny that an algorithm cannot make predictions that are roughly as accurate. Models can utilise proxies, or properties that are correlated with a property of interest, that are almost as predictively fruitful as the property of interest. Common examples of proxies are ZIP code in the United States, which can be a proxy for race, or attending a gender-segregated school as a proxy for gender. A property does not have to be represented by a model in order for the model to make predictions on the basis of this property because of proxies. This is one factor that makes bias in models difficult to detect (Johnson 2020).

Even if we granted, however, that exercises of discretion sometimes promote epistemic value, there is good reason to doubt that they reliably do so amongst street-level bureaucrats, especially in oppressive societies. That is because cognitive bias and discriminatory bias often prevent street-level bureaucrats from responding well to evidence. Cognitive biases were discussed in Section 2, with the example of judicial decision-making. Here I am referring to the kinds of cognitive biases studied by behavioural psychologists, which lead to noisy or inaccurate decision-making and are ubiquitous features of human cognition (Kahneman, Sibony, and Sunstein 2021). Because of cognitive biases, we have reason to doubt that exercises of street-level bureaucrats' epistemic capacities will produce more accurate decisions than algorithms.

In oppressive societies, cognitive biases are compounded by discriminatory biases, which further reduce the accuracy of human decision-making (Sunstein 2022). One mechanism that constitutes or produces bias is how salient particular pieces of information about someone are (Munton 2023; Whiteley 2023), or the inferential patterns that connect pieces of information (Johnson 2020). It is morally troubling, for example, if a female colleague's pastoral care work is always referenced in performance reviews, but a male colleague's research accomplishments are the focus of his meetings (Whiteley 2023). The salience of different pieces of information and differences in inferential patterns being dependent on one's social identity are particularly troubling in bureaucracies because street-level bureaucrats often make decisions based judgments about the client's character. Zacka's (2017) example of a child protective services worker who judges whether a parent is trustworthy is not an isolated example: bureaucrats often sort clients into good types and bad types in order to resolve nonroutine cases (Zacka

2017, Chapter 4). If soft evidence makes certain negative features of someone more salient because of their social identity, or if the bureaucrat tends to infer that someone of a certain social identity is of a bad type, then accessing such soft evidence does not make decisions more accurate. Furthermore, soft evidence often takes the form of narrative, with decision-makers using storytelling to add meaning to client behaviour (Kiviat 2019). But, these are precisely the type of judgments that are often influenced by or encode discriminatory biases tied to social identities (Kiviat 2019).

Therefore, any epistemic defence of discretionary decision-making is, at best, a qualified one: human discretion is not epistemically superior to an epistemic agent that lacks discretion, when human agents are subject to cognitive and psychological bias. This should prompt us to look elsewhere for a defence of the value of human bureaucratic discretion.

#### 4. The moral value of bureaucratic discretion

The alternative to an epistemic defence of the value of street-level bureaucratic discretion is a moral defence. To mount such a defence, we might focus on the morality of judgments, and argue that human street-level bureaucrats can use discretion to make morally better decisions in light of moral values. And, if algorithms cannot exercise moral judgment, then we have an argument in favour of bureaucratic discretion, and reason to limit the scope of algorithmic decision-making in government (Zimmermann and Lee-Stronach 2022).

However, I will take a different tack. Instead, I will follow Zacka (2017) and focus on the moral value of the set of dispositions that bureaucrats use in exercises of discretion. This set contains both epistemic and moral dispositions: epistemic dispositions to attend to certain information, say, but also moral dispositions to weigh values in a certain way, or practical dispositions to take certain means to one's ends.<sup>2</sup> In what follows, I will refer to this set as 'moral dispositions', but that term stands in for a set of epistemic, moral, and practical dispositions (or, just moral and practical dispositions, depending on your philosophical commitments).

---

<sup>2</sup>The arguments do not depend on a distinction between epistemic and moral dispositions. If the supposedly epistemic dispositions discussed in this section and onwards strike you as obviously moral or practical, this is more grist for my mill. But, the arguments should also be palatable to those who maintain the distinction.

In exercising discretion, argues Zacka (2017, Chapter 2), bureaucrats exercise moral dispositions. Moral dispositions ‘shape how bureaucrats perceive and frame the cases they encounter and what considerations they are inclined to prioritise when responding to them’ (2017, 66). Moral dispositions, arguably, shape perception: the moral categories that a bureaucrat tends to apply to clients, for example, may shape whether they perceive a client as angry or as insistent.<sup>3</sup> Moral dispositions also shape attention – what perceptual content or background beliefs are most salient to the street-level bureaucrat – as well as the bureaucrat’s emotional response. And, moral dispositions determine the relative priority that the bureaucrat gives to different reasons in coming to a decision about what to do.

There is a wealth of empirical evidence that street-level bureaucrats act from their moral dispositions in exercising discretion within the parameters established by bureaucratic rules and principles, especially to resolve hard cases. Much of this evidence comes from observations of how street-level bureaucrats resolve hard cases (Maynard-Moody and Musheno 2000; Zacka 2017, Chapter 4). They do not describe their reasoning about hard cases in terms of principles, nor do they often refer to principles in their interactions with other bureaucrats. Instead, bureaucrats first make judgments about a case, within the constraints set out by bureaucratic rules and procedures, and may draw on rules in order to justify the decision. And, these judgments are determined in part by bureaucrat’s moral dispositions, especially by perceived evaluative features of clients, such as being hardworking or being a freerider (Tummers and Bekkers 2014).

These moral dispositions are developed in response to the organisation’s function and its local culture. Evidence for this claim comes from the perhaps surprising extent to which bureaucrats share moral dispositions. For example, they draw on shared moral taxonomies to conceptualise and make decisions about hard cases (Zacka 2017, Chapter 4). Zacka argues that this is because moral dispositions reflect a bureaucrat’s role-ideal, or the bureaucrat’s belief about how she should best carry out her role, given the normative expectations that partly constitute that role.<sup>4</sup> As such, moral dispositions are stable over time, unless the bureaucrat’s role-

---

<sup>3</sup>This is perhaps the most contentious of the three features of moral dispositions. However, a growing literature on perception argues that cognitive and other mental structures can directly influence the content of perception (MacPherson 2012; Munton 2019; Siegel 2017). It is therefore not so far fetched that moral dispositions could shape perceptual content and attention.

<sup>4</sup>The concept of a role-ideal comes from sociology; see Zheng 2018 for discussion.

ideal changes. They are also shared, because workers in the same role tend to develop shared role-ideals, through similar experiences on the job and by learning from others. This also explains why these moral dispositions most often determine behaviour, as they are rooted in the bureaucrat's judgement of how they ought to perform their role (Maynard-Moody and Musheno 2012).

There is evidence that street-level bureaucrats act from moral dispositions, but ought they? The argument that there is a moral value to the exercise of such moral dispositions as part of discretionary decision-making is rooted in the function of bureaucracies to implement policy. In particular, the demands of policy implementation mean that there will always be three types of values that are in tension: fairly implementing policy across all members of the political community, enforcing the state's directives, and helping those in need. And, street-level bureaucrats must draw on moral dispositions in order to make trade-offs between these three values in their discretionary decision-making (Zacka 2017).

This final value is more contentious than the first two, and it is the value that creates the most tension among the set of bureaucratic values. But, there are two compelling defences of it. The first defence is that the procedural benefits of bureaucratic standardisation come at a cost. Bureaucratic rules and regulations pick out certain categories of need, or operationalise them in certain ways, that may leave out those whose needs would otherwise be justifiably recognised by the political community. And, often, it is not possible to amend the rules to correctly capture all cases: doing so will make the rules too cumbersome, and there is no set of rules that is completely action-guiding. By focusing on the needs of specific clients, rather than the rules, the bureaucrat can mitigate the justified complaint that the individual's needs are not appropriately recognised by the administrative agency. The second defence is that street-level bureaucrats play an important fiduciary role, helping clients to navigate complex systems and sometimes advocating on their behalf. Bureaucratic institutions are opaque to most clients, and for exercises of bureaucratic power to be legitimate and fair, clients must have access to translators and advocates to help them navigate these systems (Vredenburg 2022).

Even if these three values are sometimes in tension, however, there may still be an all things considered the best thing to do in every case. The task of a bureaucrat, on this account, is to use their moral judgement to determine which value trumps the others, and to implement the action that best realises that value at the minimum cost to the other values. And,

since the bureaucrat is a normative expert, due to her training and on the job experience, she can reliably choose the reasons that reflect the most weighty value. So, bureaucrats do not have to rely on any specialised moral dispositions; they can just exercise their general capacity for moral judgement.

However, it is doubtful that such reasons are easily accessible to the bureaucrat, outside of the exercises of moral dispositions developed at work. One might doubt there are such reasons in the first place. Some, especially those in the Kantian tradition, have argued that there is not a procedure-independent fact of the matter about what is morally right and wrong in politics; morality is indeterminate until there is an output of a democratically legitimate procedure (Cordelli 2020). Others, following Rawls, have argued that political justifications must be acceptable to each member of the political community; individual moral judgments are unlikely to be acceptable to each member of the political community, given the divergent conceptions of the good in every society.<sup>5</sup> A weaker claim is that even if there is a right course of action, the bureaucrat is unlikely to have epistemic access to it, outside of the exercise of those moral dispositions developed in the workplace. Organisational time pressures and a lack of resources often hinder bureaucrats from deliberating about what to do in light of full information, as discussed in Section 2. Furthermore, acting as a street-level bureaucrat requires factual expertise gained on the job, and there may not be a clear connection between one's everyday moral beliefs and that job-related knowledge. Finally, facts about the bounded rationality of human psychology means that bureaucrats will often employ heuristics and other decision-making shortcuts, as the computational costs of reasoning through every decision about what to do are too high (Simon 1957). These epistemic hurdles will be in place even in just societies with well-resourced and well-intentioned bureaucracies, and they are why street-level bureaucrats rely on the package of epistemic and moral dispositions developed in the workplace.

Even if one accepts the above, however, why think that bureaucrats ought to have stable moral dispositions? Why shouldn't the bureaucrat just choose arbitrarily among the permissible options, if there is no clear right course of action? The reason that moral dispositions are valuable, however, is that they enable consistent and explainable decision-making, which is valuable in bureaucracies. Here we do not need to posit that either consistency or explainability is valuable in itself.

---

<sup>5</sup>For an overview of these two positions and problems with each, see Viehoff 2016.

Instead, consistency and explainability are important for reasons of accountability. Street-level bureaucrats often make decisions about clients absent supervisory oversight, in a manner that is difficult to subject to top-down oversight or bottom-up scrutiny from clients (Zacka 2017, Chapter 4). If bureaucrats chose randomly, then it would be extremely difficult to distinguish problematic partiality from morally permissible random decision-making. Consistency and explainability are also better at enabling peer accountability: if bureaucrats draw on clear, shared moral dispositions that ground the justifications of their choices, then their peers can more easily criticise them for misfirings of those dispositions, or for deciding on the basis of the wrong disposition.

There is thus compelling evidence that street-level bureaucrats make discretionary decisions from a set of moral dispositions, and good moral arguments that they ought to do so. This completes the argument for the first main claim of the paper, that discretion is morally valuable.

## 5. AI as a street-level bureaucrat

Even if one grants the arguments of Section 4, one might doubt that examining the moral dispositions of human agents tells us anything interesting about the AI systems that are used in bureaucratic decision-making. This brings us to the second main claim of the paper, that AI systems often act as street-level bureaucrats. I will argue for this claim by way of two alternative positions. The first alternative is that AI systems are neutral epistemic tools. The second alternative is that AI systems act as policymakers.

One prominent view of algorithmic decision-making tools is that they are only and always epistemic aids, not tools that act on the world; therefore, they are not morally evaluable. Algorithms, one might argue, provide information that is relevant for epistemic states like belief. An algorithm that predicts healthcare costs per patient, for example, provides a decision-maker with information about the future state of the world. It is the role of the decision-maker to integrate that information with her values to decide what to do. A decision-maker who values helping those with the greatest health need may use the information to direct more healthcare resources to those with the greatest need; contrast this action with a decision-maker who values cutting costs, and may place a limit on the resources that can be spent on those with the greatest

needs.<sup>6</sup> So, the proper role of algorithms in administrative agencies is to act as epistemic tools, providing decision-makers with relevant information (Buffat 2015).

Furthermore, since algorithms are always just epistemic tools, they ought to be neutral. Many proponents of that claim understand 'neutrality' as value-freeness, or the claim that non-epistemic values ought not shape scientific processes like gathering evidence, testing hypotheses, and building models (Douglas 2009; Longino 1996; Rudner 1953). An algorithm's designer may have a social goal in mind, such as directing health resources to those most in need, but that social goal should not determine her activity in gathering training data, choosing a learning method, and testing the learned model. The claim can be extended to the technological artefacts that implement these models in code, such as a software program. An artefact is value neutral if it does not have values embedded in it (Miller 2021).

However, a wealth of recent literature in the philosophy of AI has compellingly argued that neither algorithms nor technological artefacts are value-free in all cases. There are three major strands of argument in this literature. The first strand takes inspiration from feminist philosophy of science and epistemology. It argues that non-epistemic values are necessary to resolve problems of underdetermination. There are a number of causal generalisations or future ways that the world could be that are compatible with past observations. Non-evidential assumptions, such as those influenced by moral and social values, are needed to limit the space of possibility in order to build models and make predictions (Dotan 2021; Johnson forthcoming). A second strand of argument takes inspiration from theorising about the value-ladenness of everyday artefacts (Liao and Huebner 2021; Miller 2021; Winner 1980). Artefacts, according to this strand, embody values, and thus are not value neutral. The third major strand intersects with the second. It argues that technology partly constitutes our institutional and social practices, shaping human cognition, constraining attendant behaviour, and influencing the distribution of goods, opportunity, and punishment (Aneesh 2009; Liao and Huebner 2021; Gabriel 2022). As such, it is normatively evaluable, e.g. a site of justice.

If one accepts this third line of argument, it can seem as if AI is a policymaker (Danaher 2016). Like policymakers, AI systems aim to achieve

---

<sup>6</sup>The discussion here is inspired by Henman (2010). Hellman, however, argues that whether an algorithm satisfies the fairness measure *calibration* is relevant to what we ought to believe, not do; she does not make the point about the algorithm itself.



specific value-driven ends. An AI-system may be introduced to detect tax evasion, for example, because a government is concerned that widespread tax evasion has undercut its ability to undertake necessary infrastructure projects. And, both AI systems and policy profoundly impact human cognition, institutions, organisations, and social norms, determining people's access to opportunities and their vulnerability to sanction (Gabriel 2022). This claim also fits with work in public administration research that claims that street-level bureaucratic discretion is reduced by the introduction of information technologies into governance (Bovens and Zouridis 2002).

Many AI systems, however, are better analogised to bureaucrats than to policymakers. Such a view is common in the public administration literature on bureaucratic discretion, where scholars argue that automated systems replace street-level bureaucrats with a 'system-level' bureaucracy, as these systems erode discretion (Bovens and Zouridis 2002; Zouridis, van Eck, and Bovens 2020). The statutes passed by policymakers are operationalised by bureaucrats into rules and regulations for policy implementation. Similarly, designing an AI system requires operationalising an imprecise goal, such as to hire the best job applicant, into a task that an AI system can perform and for which there is available data (Passi and Barocas 2019). Operationalising both policy statutes and imprecise goals requires further value judgments that shape the content of the policy; in that way, both bureaucrats and AI designers also make policy. Both AI systems and top-level bureaucrats tend to classify people into different abstract types or buckets that determine their access to opportunities or vulnerability to sanctions (Maynard-Moody and Musheno 2012; Fourcade and Healy 2013).

However, in many cases, an AI system acts as neither a policymaker nor a top-level bureaucrat. Instead, AI systems often act as street-level bureaucrats (Alkhatib and Bernstein 2019). There are three important parallels between some AI systems and street-level bureaucrats that merit analogising those systems to street-level bureaucrats: the importance of local context for algorithmic and street-level decision-making, peer-to-peer criticism, and decision-making that is not entirely rule-governed.

Street-level bureaucrats work in a particular community, and often work with the same clients over a period of years. As discussed above, the rules and regulations created by top-level bureaucrats may not fit well with the needs and characteristics of their clients. Discretion is an important mechanism that street-level bureaucrats use to address this mismatch between the rules and client needs in a particular population.

Similarly, many algorithms are developed with data that is importantly different from the data generated by the local context. Decision systems are often built by a company that sells them for use across a wide variety of contexts, or by a bureaucratic agency with the resources to develop new models. But, the context-specific assumptions in those models can make them unfit for use in new contexts without amendment. Organisations thus often employ workers to retrain the algorithm by providing it with structured data from the organisational context and evaluating its performance, or to augment algorithmic outputs with local knowledge (Veale and Brass 2019). A fictional example is auditors in a local welfare office, who are tasked with retraining a fraud detection algorithm by providing it with past cases of fraud to learn from and judging whether the algorithm has correctly flagged fraud. The auditors turn the algorithm into another street-level bureaucrat by providing the algorithm with the knowledge gained from their work experience. Retraining can also be automated, and is often implemented in response to a degradation in performance over time (Parisi et al. 2019). In such a case as well, the model is updated in light of learning from the local environment.

This retraining also reflects the importance of peer to peer criticism as a mechanism of accountability in street-level bureaucracy, a second feature in common. Peer criticism is an informal accountability mechanism that can partly mitigate the failures in formal accountability mechanisms (Brehm and Gates 1997; Zacka 2017). Street-level bureaucrats 'relentlessly observe and probe each other's working styles. They intervene directly to praise, criticise, mock, or confront one another. They raise questions, give advice, demand explanations, and pass judgment. Peers serve as a constant, panoptic presence' (Zacka 2017, 182–183). While peer judgement is much more limited between humans and AI systems, worker retraining of algorithms, or an algorithmic output that challenges a street-level bureaucrat's judgement, can perform a function similar to peer criticism.

The third parallel between some algorithms and street-level bureaucrats is that decision-making is not entirely rule-governed. Studies of street-level bureaucrats have found that they sometimes do not reason in terms of the rules and regulations of their organisation or moral principles, a fact which has surprised many scholars (Section 4). Their judgments about clients are certainly constrained by those rules and regulations, and, in routine cases, proceed in accordance with those rules (Lipsky 1980, Chapter 10). But, in non-routine cases, they often make decisions about individuals based on informal moral taxonomies

(Zacka 2017, Chapter 4) or judgments about the moral worthiness of clients (Maynard-Moody and Musheno 2000). Algorithmic reasoning is very different, and is not explicitly moral in this way. But, AI systems are distinct from previous expert systems because outputs are produced via a process that is not easily or aptly described in terms of simple rules. Algorithms developed through machine learning, for example, do not encode existing rules or expert knowledge; instead, learning algorithms find the best model of a set of data. These complex algorithms are useful because they make more accurate predictions; but, it is often difficult to extract a set of rules to describe how the system produces outputs (Creel 2020).

Here one might object that algorithmic outputs are determined by rules, even if those rules are not easily scrutinizable by us. This point connects to recent work by Kahneman, Sibony, and Sunstein (2021), in which they argue that algorithms make better decisions than humans because algorithms produce the same output in response to the same inputs, e.g. are not *noisy*. On this basis, one might posit that algorithms lack discretion. Discretion, according to this view, entails that decision-makers sometimes decide differently about relevantly similar cases: they may exercise problematic partiality towards one client, or may permissibly apply different rules to similar cases due to a lack of information. Furthermore, discretion is also partly constituted by the ability to take into account new evidence before a decision is made, and shift one's decision boundary. An algorithm, by contrast, does not have the flexibility to change the decision boundary between receiving new information and the decision (Alkhatib and Bernstein 2019). So, algorithms do not have discretion, therefore cannot be street-level bureaucrats.

Discretion, however, does not entail that a decision-maker actually decides differently about relevantly similar cases. That is because discretion can also be exercised prior to the set of actual decisions, in fixing an interpretation of the rules and regulations. The designer of an algorithm, for example, must decide on a particular operationalisation of the algorithm's task, one that privileges some values over others (Passi and Barocas 2019). In other words, building an algorithm requires exercising discretion in the choice of a particular understanding of the policy objective, and the algorithm inherits this discretionary choice by the designer. Discretion is also exercised in the learning process. For example, the learning algorithm identifies a subset of the relevant measurable properties to build a predictive model (commonly called *features* in the machine learning literature). Or, in some cases, the algorithm designer chooses which features

the algorithm should learn from. This choice of properties is not merely an epistemic decision; it also often reflects a commitment to what properties ought to matter, from a moral perspective. For example, in the United States, a lender is legally prohibited from making a decision on the basis of gender, and this prohibition obtains even if one's gender is predictive of creditworthiness, or the likelihood that someone will pay back a loan.

Furthermore, we should not take noise or shifting of decision boundaries as constitutive of discretion. That is because street-level bureaucrats in fact often exercise discretion in navigating conflict and underdetermination by the rules without making very noisy decisions. As we saw in Section 4, street-level bureaucrats develop robust dispositions that determine how they tend to respond to cases, influenced by their peers and experience (Zacka 2017). And, many street-level bureaucrats develop a limited set of dispositions, in the face of organisational constraints and culture. When bureaucrats handle client cases from a limited set of dispositions, they tend to respond consistently across clients. And, furthermore, street-level bureaucrats need not change their decision boundaries between receiving evidence in a case and making a decision. A cautious bureaucrat, for example, may prefer to gather evidence across a number of cases, and only then adjust her decision boundary. Neither should therefore be taken as constitutive of exercises of discretion.

Conceptualising some AI systems as street-level bureaucrats offers us a new lens on their functioning and the possibilities for interaction with other street-level bureaucrats and clients. The rest of the paper explores what type of bureaucrat AI is, and how AI systems tend to introduce certain pathologies into organisational decision-making.

## 6. AI as an indifferent bureaucrat

There is some evidence that street-level bureaucrats tend to develop moral dispositions that cluster around three poles or archetypes. I will draw on Zacka's (2017, Chapter 2) typology of three archetypal moral dispositions to argue that street-level algorithmic decision-making in administrative organisations matches the *indifferent* archetype.

Zacka proposes that there are three archetypal moral dispositions that street-level bureaucrats tend to exemplify: the indifferent bureaucrat, the caregiver, and the rule enforcer.<sup>7</sup> He characterises the three as follows:

---

<sup>7</sup>Zacka's typology is developed from an analysis of his own ethnographic research and research by Maynard-Moody and Musheno (2003). It gains support from its resonance with other typologies found in the literature, such as Wilson (1968), Brodtkin (1997), and Watkins-Hayes (2009).

	Indifferent	Caregiver	Rule Enforcer
Which facts are salient and how they are perceived ('hermeneutic grid')	Facts of administrative relevance	Facts that indicate client need, especially strain and distress	That the client is trying to take advantage of the system
Emotional response ('affective attunement')	No emotional response	Sympathy	Guarded and suspicious
How the agent weighs different factors ('normative sensibility')	Greatest weight given to processing the case quickly and efficiently	Greatest weight given to meeting client needs	Greatest weight given to handing out sanctions for rule-breaking

The *indifferent bureaucrat* is the moral type most relevant for our purposes here. The indifferent bureaucrat aims to process cases quickly and efficiently. Often, this focus on people processing comes from the limited time and resources that they have to handle client cases. They focus on facts of so-called administrative relevance, or facts that determine how a client should be treated according to administrative rules and procedures. And, the indifferent bureaucrat does not have an emotional response to the client, even when an emotional response would be fitting.

The three moral dispositions in the typology can be explained by the core values that arise from the function of bureaucratic organisations (Section 4). Zacka (2017, 89) puts the point nicely in terms of three core demands that street-level workers must satisfy:

I suspect that these dispositions recur with some frequency because they delimit the range of attitudes that one can develop toward three core demands of street-level work. Indifference, caregiving, and enforcement involve a commitment, respectively, to people processing (which calls for rapidly sorting clients into predefined categories), service provision (which requires more tailored attention to their specific needs), and regulation (which demands stricter scrutiny with regard to eligibility criteria and compliance with program requirements).

People processing is a core demand of street-level work because it helps achieve fairness and efficiency: treating like cases alike is important for fairness, and matching clients with appropriate benefits or sanctions quickly is important for efficiency. Service provision promotes client welfare, as it matches resources to individuals' needs, guided first and foremost by those needs rather than the rules. And, enforcement is also important for fairness and efficiency: it ensures that only those that qualify under existing rules receive benefits, and that resources are allocated according to the rules. So, we should expect the archetypes of

the indifferent, caregiving, and rule enforcing bureaucrat to be robust across bureaucracies that have those three core values.

AI systems that are used in street-level decision-making act as indifferent bureaucrats. This fact is grounded in properties of AI systems and socio-political facts about how they are designed and deployed (see Section 7 on the latter). As was argued above, AI systems encode a certain operationalisation of the bureaucratic task, and commitment to which properties of people are relevant to accomplish that task. In other words, AI systems encode a commitment to what the administratively relevant facts are. And unless the system is retrained, those facts stay fixed. Algorithmic decision-making systems do not have emotions in response to client cases. And, algorithmic systems can transform inputs into outputs at a speed and scale that far surpasses human decision-makers. More generally, automating a decision process using AI can save time and resources because of faster people processing (Busch and Henriksen 2018). All of these properties of AI systems make for good indifferent bureaucrats.

Furthermore, AI systems are not apt to play the role of a caregiver or rule enforcer. This latter claim may be surprising. Many popular visions of bureaucrats as rule enforcers portray them as enforcing the rules no matter their impact on the person in front of them, a role for which an AI system seems well-suited. But, the moral disposition of a caregiver and a rule enforcer both require a moral sensitivity to an individual client that algorithmic systems lack. In non-routine cases, administratively relevant facts about clients do not tend to capture facts about client need or about rule-breaking: a client may have the same health status as another client, but a less robust social network of potential caretakers; or, two clients may both need caregivers, but one may ask the caregiver to perform tasks beyond their remit. Facts about a client's social network and their treatment of caregivers are not standard inputs to bureaucratic decisions. But, it is in these hard or non-routine cases that the caregiver or the rule enforcer will invest more resources to investigate further, which algorithmic systems do not do. Algorithmic decision-making systems thus make for much better indifferent bureaucrats than regulators or caregivers.

## **7. AI and organisational pathologies**

In this final section, I will argue that AI's role as an indifferent bureaucrat is morally troubling when and because it introduces organisational

pathologies, or an imbalance in the moral dispositions in an organisation (the third main claim of the paper).

One might think that the problem raised by AI's performance of an indifferent bureaucrat is that algorithmic decision-making systems can only have one moral disposition. But, the concern about AI systems goes, a bureaucrat ought to have a plurality of moral dispositions. The ideal bureaucrat has a plurality of dispositions because bureaucrats ought to be sensitive to all the moral values at play in a context and the attendant facts, and enact their role in light of all those values. And, a plurality of moral dispositions help street-level bureaucrats do so.

However, it is not a serious mark against algorithmic decision-making systems that they can only possess one of the three moral dispositions. That is because, no matter what the type of bureaucratic agent, organisational heterogeneity of moral dispositions is more important than individual heterogeneity.

The first reason why organisational heterogeneity is more important than individual heterogeneity is that a moral disposition is a skill that is resource intensive for bureaucrats to develop. As Section 4 argued, bureaucrats' moral dispositions are unique to their role. These dispositions are developed in the course of bureaucrats' work, by learning the values behind the rules and regulations of the agency, interacting with colleagues and clients, and so on (Zacka 2017). And, because they are costly to develop, bureaucrats usually lack the time and resources to develop all three dispositions equally well. One bureaucrat, for example, may be skilled at the empathetic elicitation of sensitive information that bears on the decision, such as information about income, health, disability, or family abuse (Zacka 2017, 167). Another may have greater facility with the rules, enabling faster people processing. And, there are reasons of comparative efficiency for bureaucrats to specialise. I am unlikely to be equally proficient at carpentry, the piano, football, and animation. Some of these skills will take more time and resources for me to develop. And so, if resources are scarce, or it is important that, for each skill, there is a good proportion of agents who have mastered that skill, then agents have reason to specialise. Both of those facts are true of bureaucratic organisations.

Of course, it is important that bureaucrats don't miss the right thing to do when some value clearly trumps the others, and that clients don't have vastly different experiences depending on which type of bureaucrat they happen to end up in front of. The difference between skillful and pathological specialisation is that, when bureaucrats are skillful, they are able to

recognise when a moral disposition is fitting for a particular context, to amend their decision in light of peer or client criticism, and to learn to do both of those tasks better over time. A bureaucrat's dispositions are pathological when they prevent her from doing those things. And, this distinction between skillful and pathological dispositions applies at the level of the organisation as well. A distribution of moral dispositions in an organisation is pathological when there are too few agents with the relevant dispositions to respond to a situation, and when there is insufficient peer criticism or organisational learning over time.

An organisation where there are roughly equal numbers of agents that specialise in different dispositions can avoid organisational pathologies because of the epistemic and communicative structures these subgroups give rise to. Having enough bureaucrats that specialise in each disposition better enables individual bureaucrats to develop the kind of expertise discussed above, as peers can learn from other's moral expertise on the job to better develop their own specialised moral dispositions (Zacka 2017, Chapter 4). The second reason is that evenly balanced subgroups create robust structures of peer criticism. In order for individual bureaucrats to correct mistaken decisions and amend their moral dispositions, they must have opportunities for debate about particular decisions and organisational principles. So, they must be likely to encounter those with different judgments about cases and different affective sensibilities, grounded in a different specialisation in a moral disposition. And, they need to be likely to disagree with each other, rather than let disagreements slide. And bureaucrats are more likely to disagree if they have a supportive community of enough like-minded peers to back their judgments (Zacka 2017, Chapter 4).

We are now in a position to argue for the final claim of the paper: that algorithmic decision-making systems tend to introduce a pathology in organisations. This pathology can arise under at least two common scenarios. The first scenario is that algorithmic decision-making systems come to predominate in an organisation, replacing many of the street-level bureaucrats. An unappreciated effect of such replacement is that it homogenises the moral dispositions in an organisation. And, often, the roles of the remaining bureaucrats are redefined such that they become client-facing problem solvers for the algorithmic system (Eubanks 2018). This leaves bureaucrats with less time and role-based scope to develop and exercise other moral dispositions.

This scenario is a pervasive danger because of policymakers' and top-level bureaucrats' role-based powers and incentives. Policymakers and



top-level bureaucrats are tasked with achieving substantive justice for the political community, in a manner that is legitimate and procedurally fair. And, policymakers have an incentive to make the welfare and justice-promoting impacts of their policies visible to the political community (Hollyer, Peter Rosendorff, and Vreeland 2011). Furthermore, policymakers and top-level bureaucrats have limited time and resources, which is why they formulate general statutes and implement broad brush policy rules and procedures. So, given the obligations and incentives of their role, top-down regulators and policymakers aim to implement efficient people processing in administrative agencies, and to put in top-down controls to ensure that street-level bureaucrats implement policy effectively and legitimately (Tummers and Bekkers 2014). Previously, there was a limit to their ability to promote efficient people processing, as street-level bureaucrats make decisions about clients in face to face interactions that were beyond the reach of their managers (Lipsky 1980). Algorithmic decision-making systems, by contrast, can respond to the particulars of a client case from a fixed model that top-level bureaucrats can better control. So, algorithms present a danger of over-standardization, and a cementation of indifferent people processing dispositions.<sup>8</sup> And, there is initial empirical evidence that automated governance has reduced street-level bureaucrat's ability to make decisions on the basis of the specifics of someone's case (Busch and Henriksen 2018).

However, an organisational pathology can arise even if algorithmic decision-making systems do not predominate in an organisation. And that is because they can hinder individual bureaucrats from exercising caregiving or rule enforcing moral dispositions, due to how these tools influence human cognition. They can also interrupt epistemic and communicative networks within the organisation.

There are psychological reasons why algorithmic decision-making systems can hinder street-level bureaucrats from exercising caregiving or rule enforcing moral dispositions. One such reason is that the classificatory explanations that algorithms enable leave out certain kinds of facts that are important for exercises of such dispositions. Sociological research, for example, has found that decision-makers often use facts about moral responsibility in making decisions about, e.g. loans or

---

<sup>8</sup>Interestingly, administrative law in some jurisdictions already recognizes the danger of what I have called over-standardization. Furthermore, the solution is to encourage street-level bureaucrats to exercise discretion. A UK ruling, for example, encourages bureaucrats to "keep their ears open" when applying policy to a specific case (*British Oxygen v. Minister of Technology* [1971] AC 610.)

insurance pricing, but that algorithms, which classify individuals based on their similarity to others, leave out such facts (Kiviat 2019). Narrative explanations, by contrast, explain why some event happened, by situating someone in a social context and giving the event a meaning that is determined by past events but is over and above the causes of the event (Velleman 2003). Narratives, furthermore, are often marked by an emotional cadence, in which emotions are elicited and resolved (Velleman 2003). But, algorithmic decision-making systems are not suitable bases for narrative explanations, as they do not explain why people have the properties that they do, nor do they involve emotions. Instead, algorithmic decision-making systems enable classificatory explanations, or explanations about why individuals belong to a group based on standardised axes of similarity and difference.

So, classificatory explanations do not furnish the street-level bureaucrat with the raw materials to exercise caring or rule enforcing dispositions. And, furthermore, algorithms tend to make classification-based ways of thinking salient, which can encourage street-level bureaucrats that use such algorithmic systems to be indifferent people processors as well. This defaulting to a disposition of indifference can be further strengthened by the psychological phenomenon of automation bias, or the tendency to defer to the outputs of automated systems.<sup>9</sup>

Algorithms can also interrupt epistemic and communicative networks within an organisation. This interruption is especially likely to occur when algorithms are opaque to the street-level bureaucrats interacting with them. Algorithms developed through machine learning often rely on categories that are inscrutable to human users. They also rely on correlations for prediction and classification that do not intuitively relate to the outcome of interest, in light of existing organisational knowledge, nor can be easily mapped to the phenomena (Barocas and Selbst 2018; Burrell 2016; Sullivan 2022). Finally, the details of algorithms are often protected by intellectual property law (Burrell 2016).

Street-level bureaucrats, however, rely on shared categorisations and rules of thumb in order to make decisions and communicate to others about those decisions. A case manager, for example, may explain to another case manager why they prioritised someone for assistance by explaining that the client has a 'situation', or an urgent problem that

---

<sup>9</sup>See Goddard, Roudsari, and Wyatt (2012) for a review of the experimental evidence of automation bias across domains. Of course, this does not establish that automation bias will be present in a bureaucracy, only that it is likely to be (see, e.g. Alon-Barkat and Busuioc 2023 for evidence that automation bias is not a major driver in bureaucracies).

ought to be prioritised (Zacka 2017, Chapter 4). If algorithms do not use these categories, it is more difficult for bureaucrats to use existing knowledge to reason about the case. And, in cases where there is a high degree of algorithmic opacity, street-level bureaucrats are sometimes left in the situation of Pat, described in the introduction, who does not understand the reasons for an algorithm's determination. Such opacity can also interrupt peer criticism, which is crucial for accountability. One hope is that the integration of algorithms into administrative agencies will improve the accuracy of decision-making (Sunstein 2022). But, if street-level bureaucrats do not understand the reasons behind an output, they cannot find mistakes in their own, or the algorithm's, reasoning. Instead, they either must defer to the algorithm, or trust their own judgement. This undermines the communicative structure of organisations, which usually encourages street-level bureaucrats to disagree about cases.

All of the above tends to drive out caregiving and rule enforcing dispositions, and introduce a problematic skew in the distribution of moral dispositions in an organisation. And, clients are left without bureaucrats whose moral dispositions are attuned to responding to particular cases. But, responding to particular cases is one of the special roles that street-level bureaucrats play within bureaucratic systems. And so, we are left with organisations that have become too standardised.

## 8. Conclusion

Some algorithms act as street-level bureaucrats (Section 5). This surprising claim becomes more compelling when one digs down into the value of bureaucratic discretion (Sections 2-4), and views bureaucratic discretion as a moral disposition (Section 4). And, conceptualising algorithms as street-level bureaucrats puts us in a better position to appreciate further surprising, but important, moral facts. The first is that algorithms often act as indifferent bureaucrats (Section 6). And the second is that algorithmic administration raises concerns about the distribution of moral dispositions in public administrations, and can lead to organisational pathologies (Section 7).

While this paper focused on administrative agencies, similar arguments could be made about algorithmic decision-making in other types of organisations, such as firms, schools and universities, or civil society organisations. The moral value of discretion will partly depend on the function of those organisations. But, a few general points will carry over, such as the ways in which algorithmic decision-making systems

can shift the distribution of moral dispositions in organisations and change epistemic and communicative structures. Viewing AI systems as agents that exercise certain moral dispositions makes a host of important and overlooked questions salient, whose answers ought to inform the design of AI systems and the organisational structures in which they are deployed.

## Acknowledgements

I'm grateful to Milo Phillips-Brown, Clinton Castro, Kathleen Creel, David Gray Grant, Atoosa Kasirzadeh, Silvia Milano, and the editorial board of Inquiry for their comments and suggestions. Special thanks to Marco Meyer for help in generating the central ideas of the paper.

## Disclosure statement

No potential conflict of interest was reported by the author(s).

## References

- Alkhatib, Ali, and Michael Bernstein. 2019. "Street-Level Algorithms: A Theory at the Gaps Between Policy and Decisions." In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–13.
- Alon-Barkat, Saar, and Madalina Busuioc. 2023. "Human–AI Interactions in Public Sector Decision Making: 'Automation Bias' and 'Selective Adherence' to Algorithmic Advice." *Journal of Public Administration Research and Theory* 33 (1): 153–169. <https://doi.org/10.1093/jopart/muac007>.
- Aneesh, A. 2009. "Global Labor: Algocratic Modes of Organization." *Sociological Theory* 27 (4): 347–370.
- Barocas, Solon, and Andrew D. Selbst. 2016. "Big Data's Disparate Impact." *California Law Review* 104: 671–732.
- Barocas, Solon, and Andrew D. Selbst. 2018. "The Intuitive Appeal of Explainable Machines." *Fordham Law Review* 87: 1085–1139.
- Bovens, Mark, and Stavros Zouridis. 2002. "From Street-Level to System-Level Bureaucracies: How Information and Communication Technology Is Transforming Administrative Discretion and Constitutional Control." *Public Administration Review* 62 (2): 174–184.
- Brehm, John, and Scott Gates. 1997. *Working, Shirking, and Sabotage: Bureaucratic Response to a Democratic Republic*. Ann Arbor: University of Michigan Press.
- Brodkin, Evelyn Z. 1997. "Inside the Welfare Contract: Discretion and Accountability in State Welfare Administration." *Social Service Review* 71 (1): 1–33.
- Buffat, Aurélien. 2015. "Street-Level Bureaucracy and E-Government." *Public Management Review* 17 (1): 149–161.

- Burrell, Jenna. 2016. "How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms." *Big Data & Society* 3: 1–12.
- Busch, Peter, and Helle Henriksen. 2018. "Digital Discretion: A Systematic Literature Review of ICT and Street-Level Discretion." *Information Polity* 23 (1): 3–28.
- Cordelli, Chiara. 2020. *The Privatized State*. Princeton: Princeton University Press.
- Creel, Kathleen. 2020. "Transparency in Complex Computational Systems." *Philosophy of Science* 87 (4): 568–589. <https://doi.org/10.1086/709729>.
- Danaher, John. 2016. "The Threat of Algocracy: Reality, Resistance and Accommodation." *Philosophy and Technology* 29 (3): 245–268.
- de Boer, Noortje, and Nadine Raaphorst. 2023. "Automation and Discretion: Explaining the Effect of Automation on How Street-Level Bureaucrats Enforce." *Public Management Review* 25 (1): 42–62.
- Dotan, Ravit. 2021. "Theory Choice, Non-Epistemic Values, and Machine Learning." *Synthese* 198 (11): 11081–11101.
- Douglas, Heather E. 2009. *Science, Policy, and the Value-Free Ideal*. Pittsburgh: University of Pittsburgh Press.
- Eubanks, Virginia. 2018. *Automating Inequality*. St. New York: Martin's Press.
- Fourcade, Marion, and Kieran Healy. 2013. "Classification Situations: Life-Chances in the Neoliberal Era." *Accounting, Organizations and Society* 38 (8): 559–572.
- Gabriel, Iason. 2022. "Toward a Theory of Justice for Artificial Intelligence." *Daedalus* 151 (2): 218–231.
- Garson, David. 2006. *Public Information Technology and E-Governance: Governing the Virtual State*. London: Jones and Bartlett Publishers Int.
- Goddard, Kate, Abdul Roudsari, and Jeremy C. Wyatt. 2012. "Automation Bias: A Systematic Review of Frequency, Effect Mediators, and Mitigators." *Journal of the American Medical Informatics Association* 19 (1): 121–127.
- Henman, Paul. 2010. *Governing Electronically: E-Government and the Reconfiguration of Public Administration, Policy and Power*. Basingstoke: Palgrave Macmillan.
- Hollyer, James R., B. Peter Rosendorff, and James Raymond Vreeland. 2011. "Democracy and Transparency." *Journal of Politics* 73 (4): 1191–1205.
- Johnson, Gabrielle. 2020. "Algorithmic Bias: On the Implicit Bias of Social Technology." *Synthese* 198: 9941–9961.
- Johnson, Gabrielle. *Forthcoming*. "Are Algorithms Value-Free? Feminist Theoretical Virtues in Machine Learning." *Journal of Moral Philosophy*.
- Kahneman, Daniel, Olivier Sibony, and Cass Sunstein. 2021. *Noise: A Flaw in Human Judgment*. New York: Little, Brown, and Company.
- Kiviat, Barbara. 2019. *The Art of Deciding from Data: Evidence from How Employers Translate Credit Reports into Hiring Decisions*.
- Kleinberg, Jon, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan. 2018. "Human Decisions and Machine Predictions." *Quarterly Journal of Economics* 133 (1): 237–293.
- Kolodny, Nico. 2023. *The Pecking Order: Social Hierarchy as a Philosophical Problem*. Cambridge, MA: Harvard University Press.
- Liao, Shen-yi, and Bryce Huebner. 2021. "Oppressive Things." *Philosophy and Phenomenological Research* 103 (1): 92–113.

- Lipsky, Michael. 1980. *Street-Level Bureaucracy: Dilemmas of the Individual in Public Service*. New York: Sage Foundation.
- Longino, Helen E. 1996. "Cognitive and Non-cognitive Values in Science: Rethinking the Dichotomy." In *Feminism, Science, and the Philosophy of Science*, edited by Lynn Hankinson Nelson, and Jack Nelson, 39–58. Dordrecht: Kluwer.
- MacPherson, Fiona. 2012. "Cognitive Penetration of Colour Experience: Rethinking the Issue in Light of an Indirect Mechanism." *Philosophy and Phenomenological Research* 84 (1): 24–62.
- Mashaw, Jerry L. 1983. *Bureaucratic Justice*. New Haven: Yale University Press.
- Maynard-Moody, Steven, and Michael Musheno. 2000. "State Agent or Citizen Agent: Two Narratives of Discretion." *Journal of Public Administration Research and Theory* 10: 329–358.
- Maynard-Moody, Steven, and Michael Musheno. 2003. *Cops, Teachers, Counselors: Stories from the Front Lines of Public Service*. Ann Arbor: University of Michigan Press.
- Maynard-Moody, Steven, and Michael Musheno. 2012. "Social Equities and Inequities in Practice: Street-Level Workers as Agents and Pragmatists." *Public Administration Review* 72 (s1): 16–23.
- Miller, Boaz. 2021. "'Is Technology Value-Neutral?'" *Science, Technology, and Human Values* 46 (1): 53–80.
- Munton, Jessie. 2019. "Perceptual Skill and Social Structure." *Philosophy and Phenomenological Research* 99 (1): 131–161.
- Munton, Jessie. 2023. "Prejudice as the Misattribution of Salience." *Analytic Philosophy* 64 (1): 1–19.
- Noble, Safiya Umoja. 2018. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: New York University Press.
- Parisi, German I., Ronald Kemker, Jose L. Part, Christopher Kanan, and Stefan Wermter. 2019. "Continual Lifelong Learning with Neural Networks: A Review." *Neural Networks* 113: 54–71.
- Passi, Samir, and Solon Barocas. 2019. "Problem Formulation and Fairness." In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT\* '19)*, edited by Sorelle A. Friedler, and Christo Wilson, 39–48. New York: ACM Press.
- Rudner, Richard S. 1953. "The Scientist Qua Scientist Makes Value Judgments." *Philosophy of Science* 20 (1): 1–6.
- Siegel, Susanna. 2017. *The Rationality of Perception*. Oxford: Oxford University Press.
- Simon, Herbert A. 1957. *Models of Man*. New York: John Wiley.
- Sullivan, Emily. 2022. "Understanding from Machine Learning Models." *Philosophy of Science* 73 (1): 109–133.
- Sunstein, Cass R. 2022. "Governing by Algorithm? No Noise and (Potentially) Less Bias." *Duke Law Journal* 71 (6): 1175–1205.
- Tummers, Lars, and Victor Bekkers. 2014. "Policy Implementation, Street-Level Bureaucracy, and the Importance of Discretion." *Public Management Review* 16: 527–547.
- Veale, Michael, and Irina Brass. 2019. "Administration by Algorithm? Public Management Meets Public Sector Machine Learning." In *Algorithmic Regulation*, edited by K. Yeung, and M. Lodge, 121–149. Oxford, UK: Oxford University Press.

- Velleman, J. David. 2003. "Narrative Explanation." *The Philosophical Review* 112 (1): 1–25.
- Viehoff, Daniel. 2016. "Authority and Expertise." *Journal of Political Philosophy* 24 (4): 406–426.
- Vredenburg, Kate. 2022. "The Right to Explanation." *The Journal of Political Philosophy* 30: 209–229.
- Watkins-Hayes, Celeste. 2009. *The New Welfare Bureaucrats: Entanglements of Race, Class, and Policy Reform*. Chicago: University of Chicago Press.
- Whiteley, Ella. 2023. "'A Woman First and a Philosopher Second': Relative Attentional Surplus on the Wrong Property." *Ethics* 133 (4): 497–528.
- Wilson, James Q. 1968. *Varieties of Police Behavior: The Management of Law and Order in Eight Communities*. Cambridge: Harvard University Press.
- Winner, Langdon. 1980. "Do Artifacts Have Politics?" *Daedalus* 109 (1): 121–136. <http://www.jstor.org/stable/20024652>.
- Yarkoni, Tal, and Jacob Westfall. 2017. "Choosing Prediction Over Explanation in Psychology: Lessons from Machine Learning." *Perspectives on Psychological Science* 12 (6): 1100–1122. <https://doi.org/10.1177/1745691617693393>.
- Zacka, Bernardo. 2017. *When the State Meets the Street: Public Service and Moral Agency*. Cambridge, MA: Harvard University Press.
- Zheng, Robin. 2018. "What is My Role in Changing the System? A New Model of Responsibility for Structural Injustice." *Ethical Theory and Moral Practice* 21 (4): 869–885.
- Zimmermann, Annette, and Chad Lee-Stronach. 2022. "Proceed with Caution." *Canadian Journal of Philosophy* 52 (1): 6–25. <http://doi.org/10.1017/can.2021.17>.
- Zouridis, Stavros, Marlies van Eck, and Mark Bovens. 2020. "Automated Discretion." In *Discretion and the Quest for Controlled Freedom*, edited by Tony Evans and Peter Hupe, 313–329. Palgrave Macmillan.