Effectiveness of Algorithm-Based Decision-Making Tools in Street-Level Bureaucracy: Evidence from a Randomized Control Trial and Interviews with Caseworkers

Working Paper for the DPSA 2023

Jonas Krogh Madsen
jkroghm@ruc.dk
Roskilde University

Matthias Döring
mdoering@sam.sdu.dk
University of Southern Denmark

Kim Sass Mikkelsen
ksass@ruc.dk
Roskilde University

Kristian Bloch Haug
kbhaug@ruc.dk
Roskilde University

Disclaimer: *This is early work-in-progress. We tried to add first analyses and results and emphasize the variety of data. However, substantial parts and depth will be added later. We would especially appreciate comments on potential writing strategies regard the scope and structure of such a paper covering various types of data.*

Introduction

The advent of AI-induced technologies, big data, and other related techniques provide organizations with tools that promise substantial benefits regarding efficiency, processing capacity, but also new approaches to increase the effectiveness of public services (Young et al., 2019). Along with the rise of advanced algorithmic tools in public service provision, there has also been a surge of scientific literature that highlights their potential as well as dangers (Andrews 2019; Busuioc 2021, Vogl et al. 2020). Simultaneously, studies addressed how citizen think about these new tools. Various studies indicate general reluctance towards automated or semi-automated processes and artificial intelligence (Aoki, 2020; Caswell et al., 2010; Eubanks, 2019; Zejnilović et al., 2020). While algorithms are often criticized for reinforcing discriminatory patterns underlying the data they are built on (Andrews, 2019; Busuioc, 2020), there is also evidence that minority groups may also view them as a potential remedy for existing discrimination as human biases may be avoided (Miller & Keiser, 2021).

As part of the wider literature of digitalization (Buffat, 2015; Busch & Henriksen, 2018), researchers have examined how employees, such as street-level bureaucrats, perceive and react to these new tools and how various individual and contextual factors shape their implementation of AI-induced technology (Meijer et al., 2021; Assadi & Lundin, 2018; Gillingham & Humphreys, 2010), but also how such tools may challenge the employees' behaviors and routines (de Boer & Raaphorst, 2021; Ranerup & Henriksen, 2020). Particularly relevant in this regard are the decision-making support tools that are supposed to complement professionals' discretion (Meijer et al. 2021; de Boer & Raaphorst 2021; Assadi & Lundi 2018; Gillingham & Humphreys 2010).

However, most existing studies focus on employee outcomes alone, and rarely assess these tools' impact on public services. As a consequence, there are surprisingly few empirical studies investigating how the implementation of such tools matter to organizational as well as citizen outcomes. This leaves open the question of whether these decision-making tools actually matter to public service delivery, and how such effects (or lack of the same) are mediated by the attitudes and contextual conditions by the frontline employees who implement them in practice.

This study combines a pre-registered randomized field experiment, a pre-registered conjoint experiment, a survey among frontline caseworkers, participant observation data during the preliminary rollout phase, and 20 interviews with frontline caseworkers to investigate how the introduction of such a tool affects the service provision and situation of unemployed benefit recipients. This combination of methods and data allows for the evaluation of the tool's effectiveness while granting insights into the underlying micro-mechanisms that explain the overall findings. By exploiting this rich multi-method data, we are able to investigate the implementation process in theoretical and empirical depth as well as testing causal effects for relevant outcome variables in the unemployment administration sector. The data was collected before, during, and after the rollout of an algorithm-based decision-making support tool in a large Danish unemployment fund including their job centers.

This study provides three major contributions. First, we offer seminal evidence on the effects of implementing algorithm-based decision-making support tools in public organizations. Thus, we move beyond a pure intra-organizational perspective that most studies in the public administration literature have focused on so far. Second, we provide an in-depth analysis of the rollout process and the challenges encountered during and after the official rollout. Thus, we are able to provide qualitative evidence of the caseworkers' motives and norms that clashed with the newly implemented tool. Third, we contribute to the literature on the relationship between algorithm-based tools and professionals emphasizing that street-level bureaucrats may use such tools in various ways, for example by giving their judgment a seemingly objective support 'by the data' as well as by facilitating communication with their clients.

Theory

Street-level bureaucracy is typically characterized by professionals implementing public policies using their discretion to translate broader legal frameworks to the complexity of reality (Baviskar & Winter, 2016; Brodkin, 2015; Lipsky, 2010; Maynard-Moody & Musheno, 2003; Tummers et al., 2012). For that purpose, they often build on their specialized education, norms, and values (Brodkin, 2011). However, as Lipsky (2010) points out, street-level bureaucrats, such as caseworkers, suffer from a systematic struggle for resources, as there is a sheer unlimited demand for their services. At an individual level, caseworkers will make use of various coping strategies to solve this struggle. At the organizational level, digitalization and automation has been perceived as a remedy and therefore widely discussed. As caseworkers have to work with limited information as well as limited processing capacities, they will have to satisfice for deriving appropriate decision on how to handle individual cases (Tummers et al., 2015). Which case needs more attention, which case will have no trouble in solving the particular situation? Algorithm-based decision-making support tools are ideally supposed to provide a solution to process larger amounts of data to make better decisions. However, the role of digital tools in the context of street-level bureaucracy is a topic of dispute. Buffat (2015) summarizes the discussion using two hypotheses: the enablement and curtailment hypothesis. The former emphasizes the opportunities to automatically process easy cases that require little professional expertise but would normally still bind a substantial number of resources. Hence, by relieving caseworkers of such cases, they are able to focus on more critical or complex cases. The curtailment hypothesis, however, takes a critical stance that digitalization is often accompanied by standardization which in turn reduces discretion that is critical for effective implementation of policies.

This conflict becomes even more critical when putting in the context of AI-induced technologies in which algorithms are often too complex for the caseworkers to comprehend. This may lead to three alternative reactions to algorithm-based decision-making support systems. First, caseworker might lack the trust in the respective tool due to a “black-box effect” (Castelvecchi, 2016). If not forced to use the tool, caseworkers might find workarounds or simply not use the tool purposefully (Gillingham & Humphreys, 2010). Second, caseworkers may trust the tools results, suggestion, or advice more than their own reasoning or simply take the tool’s result for granted. This phenomenon might especially be prevalent with caseworkers that lack experience (Assadi & Lundin, 2018). Third, and the best-case scenario, both caseworkers and the algorithm-based support tool live work complementary as caseworkers will use the tool as another source of information for especially critical or vague cases.

For our specific case setting of unemployment counselling, we assume that the risk-assessment tool will be able to effectively and efficiently direct attention to cases that are more in need of attention. Thus, caseworkers will spend more time on counselling critical cases that might have a harder time finding a new job. Furthermore, caseworkers might be able to better personalize their counselling efforts so that they can identify shortcomings in their clients’ competences and suggest the appropriate measures to remedy these shortcomings. Overall, the introduction of such an algorithm-based decision-making support tool is supposed to increase the organization’s effectiveness by supporting its caseworkers. In the case of unemployment counselling, we would expect that clients will have an easier time returning into new work – especially if they are identified as critical cases.

While these tools are primarily used for informing the caseworker, they may also affect how clients perceive their situation and measures that are derived from such tools. Several studies have examined how citizens may view algorithm-based or AI-induced technologies as potential threats to equality (Andrews, 2019; Busuioc, 2020), but also as potential solutions to decrease discrimination (Miller & Keiser, 2021). As our case setting emphasizes the role of co-production – whether unemployed will find a new job is primarily up to their own actions – self-reflection may play a critical role. If clients identify themselves as easy cases that will have little trouble finding a new job, being told that this might not be true will trigger resistance due to cognitive dissonance (Festinger, 2001). However, if the assessment is supported by an algorithm-based support tool, this 'reality shock' might be easier to accept as the information is (seemingly) not only based on the personal assessment from the caseworker. Hence, algorithm-based decision-making support tools may be seen as a more objective adviser compared to the job consultant.

Hence, we formulate the following hypotheses:

Hypotheses

H1: The use of the risk-assessment tool decreases the time that unemployed clients need to re-enter the job market.

H2: The use of the risk-assessment tool increases the client's perception that the job counselling is fitted to their individual needs.

H3: The use of the risk-assessment tool decreases the number of official meetings between counsellors and clients.

H4: The use of the risk-assessment tool increases the uptake of courses for those being rated at high-risk.

[in its current state, we provide evidence for H1 as well as partially H4. The future analysis will be more comprehensive including H2 and H3]

All hypotheses as well as their analyses have been pre-registered (https://osf.io/n4p3f).

Data and methods

Setting

This study focuses on the Danish unemployment insurance sector. Compared to other countries, the Danish unemployment administration system is traditionally decentralized. Instead of having a single governmental agency providing unemployment benefits, job counseling and educational support are partially tasked to unemployment insurance funds, many of which have traditional bonds to the Danish unions. The funds are legally private organizations that provide official services delegated by the state. Their work is regulated by government, while they enjoy substantial discretion in their organization of operations. Often funds represent clients from specific work sectors. The fund we worked with for this paper, from here simply "the fund", focuses on blue-collar workers.

We focus on the implementation of an algorithmic risk assessment tool that was rolled out across 60 departments the fund. The general obligation to implement such a tool was set by the Danish legislator. However, individual funds were free to develop their own tools. While a government developed tool (developed by the government Agency for Labor Market and Recruitment) was widely adapted by other funds, the fund opted to rely on a tool it developed for itself.

The implemented risk assessment tool draws on data about the insured members of the fund (e.g., education, previous employment, age, and geographical information) to develop a statistical model that could predict risks of clients' long-term unemployment during their first meeting in the unemployment insurance fund at the very beginning of their unemployment spell. Here, caseworkers were asked to use a graphical user interface of model output that showed the client's risk of long-term unemployment and as well as the key metrics behind the prediction. The caseworker would then see a traffic-light system, providing simplified information about the client's risk of unemployment. The tool was designed to help underpin caseworkers' job counselling and service provision and help them target resources in a more meaningful way by providing additional support for those at high risk of long-term unemployment, while reducing efforts for 'easy cases'.

Experimental design

In liaison with the fund's main office, our research team were given the opportunity to support and examine the rollout of the decision support tool in a randomized controlled trial. Based on our input, the fund randomly assigned half of the departments to a treatment group where caseworkers had to use the tool during every first meeting with newly unemployed benefit recipients. In the control group, the caseworkers were not able to use the tool until the investigation period was over. As our research group was given full control over the randomization process, we can test causal claims about the impact of a single change episode on our variables of interest: change cynicism and change fatigue.

The study thereby follows a between-subjects design in which we manipulate one-factor (the introduction of the risk-assessment tool during the first meeting) with two levels. To help ensure that departments were matched on size, we matched departments in pairs based on their number of benefit recipients, and used matched-pair cluster randomization at the level of the fund's departments to randomly assign one department of each pair to the treatment group and the other to the control group. Moreover, to minimize the risk of treatment contagion, we treated two sets of large-city departments with overlapping jurisdictions (three in Odense and two in Copenhagen) as single departments. The rollout of the experimental treatment started in February 2022 and ended in November 2022.

Data

The declared goal of the introduction of these risk assessment tools is to efficiently use existing resources in welfare services, particularly the unemployment sector. While this is no new ideal in the paradigm of the "activating welfare state" (Arts & Van Den Berg, 2019; Dingeldey, 2007), algorithmic decision-making tools are considered a new opportunity that goes beyond professional discretion enacted by the caseworkers. Accordingly, clients that are able to handle themselves and their own cases should be left alone as consultation and training efforts are not needed and may even have negative effects. Rather, consultation efforts should focus on clients in need. However, the identification of what is considered to be an easy case may not always be as evident as one might assume. While proxies, such as education, gender, or age, may serve as useful heuristics used by experienced job counsellors, it is quite possible that potential risks of unemployment may be overlooked or falsely attributed. Hence, these decision-making tools have been introduced to complement caseworkers' professional expertise. Our respective unemployment fund decided to develop an assessment tool that is self-sufficient regarding necessary data. Hence, the tool only uses data already available to the fund. Therefore, a tool used by various other funds has been rejected based on this requirement. The implemented tool uses logistic regression to assess the odds of becoming long-term unemployed based on various characteristics, e.g., age, gender, previous unemployment episodes, region, and trade. The individual coefficients are aggregated and displayed to the caseworker in the form of traffic lights to facilitate the interpretation.

To ensure an encompassing picture of the tool's introduction, implementation, and effects, we collected data through multiple sources both prior to and during the rollout. In total, we possess data from five different sources.

First, administrative data on actual use of the decision support tool was collected by the fund as part of its normal benchmarking procedures. This data contains data about each meeting during the rollout phase in which the tool was to be used, whether or not the client gave consent, and longitudinal data about the clients' employment and educational outcomes. Moreover, it contains background data on client characteristics (such as gender, age, and unemployment length). The final dataset contains a total of 6.306 meetings during the rollout phase. With this data, we are able to estimate the effects of the rollout.

Second, a survey was fielded to the frontline caseworkers across both treatment and control groups during the rollout phase between late September and mid-October 2022. This data contains several relevant attitudinal measures, such as their change scepticism, fatigue, perceived usefulness of the tool, and policy (in)consistency, in addition to several different background characteristics of the caseworkers. The survey had a response rate of 43% and contains 120 full responses. With this data, we are able to estimate whether caseworker attitudes and characteristics moderate the implementation and effects of the tool.

Third, the survey to the frontline caseworkers also contained a pre-registered conjoint experiment. Here, the caseworkers were presented with four different vignettes containing client descriptions that randomly varied on central characteristics (such as with a series of age, unemployment length, educational level, mental health, and unemployment responsibility attribution). After each vignette, the caseworkers were asked to assess to what extend the tool they intended to use the tool on the client, how the tool would affect the quality of their counselling, and, finally, whether using the tool would make it easier to counsel the clients.

For this conjoint experiment that has been pre-registered (https://osf.io/nufcq), we compared four attributes:

- **Name**: Martin, Thomas, Michael, Morten, Jesper, Henrik, Anders, Lars, Christian, Søren
- **Age**: 30, 33, 36, 39, 42, 45

- **Unemployment background**: a) been fulltime employed, b) been fulltime employed excluding a one unemployment episode of three months in the past, c) been fulltime employed excluding three unemployment episode of three months in the past
- **Education**: a) no further education after "folkeskolen", b) short additional education
- **Health**: a) no health-related challenges, b) diagnosed with ADHD five years ago
- **Literacy**: a) no problems with reading and writing, b) some challenges with reading and writing
- **Prior encounters**: a) has not been counselled by you, b) has been counselled by you in the past
- **Learning costs**: a) [no information], b) during the meeting they state that they had difficulties in assessing and understanding information regarding unemployment benefits, c) during the meeting they state that they found it easy to assess and understand information regarding unemployment benefits
- **Responsibility**: a) [ingen information], b) during the meeting, they state that it's primarily your responsibility as a job counsellor to get them back into work, c) during the meeting, they state that it's primarily their responsibility to get them back into work

After being shown a vignette with randomized levels of the given attributes, respondents were asked the following question: "How likely is it that you will offer to let the statistical profiling tool assess his risk of long-term unemployment?" [Slider: 0 Completely unlikely – 100 Completely likely]. Each respondent was asked to complete this task four times in succession with randomized attribute levels.

This data contains a total of 420 trials. With this data, we are able investigate whether certain client characteristics impact the caseworkers view of the tool's quality, usefulness, and intend-of-use.

Fourth, every second week during the rollout phase we fielded a survey to all clients that had attended a meeting within the prior two weeks. Here, we measured the clients' job search self-efficacy, job search behavior, job search clarity, evaluation of counsellor and the fund, trust in the fund, and assessment of the tool. Although we fielded the survey across 13 waves, the survey had a response rate of less than 10%, giving us a total of 3xx responses. With this data, we are able to investigate how clients perceive the counselling, the use of the tool, and whether the rollout affected their self-reported job search behaviour and attitudes.

Finally, we conducted 20 interviews with caseworkers during the pilot and preliminary phase of the rollout as well as observational data of actual meetings in which the tool was used. With this data, we get a deeper insight into the caseworkers' perception of the tool, and its usefulness.

Results & Discussion

Our results show that the introduction of the tool generally had a very limited impact on clients' unemployment. When comparing the average days of unemployment among clients of the control group and the treatment group (subtracting treatment from control group), we find no significant difference between the two (see Figure 1).

Figure 1: Effect on days of unemployment

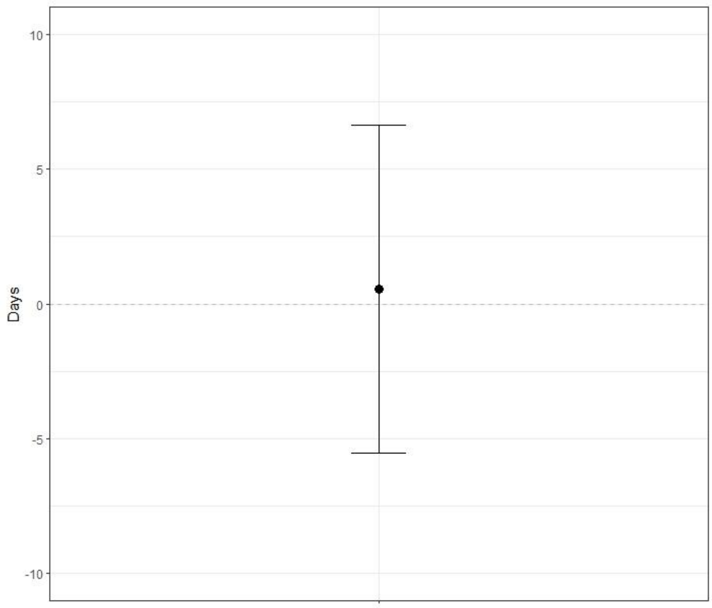

Accordingly, assessing the likelihood of occupational status changes among members did not differ when comparing control and treatment group (see Figure 2).

Figure 2: Effect on unemployment status

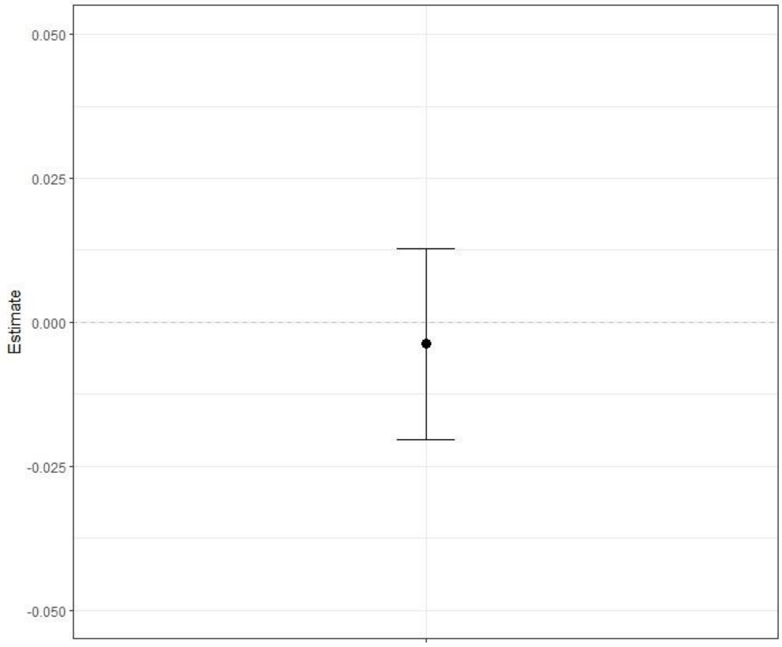

Figure 3: Effect on educational activity

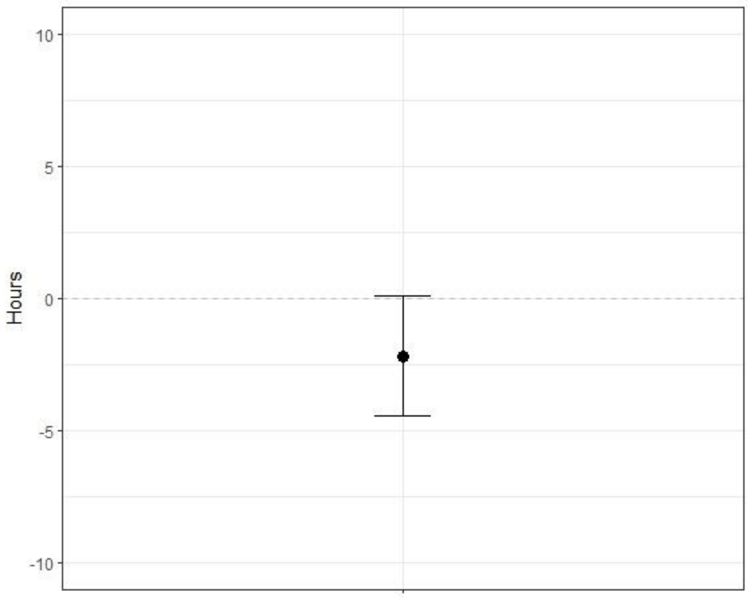

Finally, assessing the number of educational hours prescribed to newly unemployed, we see a weak effect of the treatment, implying that clients in the treatment group received fewer hours of educational courses. [We will follow-up on this with an analysis of (number of hours) in relation to their risk score]. The data show a small negative effect on the hours spent on occupational education for the unemployed (e = -2.17, p = 0.0649).

The identified lack of effect can be attributed to three potential reasons. First, the tool itself is generally not effective in achieving these goals. Second, a larger time gap is necessary before changes create a meaningful impact and effects become detectable. Third, there was a substantial issue with non-compliance with the experimental treatment (using the algorithmic risk assessment), thus minimizing its effectiveness in the treatment group. Indeed, we find convincing evidence for the latter. Figures 4 and 5 indicate that only a minor share of departments and caseworkers actually employed this tool despite being officially mandatory. Before caseworkers were allowed to use the tool, clients would have to give formal consent. Accordingly, the small and varying share of consent given within different departments and among different caseworkers indicates substantial non-compliance. The tool has widely not been used as much as it should have been. Here, our qualitative data as well as the conjoint experiment provide additional insights into the reasons for this non-compliance.

Figure 4: Share of consent per meeting by department

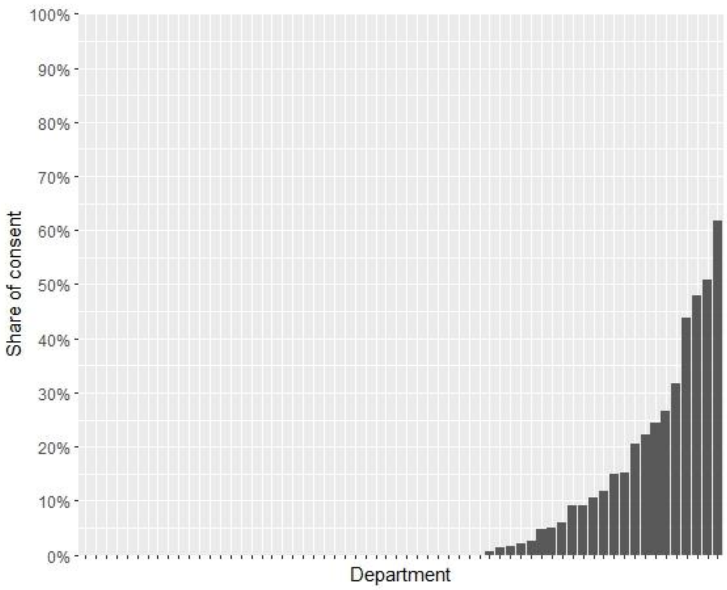

Figure 5: Share of consent per meeting by caseworker

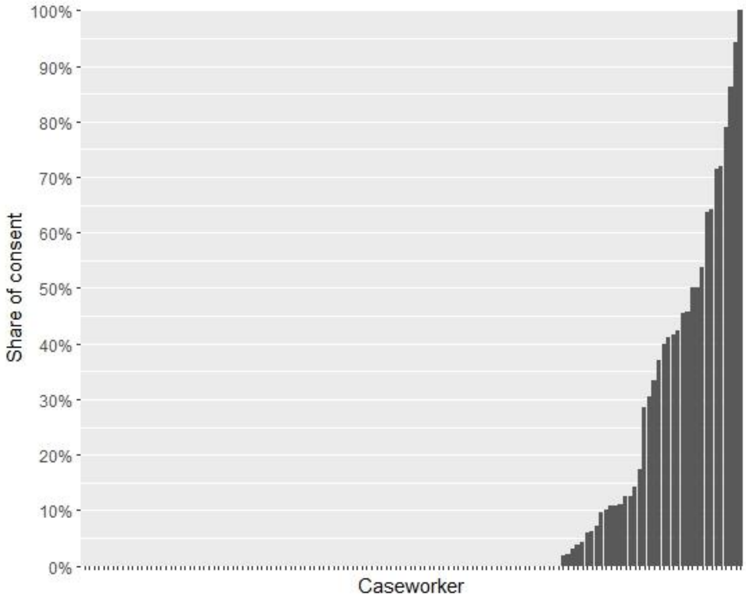

Our qualitative evidence from the pilot and the preliminary phase of the rollout suggests that, although the tool was largely perceived as easy to use, the caseworkers had various concerns about introducing the tools in their counselling of clients, and particularly during their first meeting. For instance, several caseworkers believed that the tool would lack effectiveness, that it stood in the way of meaningful conversations and trust-based relation building with the client, that their own risk assessment ability were – at least – equally as good as the tool's.

"det at risikovurdere, det har vi på rygraden" … "vi gør det i forvejen baseret på vores erfaring og intuition"

Furthermore, and perhaps most importantly, that the tool was to be employed in meetings that were already saturated with compulsory communication about the many rules and demands in the benefit system.

"Det kan let blive til endnu en arbejdsgang – samtykke skal scannes ind og gemmes (…) og er værktøjet tidsforbruget værd?"

Moreover, caseworkers feared in advance of the rollout that the tool might be too formalistic, potentially risking the trust-based and intimate relationship between client and counsellor.

"I dialogen med medlemmerne skal vi skabe tillid og tryghed først. Værktøjet kan risikere at stemple."

"Det der betyder noget er medlemmet, ikke IT-systemet og hvad det siger. Jeg frygter, at det kan blive styrende for samtalen, hvis scoren sættes over medlemmet"

The caseworkers did however believe that the tool's output could help communicate risks and convince certain high-risk clients of the importance of changing their job search behavior.

"Det visuelle, at man kan vise medlemmet det – det er den eneste åbenbaring"

The latter finding that was also supported by the results of our conjoint experiment. Here, we employed a survey among all caseworkers to indicate clients' characteristics and circumstances that would affect their likelihood of using the tool in the first place.

The results are shown in Figure 6 below.

Figure 6: Likelihood for using the tool

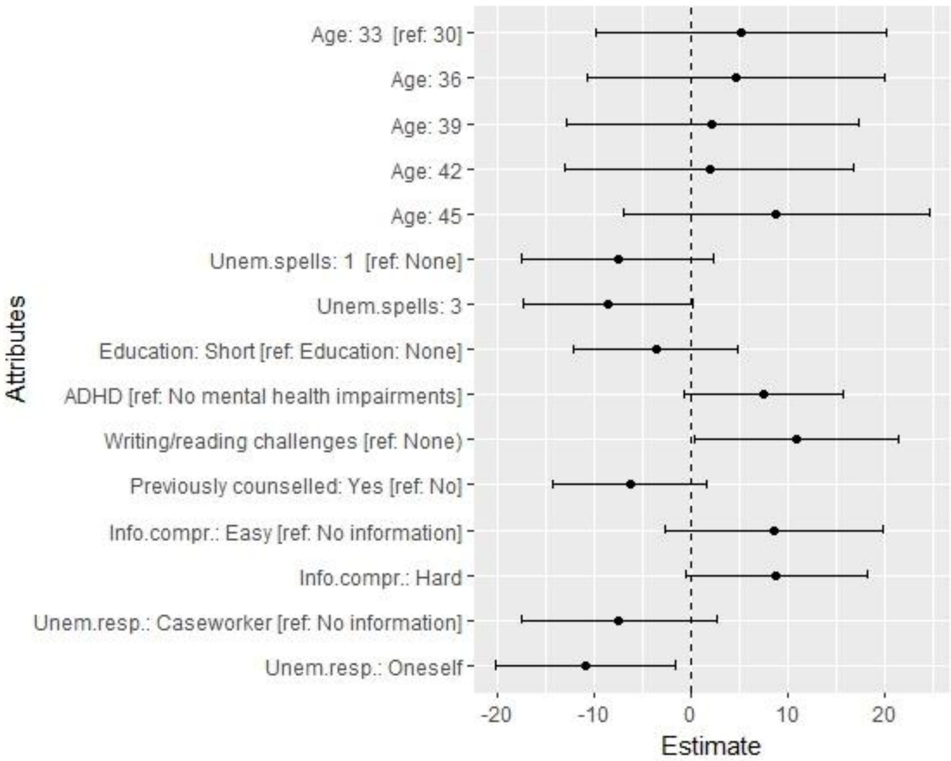

The results indicate that three factors increase the likelihood of tool usage: ADHD diagnosis, reading and writing skills, as well as difficulties in understanding information regarding unemployment benefits. These factors can be summarized as potential traits of 'tough cases' in which standardized written information pieces and material is not sufficient to inform the client. Due to the general time-constraints that caseworkers suffer from, extensive elaborations and explanations are often impossible to include within first meetings. Thus, the tool can be considered as an additional communication tool to make clients of their situation and challenges.

Simultaneously, if clients shirk responsibility for leaving unemployment, caseworkers are less likely to use the tool. This could be explained by a general frustration caseworkers may encounter in such instances in which 'convincing the client' might not be an option. Alternatively, the caseworker might want to use more time on emphasizing the client's responsibility and not 'wasting time' with the tool. In conclusion, the tool is considered an occasionally useful communication tool to deliver messages the caseworkers have already made up their mind about – and less as a tool to solve ambiguity. This communication, however, seems to be highly context specific depending on the concrete challenges the caseworkers face within the public encounter.

Conclusion

In sum, the study underlines the importance of comprehensive change management while highlighting that such tools should only be used in settings where the application also leads to substantial resource savings. Here, analyses of the demand on street-level bureaucrats may prove essential. At the same time, even if tools are supposed to serve clients (and not necessarily caseworkers), street-level bureaucrats can serve as gatekeepers that regulate access to such tools.

[…]

Literature

Androutsopoulou, A., Karacapilidis, N., Loukis, E., & Charalabidis, Y. (2019). Transforming the communi-cation between citizens and government through AI-guided chatbots. Government Information Quarterly, 36(2), 358–367. https://doi.org/10.1016/j.giq.2018.10.001

Aoki, N. (2020). An experimental study of public trust in AI chatbots in the public sector. Government Information Quarterly, 37(4), 101490. https://doi.org/10.1016/j.giq.2020.101490

Assadi, A., & Lundin, M. (2018). Street-level bureaucrats, rule-following and tenure: How assessment tools are used at the front line of the public sector. Public Administration, 96(1), 154–170. https://doi.org/10.1111/padm.12386

Baviskar, S., & Winter, S. C. (2016). Street-Level Bureaucrats as Individual Policymakers: The Relationship between Attitudes and Coping Behavior toward Vulnerable Children and Youth. In International Public Management Journal (Vol. 20, Issue 2, pp. 316–353). https://doi.org/10.1080/10967494.2016.1235641

Brodkin, E. Z. (2011). Policy Work: Street-Level Organizations Under New Managerialism. In Public Administration Review (Vol. 21, Issue Supplement 2, pp. 253–277). https://doi.org/10.1093/jopart/muq093

Brodkin, E. Z. (2015). The inside story: Street-level research in the US and beyond. In P. Hupe, M. Hill, & A. Buffat (Eds.), Understanding Street-Level Bureaucracy (pp. 25–42). Policy Press.

Buffat, A. (2015). Street-Level Bureaucracy and E-Government. In Public Management Review (Vol. 17, Issue 1, pp. 149–161). https://doi.org/10.1080/14719037.2013.771699

Busuioc, M. (2020). Accountable Artificial Intelligence: Holding Algorithms to Account. Public Admin-istration Review. https://doi.org/10.1111/puar.13293

Busch, P. A., & Henriksen, H. Z. (2018). Digital discretion: A systematic literature review of ICT and street-level discretion. Information Polity, 23(1), 3–28. https://doi.org/10.3233/IP-170050

Caswell, D., Marston, G., & Larsen, J. E. (2010). Unemployed citizen or ‘at risk’ client? Classification systems and employment services in Denmark and Australia. Critical Social Policy, 30(3), 384–404. https://doi.org/10.1177/0261018310367674 de Boer, N., & Raaphorst, N. (2021). Automation and discretion: Explaining the effect of automation on how street-level bureaucrats enforce. Public Management Review, 1–21. https://doi.org/10.1080/14719037.2021.1937684

Dingeldey, I. (2007). Between workfare and enablement ? The different paths to transformation of the welfare state: A comparative analysis of activating labour market policies. European Journal of Political Research, 46(6), 823–851. https://doi.org/10.1111/j.1475-6765.2007.00712.x

Eubanks, V. (2019). Automating inequality: How high-tech tools profile, police, and punish the poor (First Picador edition). Picador St. Martin’s Press.

Festinger, L. (2001). A theory of cognitive dissonance (Reissued by Stanford Univ. Press in 1962, renewed 1985 by author). Stanford Univ. Press.

Gillingham, P., & Humphreys, C. (2010). Child Protection Practitioners and Decision-Making Tools: Ob-servations and Reflections from the Front Line. British Journal of Social Work, 40(8), 2598–2616. https://doi.org/10.1093/bjsw/bcp155

Kolkman, D. (2020). The usefulness of algorithmic models in policy making. Government Information Quarterly, 37(3), 101488. https://doi.org/10.1016/j.giq.2020.101488

Larsson, K. K. (2021). Digitization or equality: When government automation covers some, but not all citizens. Government Information Quarterly, 38(1), 101547. https://doi.org/10.1016/j.giq.2020.101547

Meijer, A., Lorenz, L., & Wessels, M. (2021). Algorithmization of Bureaucratic Organizations: Using a Practice Lens to Study How Context Shapes Predictive Policing Systems. Public Administration Review, 81(5), 837–846. https://doi.org/10.1111/puar.13391

Miller, S. M., & Keiser, L. R. (2021). Representative Bureaucracy and Attitudes Toward Automated Decision Making. Journal of Public Administration Research and Theory, 31(1), 150–165. https://doi.org/10.1093/jopart/muaa019

Ranerup, A., & Henriksen, H. Z. (2020). Digital Discretion: Unpacking Human and Technological Agency in Automated Decision Making in Sweden’s Social Services. Social Science Computer Review, 089443932098043. https://doi.org/10.1177/0894439320980434

Tummers, L., Bekkers, V., Vink, E., & Musheno, M. (2015). Coping During Public Service Delivery: A Conceptualization and Systematic Review of the Literature. In Journal of Public Administration Research and Theory (Vol. 25, Issue 4, pp. 1099–1126). https://doi.org/10.1093/jopart/muu056

Tummers, L., Vermeeren, B., Steijn, B., & Bekkers, V. (2012). Public Professionals and Policy implementation: Conceptualizing and measuring three types of role conflicts. In Public Management Review (Vol. 14, Issue 8, pp. 1041–1059). https://doi.org/10.1080/14719037.2012.662443

Young, M. M., Bullock, J. B., & Lecy, J. D. (2019). Artificial Discretion as a Tool of Governance: A Framework for Understanding the Impact of Artificial Intelligence on Public Administration. Perspec-tives on Public Management and Governance, gvz014. https://doi.org/10.1093/ppmgov/gvz014