A Framework of High-Stakes Algorithmic Decision-Making for the Public Sector Developed through a Case Study of Child-Welfare

DEVANSH SAXENA, Marquette University, USA

KARLA BADILLO-URQUIOLA, University of Central Florida, USA

PAMELA J. WISNIEWSKI, University of Central Florida, USA

SHION GUHA, University of Toronto, Canada

Algorithms have permeated throughout civil government and society, where they are being used to make high-stakes decisions about human lives. In this paper, we first develop a cohesive framework of algorithmic decision-making adapted for the public sector (ADMAPS) that reflects the complex socio-technical interactions between *human discretion*, *bureaucratic processes*, and *algorithmic decision-making* by synthesizing disparate bodies of work in the fields of Human-Computer Interaction (HCI), Science and Technology Studies (STS), and Public Administration (PA). We then applied the ADMAPS framework to conduct a qualitative analysis of an in-depth, eight-month ethnographic case study of algorithms in daily use within a child-welfare agency that serves approximately 900 families and 1300 children in the mid-western United States. Overall, we found that there is a need to focus on strength-based algorithmic outcomes centered in social ecological frameworks. In addition, algorithmic systems need to support existing bureaucratic processes and augment human discretion, rather than replace it. Finally, collective buy-in in algorithmic systems requires trust in the target outcomes at both the practitioner and bureaucratic levels. As a result of our study, we propose guidelines for the design of high-stakes algorithmic decision-making tools in the child-welfare system, and more generally, in the public sector. We empirically validate the theoretically derived ADMAPS framework to demonstrate how it can be useful for systematically making pragmatic decisions about the design of algorithms for the public sector.

CCS Concepts: • **Human-centered computing** → **Human-computer interaction (HCI)**; *Empirical studies in HCI*; • **Applied computing** → Computing in government.

Additional Key Words and Phrases: algorithmic decision-making, discretion, bureaucracy, child-welfare system

ACM Reference Format:

Devansh Saxena, Karla Badillo-Urquiola, Pamela J. Wisniewski, and Shion Guha. 2021. A Framework of High-Stakes Algorithmic Decision-Making for the Public Sector Developed through a Case Study of Child-Welfare. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 348 (October 2021), 41 pages. <https://doi.org/10.1145/3476089>

1 INTRODUCTION

The influence of neoliberal politics and theories of New Public Management (NPM) [68] throughout most modern societies over the past two decades has sought to reform public services by emulating

Authors' addresses: Devansh Saxena, Marquette University, Cudahy Hall, 1313 W Wisconsin Avenue, Milwaukee, WI, 53233, USA, devansh.saxena@marquette.edu; Karla Badillo-Urquiola, University of Central Florida, 4000 Central Florida Blvd, Orlando, FL, 32816, USA, kcurquiola10@knights.ucf.edu; Pamela J. Wisniewski, University of Central Florida, 4000 Central Florida Blvd, Orlando, FL, 32816, USA, pamwis@ucf.edu; Shion Guha, University of Toronto, 140 St. George Street, Toronto, Ontario, Canada, shion.guha@utoronto.ca.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2573-0142/2021/10-ART348

<https://doi.org/10.1145/3476089>

corporations to improve efficiency [70]. One way to achieve this goal for public sector services (e.g., child-welfare, labor, criminal justice, and public education) is through the adoption of automated processes (e.g., decision-making algorithms), as they purportedly promise to increase efficiencies, lower costs, and provide better outcomes for citizens. [102]. As such, algorithms in the public sector have become pervasive and, in turn, well-studied in recent years [28, 42, 81]. Consequently, they have also been scrutinized for achieving worse outcomes, exacerbating racial biases, and strengthening structural inequalities [42, 101, 107] within systems that are overburdened and under-resourced, yet critically needed [33, 75].

As a case in point, 423,997 children were in the U.S. Child Welfare System (CWS) in September 2019 which represents a steady increase in the past decade [89]. This number is only expected to grow in upcoming years barring major structural reforms. This has created an ever increasing burden for CWS workers to make decisions about children that provide positive outcomes for them. Policymakers have decided that one avenue to address this issue is to implement algorithmic decision-making within CWS [1]. As such, algorithmic decision-making tools are now being used in high-stakes CWS situations, including making risk assessments of child abuse [34] and determining placement stability [86]. Brown et al. [20] conducted co-design workshops with stakeholders within the CW community (e.g., families, frontline providers, and specialists) and found that such algorithms largely bolstered distrust, perpetuated bias, and created black-boxed systems, which accelerated concerns about how these tools may negatively impact child-welfare workers' decisions. Thus, any technological solution cannot be inherently deemed 'fair' or 'just' and require complementary policy changes to affect community perceptions [20]. In a recent comprehensive review of the literature, Saxena et al. [107] highlighted the lack of human-centeredness [12] in the design and implementation of these algorithms and the need for more empirical work on how these algorithms are embedded in the daily work practices of child-welfare caseworkers. Thus, the gaps identified in these prior works led to the following over-arching research questions:

- **RQ1:** *What are the high-stakes outcomes for which algorithmic decision-making is leveraged within the child-welfare system?*
- **RQ2:** *How does the implementation of a given algorithm impact algorithmic decision-making, human discretion, and bureaucratic processes?*
- **RQ3:** *What are the potential benefits and drawbacks when balancing the trade-offs between these three elements?*

To address these questions, first, we synthesized prior literature to develop a theoretical framework for Algorithmic Decision-Making Adapted for the Public Sector (ADMAPS). While SIGCHI researchers have attempted to formalize the dimensions of algorithmic decision-making in various contexts [2, 62, 91], we argue that the high-stakes decisions being made within the public sector necessitate the critical need for a distinctly unique framework for algorithmic systems that accounts for the complexities of public sector bureaucratic processes [43, 46, 77] and the delicate application of human discretion that has historically been a cornerstone of social services [74, 116]. We did this by synthesizing related works from the fields of Human-Computer Interaction (HCI), Science and Technology Studies (STS), and Public Administration (PA).

Next, we leveraged the ADMAPS framework as a theoretical lens in which to analyze the qualitative data collected from an eight-month in-depth ethnographic case study of a child-welfare agency in the mid-western United States. We attended 55 agency meetings and conducted 20 individual interviews over the course of eight months, which resulted in daily interactions with approximately 120 CWS agency employees and external consultants. To answer (RQ1), we first identified the algorithms used within the CW agency and the relevant data and outcomes they addressed within the system. For (RQ2), we assessed whether and how each algorithm affected each

dimension of **human discretion** (*professional expertise, value judgments, heuristic decision-making*), **bureaucratic processes** (*resources and constraints, administration and training, laws and policies*), and **algorithmic decision-making** (*relevant data, types of decision-support, degree of uncertainty*). Finally, for (RQ3), we compared the four algorithms identified (i.e., CANS, 7ei, AST, and LPS) to show how ADMAPS can help balance the trade-offs in algorithmic decision-making to optimize the benefits and minimize the drawbacks associated with the high-stakes outcomes within the CWS.

Overall, we found that there is a need to refocus on strength-based outcomes centered in social ecological frameworks [19] (RQ1). We define strength-based outcomes as those that draw upon a person's assets and strengths rather than their deficits and weaknesses [9, 131]. In addition, an over-reliance on algorithmic decision-making to support bureaucratic processes can be detrimental to human discretion, however, algorithmic decision-making can support human discretionary work if they are fully supported by bureaucratic processes (RQ2). Finally, algorithmic decision-making should augment human discretion (by building theory-driven algorithms centered in practice) rather than attempt to replace it; algorithm decision-support systems and bureaucratic processes need to be aligned (lack of balance creates utility issues) and collective buy-in in such systems requires trust in algorithmic outcomes at both the caseworker and bureaucratic levels (RQ3). Thus, this paper makes the following unique research contributions:

- We conducted an in-depth ethnographic case study to understand daily algorithmic decision-making practices of caseworkers in CWS.
- We go beyond existing recommendations for AI/ML to provide specific heuristic guidelines for algorithmic decision-making in CWS that can be of use to caseworkers, supervisors, government bureaucrats and policymakers.
- We developed a theoretical framework (ADMAPS) of algorithmic decision-making in the public sector that synthesizes prior work on algorithmic decision-making in non-public sector settings with the unique challenges, limitations and opportunities in the public sector. ADMAPS is generalizable to a wider range of public sector domains such as the criminal justice system, unemployment services, and public education.

In the sections that follow, we first highlight some of the the high-stakes decisions made within the public sector. Then, we introduce our framework for **Algorithmic Decision-Making Adapted for the Public Sector (ADMAPS)**. Finally, we use ADMAPS to present an in-depth ethnographic case study of four algorithms used daily in CWS to determine high-stakes outcomes for foster children, including trauma-informed care, placement stability, and sex-trafficking risk.

2 THE HIGH-STAKES DECISIONS MADE WITHIN THE PUBLIC SECTOR

Algorithmic systems are being used to make high-stakes decisions about human lives and welfare in the public sector ranging across child-welfare, criminal justice, public education, job placement centers, welfare benefits, and housing among others. For example, the criminal justice system employs algorithms to determine length of sentencing [56], allocate resources to neighborhoods [26], and predict the likelihood of recidivism (i.e., recommitting a crime) [39]. In the child-welfare system, decisions are being made about whether to remove a child from the care of their parents based on the risk of future maltreatment [32], who should be raising a child [86], and what types of services should be offered to families [53]. The public education system also uses algorithms to assign students to public school zones [103, 104] and determine student performance [126]. Job placement centers profile job seekers and make job placement decisions using algorithms as well [3, 62]. Algorithms are also used to establish eligibility criteria for receiving benefits and offer these benefits to families in need [42]. In short, many of the ways in which algorithms are being implemented in the public sector result in life altering, if not life and/or death consequences.

The public sector differs uniquely from the private sector in terms of algorithmic decision-making in two distinct ways. First, outcomes in the public sector like assessing risk of recidivism in criminal justice or assessing need in welfare services are poorly and inconsistently defined [11, 42, 55, 107, 129]. Moreover, an individual's personal situation can (de)stabilize several times making it hard to assess what constitutes a successful outcome or intervention [62]. Second, current practices of using aggregate administrative data that is often poorly collected [31, 53, 107] and biased to predict an individual's behavior is a complex and hard task that may lead to unfairness in decision-making outcomes and is, in most western, liberal democratic systems unconstitutional and/or illegal [123]. These two factors combine to make algorithmic decision-making in the public sector, a high stakes decision-making environment that has real repercussions for the lives and liberties of people. Therefore, the algorithmic decision-making process in the public sector needs to be scrutinized with the utmost care. Thus, there is an urgent need to develop a cohesive, yet tailored framework for algorithmic decision-making within the public sector that is validated with in-depth, empirical work focusing on daily algorithmic practices and decision-making.

2.1 Algorithmic Decision-Making within the Child-Welfare System

Algorithms in child-welfare have historically relied on a narrow set of psychometric predictors that are used to assess the risks and needs of foster children and parents. However, a more comprehensive understanding about the accumulation of risk is necessary to account for the family's social support system as well as the risk posed by the system itself [107]. Saxena et al. [107] recently conducted a systematic review of algorithms employed in the U.S. child-welfare system and uncovered several discrepancies in regard to data, computational methods, and target outcomes of these algorithms [107]. There is a need for theoretically constructed algorithms centered in the nature of practice. Moreover, decisions must be made within the constraints of policies and systemic barriers; characteristics not accounted for by algorithms. The majority of algorithms in CWS are empirically constructed even though the empirical knowledge in child-welfare is quite fragmented and social science theories are needed to fill in these gaps [49]. For instance, child-welfare workers are often frustrated by algorithms because they do not account for the scarce resources in the public sector [50]. Risk assessment has also been the dominant focus of algorithms in CWS, however, there are concerns about their deficit-based nature, which only seeks to minimize risk but not improve the quality of children's lives. This is driving attention away from strength-based frameworks [9, 20, 99, 107]. Prior work has also explored the utility of algorithms designed to aid decision-making and found that they increased uncertainty and led to unreliable decisions since caseworkers were required to translate information from both the clinical and algorithmic assessments [111, 113]. More recently, researchers have also focused on the need to uncover politics, economics, and social implications of CWS algorithms and established the need to actively work with practitioners and domain experts to understand their perspectives about such systems as well as the systemic factors centered in policies, laws, and organizational culture that play a significant role in decision-making [20, 54, 101, 109]. This paper responds to these calls by conducting a deep ethnographic analysis of algorithms in use at a CWS agency and uncovers their social, technical, and political implications.

3 A FRAMEWORK OF ALGORITHMIC DECISION-MAKING ADAPTED FOR THE PUBLIC SECTOR (ADMAPS)

As shown in Figure 1, we leveraged a socio-technical perspective of algorithmic decision-making that captures the three-way interactions between: 1) **human discretion**, 2) **bureaucratic processes**, and 3) **algorithmic decision-making**. We did this by synthesizing relevant, yet disparate, bodies of work across the fields of Public Administration (PA), Science and Technologies Studies (STS),

Fig. 1. A Framework for Algorithmic Decision-Making Adapted for the Public Sector (ADMAPS)

and Human-Computer Interaction (HCI) to create a cohesive framework of Algorithmic Decision-Making Adapted for the Public Sector (ADMAPS). This framework is a core contribution of this paper and also served as a theoretical lens for grounding the qualitative analyses of our empirical case study within the domain of child-welfare.

Scholars within the PA field have extensively studied how **human discretion** by street-level bureaucrats¹ plays a central role in navigating **bureaucratic processes** and implementing policies. For example, bureaucrats act with a certain level of autonomy in how they interpret and apply professional standards when determining which clients must receive welfare benefits or services [116]. Public Administration scholars have also started to recognize the impact of information communication technologies (ICTs) on the nature of human discretion and bureaucracy with some recent attention paid to artificial intelligence [21, 23, 80, 96, 128]. Young et al. [128] introduced *artificial discretion* as a theoretical framework to help public managers assess the impact of AI and how it differs from human discretion with respect to task specificity and environmental complexity. However, several of these studies, rich in their understanding of human discretion and bureaucracy, continue to treat algorithms as peripheral end products; a new part of bureaucracy to which human discretion must adapt. Meanwhile, STS scholars conducting studies in the public sector have used Kitchen and Lauriault’s framework of *data assemblages* [66] to deeply study the intersection of **bureaucratic processes** and **algorithmic decision-making** by examining the politics of data systems [71, 101, 126]. Data assemblages perceive data systems as complex assemblages of human actors, artifacts, technical systems, institutions and ideas. This framing provides a means to consider how these systems are socially, economically, and politically constructed. Similarly, STS scholars have also used Seaver’s notion of *algorithms as culture* [112] to understand the social implications and values of algorithmic systems through an ethnographic analysis. These studies have made

¹A street-level bureaucrat is a professional service worker (e.g., social worker, police officer, teacher) who operates in the frontline of public service provision. They interact closely with clients and makes decisions about them based on how they interpret policies relating to the situations at hand [74].

significant contributions towards the community's understanding of how algorithms shape and are shaped by cultural context, how value is inscribed to algorithms, and how power is afforded to them. However, the societal and cultural perspective applied within field of STS may at times miss some of the nuance at the bureaucratic street-level, which is where human discretion is most critical. Due to the high-stakes decisions being made within the public sector, there is an urgent need to map the complex interdependencies between the core elements of human discretion, bureaucratic processes, and algorithmic decision-making that often go unnoticed in the public sector.

The SIGCHI research community is well-positioned to do this cross-disciplinary and integrative work due to our strengths in taking a human-centered [12] and value-sensitive approach [130] to the design and development of algorithms. Furthermore, CSCW is a well-suited venue for this type of research due to the collaborative nature of the work being performed by teams of CWS employees and external consultants when making critical decisions about the well-being of children. In the sections below, we describe the key dimensions of our framework.

3.1 Human Discretion

Lipsky's theory of street-level bureaucrats [74] defines *human discretion* as an individual's ability to exercise their own judgment in implementing government policies in complex and uncertain problem spaces. Scholars across multiple disciplines have recognized the importance of human discretion in developing algorithms and interpreting and administering algorithmic decisions [12, 38, 92–94] but also in interpreting and making policy decisions [22, 27, 46, 74]. However, when HCI scholars deliberate over human discretion, it generally occurs from a design perspective. That is, how can we incorporate humans' tacit knowledge, social interpretations, and values into the design process [12, 130]. Whereas, when Public Administration scholars discuss human discretion, it refers to the decision-making latitude as well as the value-laden choices that bureaucrats must make when experiencing complexity and uncertainty [22, 23, 74]. We integrate this knowledge about human discretion and close this loop by presenting the following process model where the bureaucrats use their professional expertise, engage in value judgments, and heuristic decision-making.

3.1.1 Professional Expertise. The tacit and explicit knowledge developed by workers in any given domain [12, 115]. It plays an important role with respect to the workers' confidence in their own decisions [98] as well as the level of adeptness with which they navigate bureaucratic processes [27, 46], negotiate resources [22], and seek additional supervision [41]. As novice practitioners gain confidence in their skills, they become more aware of the need to pursue additional details, supervision, and other opportunities to increase their level of domain knowledge [124]. Professional expertise, however, is also domain-specific and continually evolves with time [41, 98]. It is necessary to examine the nature professional expertise within the public sector which is rapidly evolving through the continued digitization and automation of work processes that were previously the forte of street-level bureaucrats [23, 51]. Practitioners in the public sector are continually acquiring new skills as they learn to make decisions through data systems and interpret algorithmic outputs [17, 21], however, these new skills are not being accounted for with respect to what constitutes professional expertise [62].

3.1.2 Value Judgments. Practitioners must consider and balance human and democratic values when assessing cases about citizens [47, 62]. This is another key dimension of human discretion because workers must weigh the competing motivations of different clients as well as differing notions of values and incorporate them within their decision-making processes. Value judgments play a pertinent role when a practitioner is faced with ethical dilemmas and informs their decision-making ability. Practitioner's beliefs and moral values are important factors in regard to how street-level decision-making unfolds [77, 83, 105]. Moreover, practitioner's personal values are

often mediated by organizational culture which subsequently yields results that can be significantly different than results centered in personal values [48]. Therefore, it is essential to understand the role of value judgments vis-a-vis human discretion because practitioners have assumed the role of value mediators who must weigh the needs of citizens against the demands of policymakers [100].

3.1.3 Heuristic Decision-Making. Heuristics refer to the cognitive strategies used to form judgments, make decisions, and find solutions to complex problems [52]. Gigerenzer and Gaissmaier [52] reviewed research on heuristic decision-making in business organizations, health care, and legal institutions and established the fundamental role it plays within organizations. Heuristic decision-making can lead to more accurate decisions than complex rational models and selecting information in an adaptive manner can lead to more accurate judgments than weighing all of the information [87]. Practitioners work more effectively and efficiently when their knowledge base is well-organized and centered in heuristics since it allows them to separate relevant and irrelevant knowledge for any given context [5]. Simple heuristics can be more successful especially in uncertain and complex spaces since all the information required to make a decision might not be available as a result of uncertainty [4]. Therefore, it becomes imperative for practitioners to rely on their heuristics that are acquired through experience and practice [124]. Moreover, decision-making in organizations must involve professionals' heuristics because the ideal conditions required for rational and reductive models rarely hold true in an uncertain world [95].

3.2 Bureaucratic Processes

Bureaucratic processes are the critical governance characteristics essential for policy development and implementation to serve public interests [43, 44, 46]. Both HCI and Public Administration scholars have recognized the dominant role that bureaucratic processes play both with respect to establishing the role of bureaucrats (i.e., human discretion) as well as the adoption of technology [17, 22, 76, 109]. Scholars have also emphasized that policy/bureaucratic considerations must precede technology design and professional practice considerations [63]. We include three different dimensions to bureaucratic processes consistently highlighted in the literature as described below.

3.2.1 Resources and Constraints. Availability of essential resources (administrative, financial, personnel, political) directly impact organizational performance [45, 69, 82]. Resources can also be viewed as constraints within which the organization must operate [69]. This dimension is of special importance for the public sector that is facing severely limited resources and new dilemmas in the form of burdensome workloads, high staff turnover, and a lack of experienced workers [82, 85]. Examining how these scarce resources are allocated in public services is crucial because most agencies are experiencing a push to innovate and invest in evidence-based practices to improve performance [122], however, investing in innovation can be challenging in a resource-deficit domain [90].

3.2.2 Administration and Training. Protocols, workflows, and processes established at the organization that are followed by workers in their day-to-day practice and play a significant role in decision-making [69, 74, 97]. Organizational processes offer the means to understand how an agency makes decisions within policy mandates as well as how it meets diverse public needs [16]. Processes are established to improve accountability in the form of consistent, transparent, and defensible decision-making and allow the agency to effectively communicate compliance with legal mandates as well as utilize existing knowledge routines to improve reliability [35, 97]. Moreover, processes followed by practitioners in their daily lives also continually shape policy on the ground [18, 74]. This dimension also identifies the workers' training in the public sector which plays a critical role in regard to individual, team, and organizational development [65]. New assessments and tools are continually being introduced in the public sector in a pursuit for creativity and innovation such

that it leads to standardized and evidence based decision-making [21, 25, 37], however, this also necessitates a need to examine if workers are being adequately trained to fully utilize these tools [120]. Training at the organization also establishes the basis for worker expertise by ensuring that the workers skillfully mediate both the nature of practice and bureaucratic processes [124].

3.2.3 Law and Policies. Formal actions enacted by legislatures or political executives that public administrators must comply with and implement [96]. Laws and policies establish the constraints within which all decisions (human or algorithmic) must be made as well as define the outcomes of interests themselves. For instance, law dictates which data is available for predictive modeling and how target outcomes are defined [107, 114]. This dimension is of critical importance since it directly impacts both human discretion and algorithmic decision-making. For instance, the policy decision to expand mandated reporting in child-welfare significantly increased the number of cases referred to CWS as well as broadened the definitions of child abuse and neglect (with implications for algorithmic modeling) [84]. Prior work has acknowledged the dominant role that bureaucracy or policy plays in the public sector both in regard to decision-making as well as the adoption of technology [63, 76]. The central role of bureaucratic processes is evident from prior work conducted in public services where caseworkers pushed for algorithmic systems that could help mitigate organizational contradictions and clarify organizational processes [62].

3.3 Algorithmic Decision-Making

Algorithmic Decision-Making is defined through the lens of *street-level algorithms*, a term recently coined by Alkhatib and Bernstein [2] in the HCI community. Street-level algorithms directly interact with and make on-the-ground decisions about human lives and welfare in a sociotechnical system [2]. Prior work has argued that algorithms in the public sector is a domain in its own right [34, 62, 107, 109, 121] and must be characterized differently as compared to algorithms in the private sector where the decision outcomes are well-defined. Therefore, it becomes important to examine and critique the dimensions within Algorithmic Decision-Making that impact the predicted outcomes. Algorithmic Decision-Making is the most flexible element of the framework that designers can directly impact. That is, algorithms must be developed in such a way that they balance the other two elements (human discretion and bureaucratic processes). HCI methodologies such as value-sensitive algorithm design [130] and human-centered algorithm design [12] can ensure that the algorithms account for values of stakeholders as well as theory-driven practice. Moreover, participatory design can unravel the policy mandates and institutional processes that often mediate the decision-making process and must be accounted for [76, 109].

3.3.1 Relevant Data. Necessary information about individuals and their environment must be collected to be able to adequately predict an outcome of interest. In several domains within the public sector, there is significant debate about which predictors are associated with which outcomes [39, 62, 106]. Moreover, the necessary information may not always be available or inconsistently available with contradicting factors [31]. For instance, risk assessment algorithms in child-welfare have traditionally only used a narrow set of predictors (child and parent characteristics) to assess risk [107]. However, a more comprehensive understanding of risk is necessary, including the risk posed by the system itself [49]. Therefore, algorithms need to be theoretically constructed with proper considerations from domain experts with respect to feature selection and modeling to ensure that the algorithm offers higher utility and complements theory of practice [12, 109, 130].

3.3.2 Types of Decision-Support. Two types of algorithms are predominantly used in the public sector; predictive and prescriptive algorithms. Predictive algorithms seek to predict the likelihood of the occurrence of an outcome of interest, whereas, prescriptive algorithms act as decision aids and

offer recommendations to intervene and affect the outcome of interest [36]. Examining the nature of decision-support systems is equally as important as interrogating the outcome itself. Prescriptive decision-aids are often introduced as a means to improve decision-making while not shifting agency away from workers. However, prior research shows that workers allow decision-aids to supplant their own decisions when they lack confidence and/or experience [42, 107, 113]. Moreover, the calls for human-in-the-loop might be moot if there is a lack of understanding about how algorithms impact human decision-making and how the type of decision-support (i.e.- algorithm design) impacts the practical possibilities for human intervention [96, 111, 113].

3.3.3 Degrees of Uncertainty. Decision outcomes in the public sector are not well-defined and as previously noted, a person's life can stabilize or de-stabilize several times making it hard to predict what constitutes a successful outcome or intervention [62]. Prior research has also established that an irreducible degree of uncertainty exists with respect to the outcomes in the public sector [21, 34, 57] with both humans and algorithms likely to make mistakes. Pääkkönen et al. [91] further extend this argument to state that the design of algorithmic systems must identify and cultivate important sources of uncertainty because it is at these sources where the need for human discretion accumulates since ambiguity about the operation of the algorithm persists.

This framework challenges designers, practitioners, and policymakers to rethink the core assumptions and nature of their practice which are evolving in an increasingly socio-technical public sector and need to be re-examined in light of these new challenges and opportunities. It provides a structured way to think about socio-technical problems centered in algorithmic decision-making in the public sector, study the interdependencies between the dimensions, and recognize underlying causes that impact decision-making.

4 METHODS

In this section, we describe our partnership with a child-welfare agency to address the research questions set forth in the introduction of our paper.

4.1 Study Overview

The goal of this study was to examine the algorithms that caseworkers use in their daily work lives and unpack the collaborative nature of how these algorithms were used in group settings. To accomplish this, we partnered with a child-welfare agency, which serves about 900 families and 1300 children in a large metropolitan area in the midwestern United States. The state's Department of Children and Families (DCF) contracts child-welfare services to this agency who must comply with all DCF standards including the use of mandated algorithms or decision-tools. We conducted an eight-month long in-depth ethnographic case study at the agency from August 2019 to March 2020. Before conducting observations or recruiting participants for interviews, we obtained Institutional Review Board (IRB) approval at our mid-sized private research university to conduct our study. We then emailed the participants an IRB approved consent form and obtained their verbal consent to participate in the study. During this time, the first-author observed child-welfare team meetings and conducted semi-structured interviews with key stakeholders at the agency.

4.2 Meeting Observations and Interviews

The first author conducted in-person observations of meetings to gain the necessary understanding of how algorithms were used in a team setting and how caseworkers interacted with these algorithms in their daily work practices. These observations were also helpful in understanding the collaborative work of child-welfare teams that make decisions that are mediated by policies, social-work practice, and algorithms. Next, we provide a detailed description of the team meetings and interviews.

4.2.1 45-Day Staff Meetings or Planning Meetings. The 45-day staff meetings occur within the first 45 days of a case coming into the care of the agency and are attended by child-welfare team members involved at the front-end of case planning [106]. These meetings facilitate information sharing so that consensus can be reached in regard to the child's well-being and placement stability². Each meeting is scheduled for 90 minutes, and we observed 15 meetings. These meetings are typically attended by the CWS employees that work in case management, permanency planning, family preservation, and licensing. Central to these meetings is the 7ei staffing protocol that helps the child-welfare team apply principles and practices derived from trauma-informed care (TIC) [60] to each case. The 7ei Staffing protocol is accompanied by the 7ei algorithm which acts as a framework for TIC and helps track progress with respect to each case [119]. The CANS algorithm is also used at these meetings to establish a baseline for foster children with respect to mental health well-being. We also identified two more algorithms being used by the child-welfare teams. *Legal Permanency Status (LPS)* algorithm is used to assess the quality of the current placement and recognize systemic barriers. *Anti Sex-Trafficking (AST)* algorithm is used to assess the risk of sex-trafficking for foster youth. Observing these meetings allowed us to understand how these decision-making tools were being used in practice, the benefits they offered, as well as the challenges they posed.

4.2.2 Permanency Consultation Meetings. Permanency consultation meetings are specialized meetings designed to expedite permanency³ for children placed in out-of-home care by employing innovative best practices and seeking to address any systemic or policy-related barriers. These meetings are facilitated by permanency consultants and are staffed with many of the child-welfare team members that attend the planning meetings. They regularly occur at the 5, 10, and 15+ month marks for every case until the case is closed. These on-going meetings tended to be more involved than the planning meetings as limited information is available at the onset of a case. Moreover, permanency consultations involved cases that had been with the agency for several months (if not years) and revealed the messy interaction between the complex socio-political domain of child-welfare and the algorithmic tools being used. For instance, it was interesting to observe how the child-welfare teams reached consensus during decision-making discussions when they had to account for policy and systemic barriers, social-work practice, and the algorithms. Each meeting was scheduled for an hour, and we observed 40 of these meetings.

4.2.3 Semi-Structured Interviews. Next, we used the knowledge gathered from these observations to develop our interview protocol and recruit participants who consistently attended these meetings as part of their job routines. After having first observed the child-welfare teams interact with algorithms for several months, we conducted interviews to delve deeper into the participants' understanding of these decision-tools as well as the benefits and challenges as perceived by them. We asked participants a series of questions about the nature of child-welfare work and the algorithms we observed being used at the agency. We also asked them to expand upon any interactions we had observed during the meetings. For example, we asked them to share their appreciation or dissatisfaction towards a certain algorithm or feature, as well as their team's or self frustration with the misuse of these decision tools. We conducted 20 semi-structured interviews with child-welfare staff members, which included permanency consultants, supervisors, program directors, ongoing case managers, data specialists, and clinical therapists. Seventeen interviews were conducted at the child-welfare agency in the participants' private offices or conference rooms. Due to the COVID outbreak, the last three interviews were conducted over the phone to ensure the safety of both the participants and the researcher. Interviews lasted for about 45 minutes to 1 hour in duration.

²Placement stability is defined as three or fewer placement moves for a foster child during the previous 36 months.

³Permanency is defined as reunification with biological parents, adoption or legal guardianship.

4.3 Qualitative Data Analysis

The first author took detailed observational notes during each team meeting and compiled a debriefing document with his initial insights within 30 minutes of each meeting to retain as much of his thoughts as possible. The interviews were audio-recorded and transcribed verbatim for analysis. Notes, documents, and transcripts were shared among all co-authors. Our high-level research questions guided our analyses, but within those questions, we allowed for new insights to emerge and adjusted our research questions based on emergent insights. We performed a thematic qualitative analyses [30] to answer RQ1 and RQ3. The first author read through the interview transcripts several times and created initial codes and consulted with co-authors to form a consensus around the codes, as well as resolve any ambiguous codes. Next, these codes were conceptually grouped into themes. However, in our results, we also use our observational notes to augment the insights we gained from the interviews and note potential discrepancies and nuances from the holistic insights gained from our site observations.

For RQ1, we used an open-coding process to identify the high-stakes decision outcomes associated with the four algorithms that are embedded in child-welfare practice, namely, Child and Adolescent Needs and Strengths (CANS), Seven Essential Ingredients (7ei), Anti Sex-Trafficking Response Tool (AST), and Legal Permanency Status Tool (LPS). In understanding how these algorithms were used in the daily practices of the CWS employees, we identified seven key purposes (As shown in Table 2): 1) **Compensation Calculations**: determine the monetary value to be offered to foster parents for caring for foster children, 2) **Mental Health Assessment**: conduct a mental-health screen of foster children to assess the risks and needs, 3) **Level of Foster Care**: determine a suitable placement setting capable of meeting the needs of children, 4) **Trauma-informed Care**: a trauma-responsive service model developed through an ecological understanding of adverse events and trauma experienced by children and families, 5) **Placement Stability and Permanency**: Track outcomes from the trauma-responsive service model and assess if they are leading to better outcomes, 6) **Sex Trafficking Risk Assessment**: assess the risk of sex-trafficking for a foster child, and 7) **Quality of Placements and Systemic Barriers**: Track the current quality of placement and the systemic barriers inhibiting permanency.

For RQ2, we used the ADMAPS framework to code for how each algorithm (i.e., CANS, 7ei, AST, LPS) impacted (positively, negatively, or both) dimensions of ADMAPS. For example, we found that CANS had an overall negative impact on human discretion with 80% of the interviewees indicating that it negatively impacted professional expertise by limiting the scope for value judgments and heuristic decision-making; 75% of interviewees saying that it reduced their ability to make flexible value judgments on child outcomes; and 80% of interviewees asserting that they were no longer given the discretion to make decisions on behalf of the children assigned to them because CANS made several of these decisions for them. These mappings to the ADMAPS framework allowed us to assess the role that human discretion, algorithmic decision-making, and bureaucratic processes played with respect to each of the four algorithms deployed in CWS daily practice, as well as compare the differences between them.

For RQ3, we synthesized the patterns across the four algorithms to identify emergent themes that were consistent across our analysis of RQ2 to give a bigger picture of the potential benefits and drawbacks associated when balancing the tradeoffs between human discretion, algorithmic decision-making, and bureaucratic processes.

5 RESULTS

In the following sections, we organize and present the results by our three research questions. First, we identify the high-stakes outcomes for which algorithmic decision-making is leveraged

within CWS (RQ1) and the roles that human discretion, bureaucratic processes, and algorithmic decision-making play in these decision outcomes (RQ2). We do this separately for CANS, 7ei, AST, and LPS. Next, we discuss the potential benefits and drawbacks when balancing the different dimensions of ADMAPS framework (RQ3). Interviewees profiles can be found in the appendix.

5.1 Child and Adolescent Needs and Strengths (CANS) Algorithm

The Child and Adolescent Needs and Strengths (CANS) algorithm was used at all the planning meetings (n=15) and discussed by all the interviewees (n=20). CANS algorithm is constructed using the CANS assessment; a communimetric tool designed to assess the level of need of a foster child and develop an individualized service plan [79]. With its primary purpose being communication, CANS assessment is designed based on communication theory rather than psychometric theories centered in measurement development [78]. CANS assessment was designed to support decision-making with respect to assessing a child's level of need and service planning. However, as depicted in Figure 2, CANS algorithm has been re-appropriated to measure additional outcomes discussed below. It is conducted within the first 30 days of a child entering the child-welfare system or moving to a new placement (for e.g., foster home). It is then periodically conducted every six months.

5.1.1 CANS High-Stakes Decision Outcomes (RQ1).

Mental Health Assessment and Level of Foster Care. All child-welfare workers at the agency are certified in conducting CANS for mental health screenings. As such, several participants (60%, n=12) shared that the agency uses CANS algorithm to conduct a mental health assessment of foster children based on risks posed and exhibited behaviors. Participants (60%, n=12) stated that the first CANS assessment was most useful because it helped establish a baseline for the mental-health needs of a child and the level of care that the child needed. For instance, one supervisor explained:

"We need to have starting point.. what kind of behavioral issues does a child have? What are their needs? Because we can't naively place a high-needs kid with foster parents who are not trained and certified to manage those needs. It's a recipe for disaster" -P9, Child Welfare Supervisor, MSW, 13 years

The first assessment is used to devise service plans for foster children (for e.g, behavioral therapy) based on exhibited behaviors. It also facilitates sharing this information with other parties such as legal parties and family preservation specialists who also play a role in case planning. Child-welfare teams in both the planning meetings (n=15) and permanency consultations (n=40) briefly discussed child needs and behaviors but then shifted to a more extensive conversation about trauma using trauma-informed care. Some participants (50%, n=10) also shared that CANS algorithm recommends the level of foster care that the child should be placed in (see Figure 2). Higher level foster homes are trained and certified in taking care of high-needs children. However, due to a lack of such homes, this decision often comes down to the availability of resources.

Compensation Calculations. Using the CANS algorithm for calculating foster parents' compensation was another prominent theme that emerged in 73% (n=11) of the planning meetings and 85% (n=17) of the interviews. However, it was not a dominant theme at the permanency consultations, because compensation is directly negotiated between the foster parents, case managers, and the supervisor. It does not require the input of rest of the child-welfare team. The state reimburses foster parents for the costs associated with having foster children placed in their homes. Most of the participants (85%, n=17) quickly recognized CANS as the "rate-setting tool" even though compensation calculation was not the primary purpose of the algorithm. One supervisor explained:

"Foster parents who are taking in high-needs kids should be compensated for that. They have to put in significantly more time and energy in managing those behaviors, taking kids to therapy,

Fig. 2. CANS Algorithm: Associated Decision Outcomes

setting healthy boundaries.. So we need this [CANS] standard to do that" - P12, Child Welfare Supervisor, MSW, 7 years

The state needs a *metric* to be able to determine the compensation to be offered to each foster parent. The Department of Children and Families (DCF) decided to associate this compensation with the mental health needs of a child, that is, the higher the mental needs of a child, the higher the compensation offered to foster parents. Thus, even though the primary purpose of CANS was for conducting mental health assessments and for the level of foster care, it was re-appropriated to also calculate compensation associated with caring for a foster child's mental health needs. While it logically follows that a child with more mental health issues will require a higher level of care (consequently, higher cost of care), CWS employees were also well-aware that such cost-benefit analyses tied directly to something as subjective as mental health assessments were problematic.

5.1.2 CANS and Algorithmic Decision-Making (RQ2). Overall, caseworkers were frustrated that CANS misses important context about the child but is still used in a mandated predictive capacity. In the sections below, we discuss how CANS algorithm maps onto the dimensions of algorithmic decision-making (i.e., relevant data, type of decision-support, and degree of uncertainty). We use a percentage combined with an up or down arrow to denote the percentage of participants who indicated a positive or negative impact on each dimension of the ADMAPS framework for algorithmic decision-making, human discretion, and bureaucratic processes, respectively. We follow this structure throughout the remainder of our results.

70% ↓ Relevant Data. *CANS data does not account for trauma or social interactions.* Most of the participants (70%, n=14) shared that CANS conducted the child's assessment in an isolated manner and did not account for the quality and impact of relationships in their lives which are often more important for determining the long-term well being of these children. CANS algorithm focuses on the child emotional/behavioral needs (for e.g., anxiety, anger control, substance use, behavioral regression) and child risk behaviors (for e.g., suicide risk, self-harm, delinquent behavior, runaway tendencies) to assess the mental health needs of the child. For instance, a supervisor shared –

"How do you measure empathy of others? You cant. Some foster parents are more empathetic and understanding and stand by the child no matter what. Sometimes that's all it takes. You can't put that in CANS" -P15, Child Welfare Supervisor, 9 years

Moreover, participants noted that CANS focuses on current behaviors but not the underlying trauma or traumatic triggers. The assessment is conducted based on exhibited behaviors over the

past 30 days, however, the participants explained that trauma can stay with a child for years and lead to serious emotional dysregulation from time to time.

90% ↓ Type of Decision-Support. *The predictive nature of CANS leaves no room for discretion.* This challenge emerged in 90% (n=18) of the interviews and 73% (n=11) of the planning meetings. CANS algorithm is designed to predict or measure outcomes of interest once the data has been provided. Participants shared that the predictive nature of CANS leaves little room to exercise discretion and has become a great source of frustration for them. Gaming the inputs to achieve the desired outcome is the only way through which caseworkers regain agency. Participants explained that the higher the mental health needs of a child per CANS, higher the compensation offered to foster parents. After CANS scores are entered in, the algorithm generates a monetary value to be offered to foster parents. However, if foster parents disagree with the rate and believe they should be paid more, the caseworkers manipulate the scores to produce a higher rate. Most of the participants (90%, n=18) shared their frustration regarding the inflexible and predictive nature of CANS. For instance, one supervisor explained:

"Case managers and even supervisors are being forced to.. and kind of pressured into scoring children higher in order to provide higher numbers. So foster parents get paid more. CANS is a manipulative tool being used to barter off children...Children are being exploited for payment"
-P10, Child Welfare Supervisor, MS, LPC, NCC, 9 years

70% ↓ Degree of Uncertainty. *High degree of uncertainty associated with the outcomes.* Several participants (70%, n=14) shared that CANS does not account for much of the data that they consider pertinent when assessing cases (for e.g., understanding about trauma and social support-system). This lack of relevant data leads to a high degree of uncertainty which is further exacerbated by the predictive and inflexible nature of CANS. Many felt that gaming the algorithm was the only viable option for caseworkers to exercise discretion. For instance, one supervisor shared:

"CANS has become all about the rate. Generating the right rate so foster parents are happy with little to no attention paid to mental health needs" -P13, Child Welfare Supervisor, MSW

5.1.3 CANS Severely Impedes Human Discretion (RQ2). In this section, we discuss how CANS algorithm maps onto the ADMAPS dimensions of human discretion.

80% ↓ Professional Expertise. *CANS contradicts professional expertise.* All caseworkers at the agency are required to pass the CANS certification to be able to conduct the assessment with clients. However, most participants (80% n=18) shared that their training in trauma-informed care (see Section 6.3), which offers a more comprehensive understanding of trauma and the child's environment often conflicted with CANS. CANS further inhibited professional expertise because caseworkers felt that it turned them into data brokers who must collect information about children and feed it to the algorithm to make the decisions. The lack of relevant data and lack of decision-making latitude on part of caseworkers has turned CANS into primarily a rate-setting tool to calculate compensations. For instance, one case manager shared:

"Caseworkers are doing CANS just to get it done.. to produce a good rate and reduce the conflict with foster parents. There is little to no attention paid to the mental health needs of kids." -P17, Ongoing Case Manager, 8 years

75% ↓ Value Judgments. *CANS has introduced conflicting values.* Most of the participants (75%, n=17) shared how the re-appropriation of CANS to calculate the foster parent compensations has led to several unintended consequences. CANS is re-conducted every six months to reassess the mental health needs of children, and consequentially, compensation is recalculated. However, with

a focus on exhibited behaviors and not the underlying trauma, foster parents who are helping children cope and recover can see their compensation being lowered. A supervisor explained:

"It is the complete opposite of what we want it [CANS] to do. Foster parents help minimize the behaviors and offer support so that kids can develop good coping skills. They help address the mental health needs and help kids stabilize by taking them to therapy and all their activities. But then they're punished because the kid's needs go down, and so does the rate" -P9, Child Welfare Supervisor, MSW, 13 years

Here, caseworkers are unable to prioritize properly conducting CANS because generating the adequate rate takes precedence to ensure the placement is not disrupted. Placement disruptions adversely affect foster children who develop emotional and behavioral problems and are unable to form lasting meaningful relationships with foster parents [15]. Therefore, caseworkers must prioritize supporting the current placement by any means necessary. This contradictory nature of CANS has turned caseworkers into value mediators and has left them frustrated because they are unable to adequately balance the needs of families and the demands of policymakers.

80% ↓ Heuristic Decision-Making. *CANS leaves no room for heuristic judgment calls.* This theme emerged in 80% (n=16) of the interviews and 73% (n=11) of the planning meetings. With a lack of relevant information pertinent to a case and high degree of uncertainty, it becomes imperative that caseworkers are able to turn towards heuristics and make decisions with the assistance of the child-welfare team. However, CANS does not allow for heuristic decision-making with respect to the outcomes. Participants (80%, n=16) shared that every case is contextually different with salient factors that might be central to one case but peripheral to another. This was also apparent in all the planning meetings and permanency consultations where the teams adaptively focused on information pertinent to that case.

5.1.4 CANS and Bureaucratic Processes (RQ2). Overall, CANS allocates resources but does not account for organizational constraints. In this section, we discuss how CANS algorithm maps onto the ADMAPS dimensions of bureaucratic processes.

40% ↑ 70% ↓ Resources & Constraints. *CANS has introduced new constraints in the case planning process.* There is both a benefit and drawback to how CANS algorithm accounts for resources and constraints. Some participants (40%, n=8) found value in the first assessment since it established the baseline for mental health and compensation for foster parents. Ideally, the algorithm offers an efficient way to allocate funds to foster parents based on mental health needs of the child. However, it does not account for organizational constraints and its implementation ends up introducing more constraints that frustrate caseworkers. Participants (70%, n=14) shared that properly conducting CANS requires the caseworkers to interview several people, however, they are only able to interview foster parents due to high caseloads. A permanency consultant asserted:

"You are supposed to interview foster parents and teachers as well as others that kids interact with to get a good CANS assessment. But with high caseloads, caseworkers only talk to the foster parents" -P4, Permanency Consultant, MSW, APSW, 8 years

However, as previously discussed, CANS is conducted every six months and foster parents are motivated to exaggerate behaviors to continue to be paid consistently. This is a constraint introduced by the implementation of the algorithm itself. Participants (50%, n=10) also shared that the algorithm recommends the level of foster care that a child should be placed in, however, there is a lack of good foster homes in the system and this requirement is seldom met. Moreover, two data specialists shared that the caseworkers' yearly job performance is tied to the timeliness with which they complete and submit CANS assessments since funds need to be allocated in a timely

manner. Therefore, it exacerbates the need to get the assessments completed irrespective of the mental health needs of a child, and consequentially, adds another organizational constraint.

40% 80% **Administration & Training.** *Caseworkers are trained to conduct CANS but not on managing constraints.* This theme emerged in 80% (n=16) of the interviews. CANS offers a way to support an organizational process of allocating resources. Caseworkers are trained and certified to conduct the mental health assessment, and as previously noted, some participants (40%, n=8) found value in the first CANS assessment as a means to establish a baseline for mental health needs. However, they are not trained on how to manage the systemic constraints and value conflicts introduced by the algorithm itself. Participants (50%, n=10) shared that new caseworkers socially learn from other caseworkers (and through experience) to manage these conflicts by gaming CANS.

80% **Laws & Policies.** *The use of CANS is legally mandated.* Most participants (80%, n=18) recognized (and often complained) that CANS was legally mandated by the state and offered them a convenient means to allocate resources to foster parents every six months. Participants recognized that they must comply with this policy and continue implementing CANS enough though they routinely manipulated it to generate higher compensations.

5.2 Seven Essential Ingredients (7ei) Algorithm

The agency uses trauma-informed care [60] as their core guiding principle, which is embedded in their practice and trauma-responsive service model [119]. 7ei is a theoretically constructed algorithm that is centered in the medical practice of trauma-informed care (TIC) and complements the agency's social work practice [60, 64]. We witnessed 7ei being utilized in all meetings (N=55). This algorithm complements specialized trainings developed by the agency for child-welfare workers and introduces the complexity of trauma, frameworks for understanding the effects of trauma, and the practices and principles of TIC [119]. 7ei acts as a guiding framework and is used to track and score each case from a TIC perspective. The algorithm is designed to be used in a team setting such that team members can offer their expertise, reach consensus decisions, and devise case plans. The team discusses and scores the child's and caregiver's wellness on seven domains as depicted in Figure 3 and brainstorms solutions on how to make progress on these domains. 7ei algorithm based in TIC has also been proven to improve child outcomes such as placement stability and permanence [119]. 7ei assessments are individually focused, however, unlike CANS, the results are trauma-focused and also guide family interventions, highlighting areas of child functioning upon which the caregivers and professionals should focus their attention [119].

5.2.1 7ei High-Stakes Decision Outcomes (RQ1). The seven domains of 7ei are: *Prevalence, Impact, Perspective Shift, Regulation, Relationships, Reasons to Be, and Caregiver Capacity*. As such, 7ei is not directly tied to a specific outcome such that every time 7ei is used an outcome is recommended or predicted. Instead of predicting an outcome of interest using other factors, 7ei is used to track outcomes (i.e.- 7ei domains) over time to assess the trajectory of a child-welfare case. It is primarily used as a prescriptive tool that serves as a framework for team-based brainstorming of solutions guided by a TIC framework. As such, the seven domains of 7ei are both the input variables and the output variables of the 7ei assessment process because these are both the considerations and outcomes that the team is trying to measure and improve over time. Improvement in the 7ei domains is associated with placement stability and permanency outcomes, however, the agency leadership resisted developing a singular aggregate index that would measure this outcome. Ongoing conversations with agency leadership (program directors, quality improvement leaders) revealed that assessments are more likely to be manipulated if they are tied to singular metrics. Agency leadership also uses 7ei to assess progress within the agency with respect to trauma-informed care. For instance, it helps them understand whether conversations and meetings founded in TIC are

Fig. 3. 7ei Algorithm: Associated Decision Outcomes

leading to better permanency outcomes. Case-level outcomes of 7ei that are discussed at all the planning meetings and permanency consultations are depicted in Figure 3.

5.2.2 7ei and Algorithmic Decision-Making (RQ2). We found that the prescriptive nature of 7ei helps adaptively select data and address uncertainties. In this section, we discuss how 7ei algorithm maps onto the ADMAPS dimensions of algorithmic decision-making.

90% ↑ Relevant Data. 7ei provides a comprehensive view of the child and their ecosystem. Participants shared that 7ei algorithm offers a comprehensive view of the foster child, caregivers, impact of traumatic events in their life, as well as interactions in their social ecosystem. Moreover, the child-welfare teams in all the permanency consultations (n=40) and planning meetings (n=15) were able to adaptively select factors that were most pertinent to that case. Most participants (90%, n=18) emphasized that each case carried a lot of nuance and could not be addressed based on a few broad set of predictors so they appreciated that 7ei allowed them to focus on certain factors and then brainstorm ideas on how to help the family make progress. For instance, a program director shared:

“We have tried the cookie cutter approach in the past. Assigning everyone to parenting classes, therapy, and other family support services. It failed and it is horrible to do to a family. So, with 7ei we focus on addressing core issues whether it's the parent's self-esteem, their own abandonment issues or child's emotional regulation that will really help this family” -P7, Child Welfare Program Director, MSW, 20 years

Timeline of the case established which 7ei domains the child-welfare teams converged on. The teams focused more on some domains than others based on how long the child had been in the system. For instance, in case of children who had been in the system for a few months, the team spent more time on *Prevalence* and *Impact* and then focused on *Perspective Shift* (see Figure 3). Trauma symptoms are associated with negative short-term outcomes such as placement instability [29], therefore, the team focuses on recognizing trauma early on so that proper interventions could be made that promote healing and improving outcomes. However, cases that had been in the system for a longer period or had experienced multiple placement moves, the team focused more on *Regulation*, *Perspective Shift*, and *Reasons to Be*. Prior studies show that placements are often disrupted because foster parents are unprepared to manage behaviors of traumatized children

[24, 118], therefore, the team focuses on these domains to assess how to improve self-regulation for the foster child as well as expand the caregivers' understanding of trauma.

85% ↑ Type of Decision-Support. *The prescriptive nature of 7ei allows for brainstorming and idea generation.* Most participants (85%, n=17) appreciated that 7ei allowed for open discussions and brainstorming of solutions. 7ei is designed to be used in a team setting and acts as the TIC framework such that the child-welfare team deliberates over each domain, scores it, and formulates an action plan through a trauma-informed perspective. A case manager shared:

"I like 7ei because we use it as a team, and it allows us to brainstorm ideas. We look at not just the behaviors but the underlying trauma that is causing those behaviors." -P18, Case Manager, BSW, 2 years

The 7ei domains are also the outcomes of interest that the team seeks to track and affect over time. Improvements in these domains are directly associated with the outcomes of placement stability and permanency which are tracked at the agency level.

80% ↑ Degree of Uncertainty. *Adaptively selecting relevant data helps address high degree of uncertainty.* Most participants (80%, n=16) felt that every case was contextually different and required their individual attention. That is, it was imperative they adequately weigh nuances and factors pertinent to that case. 7ei allows the child-welfare team to adaptively select the domains pertinent to a case and spend significantly more time on them and brainstorm solutions. Tracking outcomes over time and not using the tool in a predictive capacity allows the team to mitigate the high degree of uncertainty that would otherwise be associated with the predictions.

5.2.3 7ei Augments Human Discretion (RQ2). In this section, we discuss how the 7ei algorithm maps onto the ADMAPS dimensions of human discretion.

85% ↑ Professional Expertise. *7ei helps develop professional expertise.* The algorithm is theoretically constructed and is centered in trauma-informed care. Participants (85%, n=17) believed that it allowed them to brainstorm ideas from a TIC perspective and develop plans specific to that family. Continuous engagement with 7ei in trauma-informed meetings ensures that caseworkers are always thinking through TIC frameworks. A permanency consultant asserted:

"7ei isn't just a 'thing' that we do. It is centered in everything that we do. It ensures that caseworkers are always thinking through TIC frameworks" -P1, Permanency Consultation Supervisor, MSW, APSW, 22 years

This is especially important for CWS since the system lacks experienced caseworkers due to high turnover [24]. Continually working through 7ei under proper supervision ensures that new caseworkers are developing professional expertise.

70% ↑ Value Judgments. *7ei allows caseworkers to make informed value judgments.* This theme emerged in 70% (n=14) of the interviews and 72% (n=40) of the meetings. 7ei is centered in social work's core values of service, dignity, and worth of the person and allows caseworkers to prioritize these values and devise interventions that will directly help a child and their family cope with trauma. For instance, prioritizing the well-being of a child does not only mean sending them to therapy. It also incorporates addressing concerns within their ecosystem. A supervisor explained:

"Therapy only goes so far if nothing changes in the child's ecosystem and they feel continually triggered by others. With 7ei we try to address problems in this ecosystem and devise approaches that will help the family" -P9, Child-Welfare Supervisor, 13 years

7ei allows the child-welfare team to make value-based judgments and take steps that improve the quality of a child's relationships. This takes the form of family-level interventions or sharing information with caregivers about the impact of trauma to bring about a perspective shift.

85% **Heuristic Decision-Making.** *7ei enables heuristic judgment calls.* Most participants (85%, n=17) appreciated that 7ei offered them flexibility and autonomy in how they interact with it. This theme also emerged in all the planning meeting (n=15) and permanency consultations (n=40) where the child-welfare team adaptively selected information that was most pertinent to the case and often acted as an obstacle towards achieving permanency. For instance, cases where the child had experienced multiple placement moves, the team focused on *Regulation* and *Reasons to Be* to devise plans that would help improve the child's emotional, behavioral, and cognitive functioning.

5.2.4 7ei Supports Bureaucratic Processes (RQ2). In this section, we discuss how the 7ei algorithm maps onto the ADMAPS dimensions of bureaucratic processes.

60% **Resources & Constraints.** *7ei accounts for the resources at the organization.* Participants (60%, n=12) shared that 7ei is locally developed at the agency and accounts for the resources available at the agency in the form of supervision, expertise, and specialized trainings. Participants also shared that critical decision-making power in regard to achieving permanency sits with the legal parties (district attorneys, judges), however, 7ei operates within these constraints and allows caseworkers to help families and prepare them to be able to receive a favorable decision in court.

75% 55% **Administration & Training.** *7ei is embedded in daily work processes but requires additional oversight and training.* This theme emerged in 75 % (n=15) of the interviews and all the planning meetings (n=15) and permanency consultations (n=40) and there is both a benefit and a drawback associated with it. Participants (75%, n=15) recognized that continually engaging with the tool in meetings was helpful but some participants (55%, n=11) shared that the tool also added more tasks to an already heavy workload. The agency offers specialized trainings for trauma-informed care with 7ei acting as a tool that complements these trainings. Moreover, we observed that 7ei is embedded into daily work processes because it is utilized every time a case is discussed in a team setting. This continued engagement under supervision has helped earn the trust of caseworkers who learn not just from the algorithm but also the collective expertise of the child-welfare team. For instance, a case manager shared:

"I like 7ei because we use it as a team, and it allows us to brainstorm ideas. It just helps.. thinking out loud with everyone and knowing that I don't have make these decisions alone. Also, it helps guide my thought process but doesn't tell me what to do." -P17, Case Manager, BSW, 8 years

However, 7ei has added to the workload of both the case managers and supervisors who must discuss and complete an additional tool as part of their job requirement at the agency.

65% **Laws & Policies.** *7ei is not legally mandated and only used locally at the agency.* Several participants (65%, n=13) were frustrated by the fact that 7ei was not legally mandated and that they had to continue using CANS. Even though an independent research study showed that 7ei is leading to better permanency and placement outcomes, the algorithm still lacks legitimacy at the state and federal level.

5.3 Anti Sex-Trafficking (AST) Algorithm

An emerging high-stakes decision for which an algorithm is used is assessing the risk of sex-trafficking for a foster child over 10 years of age. We observed this in 65% (n=13) of the interviews, 33% (n=5) of the planning meetings, and 35% (n=14) of the permanency consultations. The agency has a dedicated team called HART (Human Anti-Trafficking Response Team) that manages cases where the foster child might be at high risk of being trafficked. If a child meets the criteria for risk of sex-trafficking per the AST algorithm, the case must be reported to HART. Early identification of such indicators can play a significant role in ensuring child safety.

5.3.1 AST High-Stakes Decision Outcomes (RQ1). Risk indicators are divided into three domains - "At-Risk", "High-Risk", and "Confirmed". The child-welfare team must continue to closely monitor the case if they select fewer than three "At-Risk" indicators and continue to have conversations to mitigate the risks. If the team selects three or more "At-Risk" indicators or 1 or more "High-Risk" or "Confirmed" indicators, then the case must be referred to HART.

5.3.2 AST and Algorithmic Decision-Making (RQ2). Overall, we found that AST algorithm often missed important context, which frustrated caseworkers. In this section, we discuss how AST algorithm maps onto the ADMAPS dimensions of algorithmic decision-making.

65% ↑ 75% ↓ Relevant Data. *AST assesses pertinent risk indicators but misses context about the case.* Several participants (65%, n=13) shared that the algorithm offers a new perspective on sex-trafficking by capturing risk indicators that they had not considered before. For instance, "possession of money, electronics, cosmetics, or clothes that are unexplained", "traveling out of the area or somewhere out of the ordinary", and "unwilling to provide information about an older partner" are some of these indicators. A supervisor shared:

"It's helpful to have us think about it differently. I have been doing this for a very long time and when I started, we were never thinking about trafficking. So now we are more conscious of these risks" -P8, Child Welfare Supervisor, MSW, APSW, 19 years

However, participants (75%, n=15) also shared that the algorithm misses context about each case and the presence of some of these risk indicators did not mean the child was being sex trafficked.

75% ↓ Type of Decision-Support. *Predictive nature of AST frustrates caseworkers.* Several participants (75%) who found value in the algorithm as a guide also shared their frustrations in regard to its mandatory reporting nature. AST is used in a predictive capacity such that if certain risk indicators are selected then the case must be reported to HART. For instance, a permanency consultant shared:

"If a child is at risk for sex trafficking then we are having those conversations from the very beginning and taking necessary action. This decision tool is not helpful in the way it's being used and only frustrates HART" -P5, Permanency Consultant. MSW, 12 years

Reporting to HART takes the case away from the child-welfare team who have spent significant amount of time building a relationship with the child and their caregivers. Moreover, HART is receiving an influx of calls that do not require their expertise as a result of this algorithm.

75% ↓ Degree of Uncertainty. *Lack of context about the case leads to high degree of uncertainty.* Several participants (75%, n=15) shared that even though AST was useful as a guide, the presence of risk indicators did not mean the child was at risk of being trafficked. There was still a lot of pertinent information that was necessary to make such a determination. For instance, a supervisor shared that one of their foster kids had a history of sexual abuse and met a bunch of criteria on the tool, however, the supervisor's team is actively involved with the child and their caregivers, understand their needs, and did not believe the case needed to be reported to HART.

5.3.3 AST Impedes Human Discretion (RQ2). In this section, we discuss how AST algorithm maps onto the ADMAPS dimensions of human discretion.

70% ↑ Professional Expertise. *AST has made caseworkers more aware of risk indicators.* Most participants (70%, n=14) appreciated that this decision-tool taught new caseworkers to be actively aware of indicators that are associated with sex trafficking. As stated by P5 above, AST has improved professional expertise since caseworkers are more cognizant of such risks.

75% ↓ Value Judgments and Heuristic Decision-Making. *AST does not allow caseworkers to make value judgments or engage in heuristic decision-making.* Several participants (75%, n=15)

shared that the risk indicators were just part of a bigger picture and there was still a lot context that the child-welfare team needed to collectively unpack to better assess the situation. This requires the child-welfare team to engage in heuristic decision-making and weigh the necessary information. For instance, a case manager shared:

"One of my foster teens is sexually active and has a long-term boyfriend. But looking at the tool, everyone is like... 'Is that her boyfriend or her pimp?' I know this kid and she trusts me... building trust with these kids takes time... she doesn't need to be referred to HART [Human Anti-Trafficking Response Team]" -P17, Case Manager, BSW, 8 years

Here, the case manager emphasizes that having trusting relationships was equally important for preventing foster youth from being sex trafficked. Referring the case to HART takes this case away from the case manager and the child is assigned a new case manager from HART.

5.3.4 AST and Bureaucratic Processes (RQ2). Overall, we found that AST is inadequately supported by bureaucratic processes which leads to frustrations on part of caseworkers. In this section, we discuss how AST algorithm maps onto the ADMAPS dimensions of bureaucratic processes.

60% ↓ Resources & Constraints. *AST has caused an influx of cases referred to HART.* Several participants (60%, n=12) shared that AST has added to the frustrations of HART who do not have enough resources to manage all the cases that are being reported to them because of the mandatory reporting aspect of AST. This leaves HART with significantly less resources to focus on the cases that need them. A supervisor on HART explained:

"It's become overused and is being abused. It's being used in different settings and not how it was originally intended to be used. We are getting an influx of calls that don't need to be called." -P10, Child Welfare Supervisor, MSW, LPC, NCC, 9 years

65% ↑ 60% ↓ Administration & Training. *AST is conducted as an organizational process and offers some training.* Most participants (65%, n=13) shared that AST is routinely conducted for foster youth over 10 years of age. As discussed previously, it trains new caseworkers to be more conscious about risk factors but often to the detriment of HART. Participants also consider this decision-tool to be peripheral to the case planning process for most cases and utilized it as a mandated requirement.

60% ↑ Laws & Policies. *The use of AST is mandated by law.* The Department of Children and Families (DCF) has legally mandated the use of this algorithm as a means to proactively protect foster youth as well as collect more data about associated risk factors and the number of foster youth considered at risk. Next, we discuss the Legal Permanency Status (LPS) algorithm.

5.4 Legal Permanency Status (LPS) Algorithm

Tracking performance metrics such as placement stability and permanency emerged as another decision outcome for which an algorithm is being used. Several participants (65%, n=13) discussed this decision-tool, and we observed the tool being used at all the permanency consultations (n=40). Federal legislation has established permanency as one of the primary goals of CWS and requires agencies to meet this well-defined and measurable benchmark [125]. Legal permanency is defined as reunification with the biological family, adoption or transfer of guardianship. [67].

5.4.1 LPS High-Stakes Decision Outcomes (RQ1). This algorithm is used to track an outcome over time instead of predicting an outcome of interest using input data. The agency uses it at permanency consultations to track the quality of the current placement and if those specialized meetings are leading to better outcomes. The tool is facilitated by the permanency consultant and upon completion, the team categorically rates the quality of the placement. It is also used to

track systemic barriers that are getting in the way of achieving permanency. Some participants recognized (40%, n=8) the utility of this tool and a permanency consultant explained –

"It helps us to be actively aware of where we are at in terms of permanency. We have to follow a [15-month] timeline for permanency and if parents are not showing initiative and not completing court ordered services then we need to start exploring placement options. So, this tool is kind of an extra push" -P5, Permanency Consultant, MSW, 12 years

5.4.2 LPS and Algorithmic Decision-Making (RQ2). Overall, we found that even though LPS is used in a prescriptive capacity, it still lacks utility. In this section, we discuss how LPS algorithm maps onto the ADMAPS dimensions of algorithmic decision-making.

65% ↓ Relevant Data. *Definitions of input variables are ambiguous and frustrate caseworkers.* Several participants (65%, n=13) shared their indifference towards LPS and stated that the definitions of input variables in terms of what constitutes a “Good”, “Fair”, or “Poor” placement were ambiguous. A permanency consultant who is tasked with conducting this decision-tool shared:

"This tool is confusing in itself; the definitions are very vague. We can't necessarily put families into categories. We have to use this tool and have conversations about where they [families] might fit best, but it doesn't give a clear picture of what's really going on with the placement." - P6, Permanency Consultant, BSW, 3 years

Participants (65%, n=13) also shared that how LPS was conducted depended on the perspective of the team. Different people might look at the same set of facts and reach different conclusions as to why the placement should be rated as good, fair, or poor.

60% ↓ Type of Decision-Support. *LPS is used as a prescriptive tool, but caseworkers lack agency towards affecting outcomes.* Participants (60%, n=12) shared that the tool did not predict or recommend an outcome based on the inputs, however, the utility of the tool lies in recognizing the current state of a case with respect to permanency and addressing systemic barriers, but the child-welfare team lacks agency with respect to addressing several systemic barriers. For instance, critical decision-making power in regard to permanency decisions sits with the legal parties (district attorneys and judges) who can choose to fully disregard the child-welfare team's recommendation.

65% ↓ Degree of Uncertainty. *Lack of relevant information leads to high degree of uncertainty.* Participants (65%, n=13) were most frustrated by the lack of relevant information that LPS needed to account for in order to make a proper determination about the quality of the placement. Moreover, there is high uncertainty associated with the timeliness with which some systemic barriers can be addressed making it hard to assess the quality of the placement (see P2 quote below).

5.4.3 LPS impedes Human Discretion (RQ2). In this section, we discuss how LPS algorithm maps onto the ADMAPS dimensions of human discretion.

40% ↑ Professional Expertise. *LPS builds professional expertise by establishing an urgency towards permanency.* Some participants (40%, n=8) shared that the decision-tool was useful in that it established a sense of urgency and insistence towards achieving permanency for foster children. It teaches new caseworkers that they needed to prioritize finding placement options even if that upset biological parents. The child-welfare team follows a 15-month timeline where the parents must complete court order services within this timeline to achieve reunification. Towards the end of this timeline, CWS team must begin exploring alternate placement options.

60% ↓ Value judgments and Heuristic Decision-Making. *LPS tool does not allow for value judgments or heuristics.* Participants (60%, n=12) shared that the tool does not allow for value judgments or heuristics on part of the child-welfare team since decision-making power about permanency decisions sits with the legal parties.

5.4.4 LPS and Bureaucratic Processes (RQ2). Overall, LPS supports bureaucratic processes but does not account for organizational constraints. In this section, we discuss how LPS algorithm maps onto the ADMAPS dimensions of bureaucratic processes.

65% ↓ Resources & Constraints. *LPS does not account for organizational resources and imposes new constraints.* Participants shared (65%, n=13) that ambiguity around input variables and lack of agency with respect to permanency decisions has turned this decision-tool into documentation that every child-welfare team must complete at permanency consultations. Moreover, there are arbitrary constraints placed on the decision-tool in regard to how a placement can be rated. For instance, a program director explained:

"We may have a court hearing date set and all of a sudden, the placement is "Good". We now have a "Good" rating because we have a court hearing. But it might take three years to terminate parental rights or go to a guardianship. And then you went three years without achieving permanency but somehow, we have a "Good" placement rating." - P2, Permanency Consultant, MSW, 20 years

80% ↓ Administration & Training. *Caseworkers are indifferent towards how LPS is administered.* Even though the tool is expected to be central to permanency consultations, we noticed that for majority of these meetings (87%, n=35), LPS was used towards the end as a requirement to rate the quality of the placement. Most participants (80%, n=16) were indifferent towards LPS because of the several constraints and utility issues discussed above. One supervisor shared:

"It's just another thing we have to do in the permanency meetings. I let the Permanency Consultants score it however they like" - P15, Child Welfare Supervisor, MSW, 9 years

65% ↑ Laws & Policies. *The use of LPS tool is legally mandated and helps track performance outcomes.* Participants (65%, n=13) shared that the Department of Children and Families (DCF) has mandated using this decision-tool as a way to track the important outcomes of quality of placements and permanency. Moreover, it's important to track the systemic barriers that often impact permanency to formulate policies at the state level that address these systemic barriers.

5.5 Assessing the Benefits and Drawbacks of Differing Approaches (RQ3)

Table 1 offers a summary of how the four algorithms balance the dimensions of the framework. In this section, we discuss the benefits and drawbacks that arise when trying to balance the tradeoffs between human discretion, algorithmic decision-making, and bureaucratic processes.

5.5.1 Algorithmic Decision-Making Should Seek to Augment Human Discretion, Not Supplant it. In this section, we discuss the themes around the benefits and drawbacks that emerged when balancing the ADMAPS dimensions of human discretion and algorithmic decision-making.

When aligned, algorithms augment decision-making processes, but a lack of alignment can take away autonomy and heuristic decision-making. All of our participants emphasized that every case is contextually different and that a family's circumstances often change through the life of the case. For instance, a parent seeking reunification with their child might experience a lapse while trying to maintain a stable job, maintain their sobriety, or consistently attend court ordered services. In other words, algorithms should be designed with the recognition that there will be a high degree of uncertainty associated with any relevant data and the subsequently predicted outcome. Therefore, algorithms need to not only make room for human discretion but also facilitate value judgments and heuristics in order to offer utility. Most of our participants (80%, n=16) mentioned that 7ei algorithm augmented their decision-making processes when they were making difficult decisions. It prioritizes and enhances the value judgments and heuristic decision-making that

	Human Discretion		Algorithmic Decision-Making		Bureaucratic Processes	
CANS	Professional Expertise	80% ↓	Relevant Data	70% ↓	Resources/Constraints	40% ↑ 70% ↓
	Value Judgments	75% ↓	Decision-Support	90% ↓	Admin & Training	40% ↑ 80% ↓
	Heuristic Decisions	80% ↓	Uncertainty	70% ↓	Laws & Policies	80% ↑
7ei	Professional Expertise	85% ↑	Relevant Data	90% ↑	Resources/Constraints	60% ↑
	Value Judgments	70% ↑	Decision-Support	85% ↑	Admin & Training	75% ↑ 55% ↓
	Heuristic Decisions	85% ↑	Uncertainty	80% ↑	Laws & Policies	65% ↓
AST	Professional Expertise	70% ↑	Relevant Data	65% ↑ 75% ↓	Resources/Constraints	60% ↓
	Value Judgments	75% ↓	Decision-Support	75% ↓	Admin & Training	65% ↑ 60% ↓
	Heuristic Decisions	75% ↓	Uncertainty	75% ↓	Laws & Policies	60% ↑
LPS	Professional Expertise	40% ↑	Relevant Data	65% ↓	Resources/Constraints	65% ↓
	Value Judgments	60% ↓	Decision-Support	60% ↓	Admin & Training	80% ↓
	Heuristic Decisions	60% ↓	Uncertainty	65% ↓	Laws & Policies	65% ↑

Table 1. RQ3: Tradeoffs between balancing the dimensions within human discretion, algorithmic decision-making, and bureaucratic processes, Percentages (%) represent the proportion of participants who stated that a dimension was positively (or negatively) impacted by the algorithm.

caseworkers must engage in when devising action steps to help families. Participants (85%, n=17) especially appreciated that 7ei facilitated brainstorming and idea generation instead of predicting an outcome of interest. One case manager explained:

"Of all the things we have brought up, 7ei is my favorite because it helps us think differently, understand what a family has been through and then brainstorm ideas on how to help them based on this understanding of trauma" -P3, Permanency Consultant, MSW, APSW, 9 years

Participants noted that the tool offers flexibility and autonomy in how they interact with it and which 7ei domains they focused on. Adaptively selecting information instead of analyzing all the information is a key feature of heuristic decision-making and can lead to more accurate decisions. On the other hand, tensions arise when algorithms attempts to supplant human discretion. This is the case with CANS whose predictive nature does not account for the high degree of uncertainty that accompanies each case, and consequently, does not make room for discretionary work on part of the caseworkers. With a lack of autonomy, gaming the algorithm is the only way caseworkers are able to exercise discretion and produce the desired outcome. One supervisor shared:

"CANS is all about producing a good rate so foster parents can afford the resources they need to take care of the child. I have had foster parents put in notices [to end placement] because they couldn't support the child anymore" -P14, Child Welfare Supervisor, MSW, 30 years

This inflexible and predictive nature has shifted focus away from the primary outcome of interest (i.e. – mental health screening) and towards the secondary outcome that allocates resources.

When aligned, algorithms can help embed important value judgments into the decision-making process. Most of the participants (80%, n=16) felt strongly about the need to support each family in a different capacity and through different practices. This was also a dominant theme in all the planning meetings (n=15) and permanency consultations (n=40) where the child-welfare team devised specific plans for each family through trauma-informed care (i.e. - using 7ei). 7ei algorithm is centered in some of social work's core values of service, dignity and worth of the person, and importance of human relationships [119] and informs the child-welfare team's work processes. For instance, a supervisor explained:

"7ei helps us think about how we can help every family. What can we do to help mom develop her self-worth? How can we help her build relationships with relatives or people in the community, so she has more caregiver support" -P8, Child Welfare Supervisor, MSW, CAPSW, 19 years

Several participants (65%, n=13) explained that 7ei scores translated into actionable steps that directly sought to help children and their families. In the meetings, 7ei brainstorming sessions resulted in solutions that the child-welfare team could affect directly and not simply refer children and parents to therapy. For instance, the case managers and supervisors planned activities that the family could engage in to improve child as well as family functioning, discussed information to share with foster and biological parents about impacts of trauma exposure, as well as steps caregivers could take to establish healthy boundaries with foster children and enforce positive discipline. For instance, a permanency consultant explained:

"It [7ei] helps you have a perspective shift on the family and the child. You don't need to refer everyone and their mother to therapy. Sometimes it's just as simple as having them do something as a family that's different than what they ever did before [picnics, sports].. and challenge them in different ways" -P2, Permanency Consultant, MSW, 20 years

On the contrary, algorithms that do not embed human values into its design may consequentially end up minimizing them. For instance, CANS recommends the level of foster care that the child should be placed in and based on that the child-welfare team finds foster parents who have the resources to meet those needs. However, participants explained that more financial resources do not always equate to a foster child's well-being. For instance, a permanency consultant explained:

"Everyone has a different compass for well-being. We had a child who was placed with well-off foster parents in a five-bedroom house and he wasn't doing well there because it was a culture shock for him. He wasn't used to a huge home, a great school, a big backyard... and he completely shut down. So, we found his aunt and we moved him down to Chicago. It's a two-bedroom house with five other kids and this child is thriving! So, he needed to be with his family, and he needed to have his own culture. -P4, Permanency Consultant, MSW, APSW, 8 years

Here the permanency consultant emphasizes the core human value of having trusting relationships in one's life and the importance of one's culture that play a critical role towards achieve emotional and cognitive well-being. Unlike CANS, 7ei allowed the child-welfare team to prioritize *Reasons To Be* and *Regulation* to really help children instead of simply focusing on financial resources. Caseworkers must continually negotiate values and balance the individual needs of people with the demands of policymakers. Algorithms that do not allow caseworkers to make such value judgments may lead to more frustrations and limit their utility.

Algorithms that replace the need for caseworker expertise can lead to inadequate and unreliable decision-making. Another dominant theme that emerged in 70% (n=14) of the interviews, 53% (n=8) of the planning meetings, and 50% (n=20) of the permanency consultations was that the algorithms in use were diminishing the need for caseworker expertise and family-centered care in a very contextual domain. CWS is a high-stakes domain where most caseworkers lack adequate work experience, carry high caseloads, and are under a lot of pressure from legal parties. Most of the participants (70%, n=14) expressed concerns that algorithmic systems might simply act as a safe default for most caseworkers such that questioning an algorithmic decision would add more work to their plate. A program director explained –

"We are not hiring people with a lot of experience. All new hires are recent graduates who don't quite know what this field looks like. They are happy to trust the machine to just get through the day" -P7, Child Welfare Program Director, MSW, 20 years

Professional expertise in the public sector deteriorates when algorithms limit the scope for value judgments and heuristic decision-making. Value negotiations and heuristics are an indispensable aspect of professional practice that workers must continually engage in to build expertise [77]. There are also growing concerns that algorithms such as CANS are leading to children being

referred to unnecessary services which also shifts the focus away from family-centered care. CANS scores directly translate to actionable steps in the form of services that children are referred to. Caseworkers co-opt CANS to produce higher compensations for foster parents or the foster parents might exaggerate child behaviors, however, several participants (60%, $n=12$) recognized that this also meant unnecessary services being requested for children. One case manager explained:

"Foster parent might exaggerate behaviors just to get more money. And then we put the kid in therapy and are not addressing their needs specifically. The kid in therapy is then being asked 'why are you so sad?' and the kid is like.. 'I'm not sad.' " -P17, Case Manager, BSW, 8 years

This is further problematic because unnecessary services are not only an added financial burden on CWS but they also add to the medical trauma of foster children who are continually being told that something is wrong with them. Finally, as previously noted, services assigned through an isolated view of a child might be less effective than family-level interventions developed through a trauma-informed perspective. Algorithms like CANS, however, limit child-welfare workers from using their expertise and developing more family-centered practices.

5.5.2 Algorithmic Decision-Support Systems and Bureaucratic Processes Need to be Aligned. In this section, we discuss the themes around the benefits and drawbacks that emerged when balancing the ADMAPS dimensions of algorithmic decision-making and bureaucratic processes.

When algorithmic decision-making and bureaucratic processes are aligned, it can help train caseworkers. Most participants (70%, $n=14$) recognized the value of algorithms as essential training mechanisms. Continuous engagement with 7ei in trauma-informed meetings under proper supervision ensures that caseworkers are continually having conversations centered in trauma and are coming up with solutions founded in TIC. This is essential because CWS suffers from a high turnover with lack of adequate training and supervision being some of the leading causes [24]. During the observations, experienced members of the child-welfare team such as the supervisors and permanency consultants brainstormed ideas with the caseworkers using 7ei and shared practices and approaches that had worked in the past with other families. A program director explained:

"It [7ei] makes us think differently and in the moment think through TIC [trauma-informed care].. what is the impact of that [incident]? How can we help support parents and children and help with emotional regulation? What can we do to build their relationships and how do we support the people in their life so that they show better caregiver capacity to manage what this child is going through" -P7, Child Welfare Program Director, MSW, 20 years

7ei meetings are also attended by Caregiver Support Specialists and Family Preservation Specialists who based on case circumstances also offer their expertise to the caseworkers on how to proceed. Several participants (70%, $n=14$) also recognized the Anti Sex-Trafficking algorithm as a good training tool such that caseworkers are always aware of and looking for red flags associated with sex-trafficking. This tool is also used in a team-setting and allows for the less experienced caseworkers to learn from the seasoned members of the team in how to perceive certain risks as well as how to follow up on them without confronting the foster youth. One supervisor explained:

"Its a whole change of mindset and and we're now more cognizant of some of those red flags. So I appreciate that. So, anytime in supervision, if they [caseworkers] start talking about some of these things, we pull up the tool and go through it and start discussing how to proceed" -P15, Child Welfare Supervisor, 9 years

However, participants who found great value in the Anti Sex-Trafficking algorithm as a training mechanism were also equally frustrated by the mandatory reporting nature of it.

When the algorithmic decision-support system is not aligned with the constraints of the bureaucratic system, it leads to utility issues. The utility of algorithmic decision-support systems is severely limited when they do not account for the organizational constraints within which they must operate. This challenge and frustrations on part of caseworkers was observed in 80% (n=16) of the interviews, 73% (n=11) of the planning meetings, and 82% (n=32) of the permanency consultations. Most of the participants (80%, n=16) stated that they were unable to use algorithms as intended because of several organizational constraints such as limited time availability due to high caseloads. For instance, properly conducting CANS requires the caseworkers to interview several individuals such as the foster parents, relatives, and teachers to be able to get consistent information. However, most caseworkers have high caseloads and do not have enough time to devote towards properly conducting each assessment. A permanency consultant explained:

"It can take hours to properly do CANS. Everyone is stretched too thin.. we literally do not have the time to do that." -P2, Permanency Consultant, MSW, 20 years

Moreover, the bureaucratic policy that requires CANS to be conducted every six months makes it unfeasible for caseworkers to interview anyone beyond the foster parents. However, as previously noted, the assessment is also tied to the compensation offered to foster parents which further leads to foster parents exaggerating child behaviors. A supervisor explained that the algorithm did not need to be manipulated and caseworkers could request an exceptional rate to account for several different factors such as transportation services, therapy, school activities etc. However, she also explained that high caseloads make it hard to devote any extra time to each assessment:

"The supervisor or worker can take a little extra time and put an exceptional amount in. But again, everybody has too much to get done and it's unlikely for somebody to even think about that until they are in the middle of submitting CANS and its due the next day." -P8, Child Welfare Supervisor, MSW, CAPSW, 19 years

Caseworkers who continue to feel disempowered and frustrated by CANS find it easier to manipulate CANS scores to produce a higher rate than request the addition of an exceptional rate which needs to be approved by the supervisor. Most participants (75%, n=15) also shared significant concerns about the fidelity of data being used to develop algorithms. Administrative data curated through bureaucratic processes can not be uncritically used to develop algorithms. Participants shared that the data about families is collected by caseworkers and it was hard to acquire accurate information. Oftentimes, caseworkers might receive conflicting information from parents, relatives, and neighbors. A supervisor explained:

"People aren't willing to happily share personal information about themselves. And understandably so. It's really hard to be able to get consistent information to be able to put in some data system to bust out a decision. So I find this really difficult to comprehend how we would even consider that" - P13, Child Welfare Supervisor, MSW, 12 years

Families view caseworkers as representatives of the state and are unwilling to trust them or share any information that they might consider incriminating. Therefore, much of the information is based on the caseworker's perception of the family. This problem is further exacerbated by the fact that most caseworkers are new graduates and lack adequate experience. A program director explained that during investigations, caseworkers need to ask the right questions, read situations, and follow-up to be able to derive meaningful information. However, caseworkers acquire these skills through years of experience.

If the algorithms are not explained well to the practitioners, they don't trust them. Most of the participants (85%, n=17) were unaware of the algorithms being used. They did not recognize CANS or 7ei as algorithmic systems, even though they interact with them on a daily

basis. This finding is especially interesting because the CANS training material recognizes it as an algorithm and explicitly lays out the purposes (i.e. - outcomes) that the algorithm is designed to accomplish. CWS started using algorithms as a means to track important performance metrics such as permanency and placement stability as mandated by federal legislation [127]. Moreover, several psychometric assessments that are routinely used in child-welfare to assess the risks and needs of children and parents have adopted algorithmic analogues in the form of risk assessment algorithms. They are now being used as data collectors for all cases as well as decision arbiters for future cases. The CANS algorithm is a case in point of this scenario which is an algorithmic version of the CANS communimetric assessment [78] and has been re-purposed for other outcomes. Interestingly, caseworkers were well aware of assessments used at the agency but did not actively recognize the algorithmic components unless the interviewer nudged them to think about some of the automated aspects. For instance, when explicitly discussing CANS, all the interviewees realized it to be an algorithm that plays a pertinent role in decision-making. We asked a permanency consultant how she thought CANS worked and she exclaimed:

"Yes, okay! Right! Because we put in the scores, and it generates the rate. Oh and I hate the CANS! I, I've expressed it a lot of times that CANS should not be tied to money." -P6, Permanency Consultant, BSW, 3 years

As soon as participants recognized the automated decision-making aspects of these decision tools, they started to share some other ambiguous aspects. For instance, they had trouble recognizing who managed these decision-tools and shared that different tasks associated with these tools were distributed across the agency. Most of the participants (85%, n=17) were unable to isolate CANS to a particular role or "place". CANS is embedded in several aspects of the case planning process with different departments involved. Case management team focuses on assessing risks and needs whereas the data specialists track scores, and the fiscal liaisons review and approve foster parent compensations. Participants (70%, n=14) also alluded to the fact that algorithms meant different things to people based on their intended goals. For instance, one permanency consultant asserted:

"Case management try to do a good job with risks and needs but of-course foster parents only care about the rate. I know X's team [data specialists] look at scores to see if there is improvement. But there is also Y [fiscal liaison] who just wants the scores to be turned in on time so that he doesn't have to chase people down" -P1, Permanency Consultation Supervisor, MSW, APSW, 22 years

Several participants (75%, n=15) recognized that CANS, AST, and LPS decision tools meant different things in different contexts. Lack of awareness about algorithms at the organizational level and their distributed nature augmented caseworkers' distrust.

5.5.3 Collective Buy-in Requires Trust in Outcomes at Both the Caseworker and Bureaucratic Levels. In this section, we discuss the themes around the benefits and drawbacks that emerged when balancing the ADMAPS dimension of human discretion and bureaucratic processes with respect to decision outcomes that the CWS agency implements.

When the algorithmic outcomes miss important context, practitioners don't trust them

This challenge emerged in 70% (n=14) of the interviews, 67% (n=10) of the planning meetings and 75% (n=30) of the permanency consultations. Algorithms attempt to generalize child and family characteristics and place them in certain categories to be able to make a determination. However, this inadvertently leads to a loss of information with respect to the final outcome since all the information cannot be accounted for by these algorithms. For instance, the CANS algorithmic assessment focuses on child behaviors, risks, and needs based on the last 30 days and does not

account for traumatic triggers that can lead to serious emotional dysregulation from time to time. A permanency consultant explained:

"CANS does not account for trauma in the way 7ei does. We have a child that goes into manic depression every year around holiday season and needs to be cared for 24/7. So, one of the foster parents has to quit their job. However, there is no way to account for that in CANS. This child is doing fine right now but we know that traumatic trigger is coming." -P3, Permanency Consultant, MSW, APSW, 9 years

Participants shared similar concerns about the Anti Sex-Trafficking tool. Participants (65%, n=13) shared that the tool was useful to identify risk factors, however, the presence of risk factors does not necessarily mean that the foster youth is being trafficked. A supervisor explained:

"We have a 12 year old boy who suffers from a lot of sexual trauma. He met a bunch of criteria on the tool and we had to report to HART. We have been with this boy for many years.. we know him and what his behaviors are connected to. Reporting to HART felt very odd" -P11, Child Welfare Supervisor, MSW, APSW, 9 years

Here, the supervisor alludes to the fact that presence of risk factors did not capture the full picture and missed important context about what is going on with this child. As previously discussed, child-welfare teams are having on-going conversations about the risk of sex-trafficking from the onset of a case and are able to investigate concerns without reporting to HART. Continued interactions illustrated by these examples leads to an accumulation of distrust towards algorithms that often miss important nuances and context.

There is collective buy-in when practitioners believe that better outcomes are being achieved. This theme emerged in 85% (n=17) of the interviews, 67% (n=10) of the planning meetings, and 62.5% (n=25) of the permanency consultations. Participants appreciated 7ei and trauma-informed care because it allowed them to directly help families by developing specific plans for them. Moreover, they were aware that trauma-responsive services developed through TIC lead to better permanency and placement outcomes for children. When frustrated by algorithms such as CANS, caseworkers often asked why they were still expected to use CANS when 7ei was leading to better outcomes.

The agency developed a comprehensive four-part evidence-based service program centered in trauma-informed care of which the algorithmic tool is just one component. The first component of this program involves extensive ongoing TIC trainings. The second component introduced child-welfare staff to trauma-informed assessments that resulted in family-level interventions. The third component was the availability of specialized supervision and consultations provided by a clinical supervisor, TIC program administrator, medical expert, and a caregiver support specialist. Finally, the fourth component of this program is the 7ei algorithm which allows the child-welfare team to break down each case into seven domains and apply TIC principles and practices to it. In sum, the agency significantly invested in resources to ensure that TIC practice and 7ei were fully-supported by bureaucratic processes. A better understanding of what a family might be going through at a psychological level and knowing exactly how to help them has led to the collective buy-in from caseworkers. Most participants (85%, n=17) found great value in TIC and appreciated 7ei as a tool designed to be a guiding framework for TIC. One permanency consultant explained:

"Compassion fatigue in child welfare is very real. This tool has helped me truly understand what trauma can do to the brain. So I think now that we have a better understanding of trauma... we also have more empathy, and we go the extra mile to help families." -P1, Permanency Consultation Supervisor, MSW, APSW, 22 years

This was also evident at the meetings where the child-welfare staff urged case managers to ask and think in terms of *"What happened to you and how can I help you?"* instead of *"What is wrong with you?"* when working with children and families. That is, having a perspective shift and thinking in terms of the impact of trauma and not just exhibited behaviors. Moreover, a study published by independent researchers showed that child-welfare teams who implemented TIC practices using 7ei exhibited improved permanency and placement stability outcomes for their cases. This has further deepened the collective buy-in from caseworkers who often brought up the "study" to state that their practice centered in TIC worked and lead to better outcomes, and therefore, the legal parties needed to trust their judgments more. However, interestingly, none of the participants except the program director were able to locate the study and share it with us.

Lack of trust in algorithmic outcomes, leads to concerns about unethical and unsound decision-making. This theme emerged in 80% (n=16) of the interviews, 40% (n=6) of the planning meetings, and 45% (n=18) of the permanency consultations where participants recognized that algorithms in use sometimes led them to make unethical or unsound decisions. Surprisingly, even though 90% (n=18) of the participants recognized that caseworkers were gaming the system to produce higher compensations, some participants (40%, n=8) did not consider the decisions made to be unethical or unsound. To them, this is how the system was set up to be used. With contradictory incentive structures and conflicting values, caseworkers are often put in a position where they are forced to make such decisions. As previously discussed, both caseworkers and foster parents are co-opting CANS to produce a higher compensation, however, most participants (90%, n=18) recognized that the base compensation offered to foster parents is pretty low and gaming the algorithm was the only convenient way to produce an adequate rate. A supervisor explained:

"There is a lack of good foster homes and there is no financial incentive to be doing this work... So caseworkers do whatever they can to get them [foster parents] the money they need or we risk disrupting a placement...Several foster parents have put in notices in the past because they can't financially sustain the placement" -P9, Child Welfare Supervisor, MSW, 13 years

Another challenging aspect of using CANS where it leads to unsound decision-making is that it is now being used to track progress with respect to mental health; an unintended outcome of the continuous data collection. Consequently, an improvement (or deterioration) in mental health is impacting placement decisions. A clinical therapist explained:

"Sometimes a kid isn't doing well because they are working through things. For example, when a kid starts therapy sometimes their behaviors gets worse, but it might actually be that they're working through some things in therapy. And that's not captured by the tool. When there is no context, I just can't interpret whether or not that worsening is actually a bad thing." -P16, Clinical Therapist, LCSW, 5 years

These frustrations were also consistent with the fact that CANS is tied to service-planning where CANS recommends unnecessary child-level services when family-level interventions are often more effective in helping the family heal and cope with trauma. Prior studies have found that caseworkers distrust risk assessment algorithms such as CANS because of their deficit-based nature [20]. Interestingly, participants who were averse to this deficit-based nature did find value in risk assessment in regard to understanding a family's history. One permanency consultant explained:

"The data gives you a bigger picture of what the family is going through. Because the families are not always honest with us about what they need. And so if you can see that they were in drug and alcohol services or housing authority... that would be helpful in child welfare. Not necessarily giving them a score. I don't even know what to do with the score. It tells me nothing." -P6, Permanency Consultant, BSW, 3 years

Understanding a family's past helps caseworkers assess the type of services that would be most beneficial for them. Looking at this data from a trauma-informed perspective allows caseworkers to better understand the underlying trauma in the family that might be leading to exhibited behaviors and subsequent interactions with the child-welfare system. This points to a change in mindset in regard to the intended goal of risk assessment of (i.e. - predicting the risk of future harm) towards a strength-based outcome of providing the right services to families and not necessarily labeling them as high or low risk. The permanency consultant here also invokes social work's core value of service and helping families and sees the utility in risk assessment through that perspective.

6 DISCUSSION

Our results provide implications for algorithmic decision-making in CWS, more broadly for the public sector, as well as specific design guidelines for developing such systems in the public sector.

6.1 Implications for Algorithmic Decision-Making in the Child-Welfare System

6.1.1 Identifying Gaps and Opportunities for Algorithmic Decision-Making for Child-Welfare through the ADMAPS Framework. Our results suggest that caseworkers have differing perceptions of the design, outcomes, and intended uses of the various algorithms. ADMAPS dimensions helped uncover caseworker perspectives with respect to different aspects of algorithms. For instance, examining the bureaucratic processes at the agency revealed that the most frustrating part about CANS is that it is being used to calculate financial resources for foster parents. Furthermore, deliberating over the relevant data and degree of uncertainty with the participants revealed that CANS did not account for traumatic triggers in a child's ecosystem or the lack of interpretation regarding worsening behaviors. That is, ADMAPS uncovered the multiple and conflicting roles of CANS as a compensation calculator as well as a mental health assessment tool. On the other hand, participants considered 7ei to be analogous to the trauma-informed care framework but not an algorithm that was continually collecting data and tracking outcomes over the life of a case. Differing socio-technical imaginations of algorithms in the public sector are aligned well with existing literature [20, 62, 103, 107, 112]. Moreover, while many of our participants were able to connect how CANS risk scores were used in making decisions about a child based on state guidelines, yet others were unable to make this connection since CANS is deeply embedded within their daily culture, where they did not recognize it as an algorithm anymore. This aligns with prior literature on thinking about algorithms as part of culture [112] and points to a need in raising more awareness in caseworkers about the algorithms that they use daily. One way to improve this is to support explanations of algorithms in daily use; not just interpretations of outcomes but also designing around other interactions (e.g. data collection, input, visualization etc). Recent work on explainable AI has found support for focusing on explanations of daily interactions of users with algorithms [73, 88]. In addition, transparency of data, methods, and outcomes can support collaborative algorithmic decision-making processes [61, 110]. Finally, assessing the scope of bureaucratic processes and how algorithms must function within these constraints allows for better practitioner engagement and leads to algorithms that are centered in practice.

6.1.2 Balancing Strength-based and Deficit-based Approaches. Our results show that there are differing approaches in algorithmic implementation that is mandated by the state as opposed to the non-profit agency that is contracted by the state, to provide child-welfare services. For instance, the state mandates the use of CANS, which the caseworkers are not very amenable to as they can see the deficiencies in using this algorithm to provide good services for children in reality. In response, the agency developed 7ei using best practices from health services and social work (Trauma-Informed Care) [72, 119], and continue to use this algorithm in parallel with CANS to provide a more holistic perspective for the best outcome for a child and family. As outlined in our

results, CANS and 7ei affect the three dimensions of ADMAPS differently because fundamentally, CANS is a deficit-based assessment that focuses on risk mitigation and resource allocation; on the other hand, 7ei is a strength-based assessment that focuses on improving the outcomes for children. This is not to imply that CANS is always employed in a deficit-based context while 7ei only focuses on building strengths. These kinds of assessments can be commonly found in public services beyond child-welfare. For instance, within criminal justice systems, deficit-based risk assessments are the norm [117] but are heavily criticized for being unfair [14]. An important takeaway is that we need both strength and deficit-based approaches to make better decisions in high stakes environments. An initial assessment of risk is necessary but a transition towards helping people through strength-based approaches is equally important to prevent referrals and future interactions with the system [6–9]. Moreover, regardless of the type of outcome, street-level bureaucrats in the public sector must still exercise discretion and contextualize the algorithmic results for each case and act within the constraints posed by bureaucratic processes as explained by the ADMAPS framework [2, 62, 91].

6.2 Implications for Algorithmic Decision-Making in the Public Sector

6.2.1 ADMAPS Emphasizes Managing Interdependencies and Trade-offs in Algorithmic Decision-Making Within the Public Sector. The purpose of ADMAPS is twofold : 1) to interrogate algorithmic interventions in the public sector in a way that ensures that they balance the dimensions of *human discretion* and *bureaucratic process* with *algorithmic decision-making*, and 2) to offer practical guidelines for developing algorithms that offer higher utility to practitioners. A key implication of using ADMAPS is that better algorithmic decisions are made when algorithms account for and balance the complex interdependencies within socio-technical systems, rather than operate in isolation. We found that when one aspect of the framework was optimized, the other aspects of the framework often suffered. For instance, CANS demonstrated an over-reliance on predictions to support mandated bureaucratic processes; therefore, it minimized human discretion to the point that it negated collective buy-in and created ethical dilemmas where caseworkers felt pressured to manipulate the system. Additionally, each dimension of the framework had cross-dependencies with others, demonstrating practical trade-offs at both the macro- and micro-levels of the model. For example, predictive versus prescriptive types of decision-support affected administration and training in different ways. The predictive nature of CANS tended to replace the need for training for interpreting outcomes (i.e., providing the answer), while the prescriptive nature of 7ei augmented training by helping new CWS employees learn through the collaborative brainstorming process. Thus, a core contribution of this paper is that we demonstrate how the theoretically derived ADMAPS framework (empirically validated through this case study in child-welfare) can be used within other public sector domains (e.g., criminal justice [58, 59]) to assess the strengths and weaknesses of algorithms making high-stakes decisions in the lives of people. For instance, ADMAPS can help develop algorithms for judicial decision-making designed to aid judges. Judges must make decisions based on legal justification, interpretation, and application based on relevant laws and precedents [40]. The fluid nature of legal reasoning is often at odds with the discrete predictive nature of algorithms [55]. Here, similar to the 7ei algorithm, ADMAPS can help develop analogues that augment human discretion instead of curtailing it. Similarly, ADMAPS can help design algorithms for job placement centers where the caseworkers must exercise discretion in applying complex legal frameworks, assessing resource constraints, as well as resolving organizational contradictions (i.e., *bureaucratic processes*) to extend unemployment benefits to citizens [62].

6.2.2 Holistic Assessments, Not Deterministic Scoring Lead to Improved Decisions. Through the lens of ADMAPS, the tensions between the three elements (human discretion, bureaucratic processes,

and algorithmic decision-making) become apparent. This further implies that policymakers and most practitioners do not get the opportunity to understand and appreciate the value of holistic algorithmic assessments that can augment the current mandates of univariate risk scores. These mandates exist based on state legislation and are meant to provide a legal justification for using algorithms in child-welfare. While the legal implications are out of scope for this paper, our main takeaway is that the current tension that exists between the state and the agency needs to be redressed in order to improve outcomes for children. This is not to say that we must mandate decision-making from both CANS and 7ei but on the other hand, reduce the dependency on singular metrics and increase the dependency on understanding the underlying trauma and behaviors of a particular child. Moreover, as clearly depicted through our results, using singular metrics from a risk assessment to make subsequent decisions about unrelated determinants (for e.g., calculating compensation for foster parents) only incentivises gaming the system and results in a vicious cycle for children in foster care. The benefits of strength-based approaches have been supported in prior literature [9, 99, 107, 119] as well as the pitfalls for reappropriating algorithmic outcomes for different purposes [79, 106]. More importantly, algorithmic decision-making should not be treated as inevitable; knowing when not to design [13], not to deploy [10], and to resist [108] is equally important. If systems designers and policymakers are unable to balance the dimensions of ADMAPS in a way that it serves on-the-ground practitioners then alternative, non-algorithmic approaches such as collaborative assessments or processes should be developed that support and/or streamline bureaucratic processes and augment the quality of human discretionary work.

6.3 Design Guidelines for Algorithmic Decision Support Systems in the Public Sector

Based on our findings we provide the following design guidelines for developing algorithmic decision-support systems in the public sector. Our guidelines highlight the need to support the complex interdependencies between the three dimensions of ADMAPS.

- Consider making algorithmic outputs multidimensional, rather than a singular metric. This allows for flexibility in interpreting the output through the use of human discretion.
- Make algorithmic metrics suggestions, rather than mandated decision outcomes. Creating flexibility in the organizational process will reduce bureaucratic overhead and prevent the system from getting overburdened.
- Account for the degree of uncertainty in the data and the associated outcomes and make room for value judgments and heuristic decision-making.
- Design algorithmic systems to be used collaboratively so that joint oversight and expertise can provide fairness, transparency, and accountability.
- Consult key stakeholders to form a consensus on what data should and should not be collected to inform high-stake decisions.
- Design the system to learn and adapt from expert users by being able to identify exemplar cases and the reasons why they were successful. This may lead to serendipitous data points that were not previously captured formally by the system.
- Account for the organizational resources and constraints within which all decisions (human and algorithmic) must be made and incorporate this into algorithm design.
- Avoid direct trade-offs between input variables that create ethical dilemmas. Create safeguards in the system that check for gaming behaviors, such as "what-if" analyses.
- If tradeoffs cannot be properly managed, consider alternative, non-algorithmic approaches that streamline and/or support bureaucratic processes and augment human discretion.

The objective of employing algorithms in the public sector is to improve decision-making by providing efficient, consistent, fair, and defensible decisions. However, as depicted by this case study

and the ADMAPS framework, there is a high degree of uncertainty associated with these algorithmic decisions which are further constrained by policies and organizational factors. Therefore, designing for algorithms that support practitioners' decision-making processes (instead of providing predicted outcomes) may offer higher utility and augment the quality of human discretionary work.

7 LIMITATIONS AND FUTURE WORK

Our study contributes both a generalizable framework of Algorithmic Decision-Making Adapted for the Public Sector (ADMAPS) and an in-depth ethnographic case study of sociotechnical systems used in the domain of child-welfare. However, there are several limitations that introduce opportunities for scholars to expand upon this research. First, this study solely focuses on the perspective of CWS caseworkers on algorithmic decision-making systems. We believe it is important to also include the perspective of CWS children and families as they are the affected communities of concern. Future research should focus on understanding the perspectives of families and their (lack of) agency with respect to the decisions made about them through the use of algorithmic systems. Second, this study only focused on the algorithms being used in collaborative team settings; however, there might be other decision-tools that caseworkers might use independently in their daily work. For instance, all CWS employees are required to use a comprehensive state mandated data system that has several data-driven visualizations and decision-tools built into it. But as our results indicate, caseworkers lack an adequate understanding of algorithms or decision-tools and only recognize the automated decision-making aspects when explicitly asked to think in those terms. Therefore, future research should focus on studying the state-mandated data systems, in-built decision-tools, and their impact on caseworkers' decisions more holistically.

8 ACKNOWLEDGMENTS

This research was supported by the National Science Foundation (CRII-1850517) and the William T. Grant Foundation (#187941, #190017). Any opinion, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of our sponsors or community partners. We would like to thank our collaborators at the child-welfare agency, SaintA, for allowing us to conduct this extensive ethnography. We are also thankful for our study participants as well as the anonymous reviewers whose suggestions and comments helped improve this manuscript.

9 CONCLUSION

We conducted an in-depth ethnographic study to understand the daily algorithmic practices of caseworkers at a child-welfare agency. Concurrently, we also developed a cohesive framework of Algorithmic Decision-Making Adapted for the Public Sector (ADMAPS) by systematically reviewing and synthesizing prior literature in HCI, STS, and PA. We qualitatively coded our data from the ethnography to the dimensions of ADMAPS to reveal the complex interdependencies between human discretion, algorithmic decision-making, and bureaucratic processes. Our findings show that there is a need to invest in strength-based approaches centered in ecological frameworks. This approach not only seeks to improve the lives of people but also builds collective trust in the outcome itself, and subsequently, leads to collective buy-in at both the caseworker and bureaucratic levels. Moreover, algorithms need to be designed such that they augment human discretion by allowing practitioners to engage in value judgments and heuristic decision-making. In addition, algorithms need to be fully supported by bureaucratic processes by allocating necessary resources and accounting for organizational constraints. As a result of this study, we also propose heuristic guidelines for the design of high-stakes algorithmic decision-making tools in the public sector.

REFERENCES

- [1] Senate of the United States 115th Congress. 2018. S. 3039 - Using Data To Help Protect Children and Families Act.
- [2] Ali Alkhatib 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.
- [3] Doris Allhutter, Florian Cech, Fabian Fischer, Gabriel Grill, and Astrid Mager. 2020. Algorithmic profiling of job seekers in Austria: how austerity politics are made effective. *Front. Big Data* 3: 5. doi: 10.3389/fdata (2020).
- [4] Florian Artinger, Malte Petersen, Gerd Gigerenzer, and Jürgen Weibler. 2015. Heuristics as adaptive decision strategies in management. *Journal of Organizational Behavior* 36, S1 (2015), S33–S52.
- [5] Jos ARM Arts, Wim H Gijssels, and Henny Boshuizen. 2000. Expertise Development in Managerial Sciences: The Use of Knowledge Types in Problem-Solving. (2000).
- [6] Karla Badillo-Urquiola, Jaclyn Abraham, Arup Kumar Ghosh, and Pamela Wisniewski. 2018. A Stakeholders’ Analysis of the Systems that Support Foster Care. In *Proceedings of the 2018 ACM Conference on Supporting Groupwork*. ACM, 158–161.
- [7] Karla Badillo-Urquiola, Arup Kumar Ghosh, and Pamela Wisniewski. 2017. Understanding the unique online challenges faced by teens in the foster care system. In *Companion of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. ACM, 139–142.
- [8] Karla Badillo-Urquiola, Scott Harpin, and Pamela Wisniewski. 2017. Abandoned but Not Forgotten: Providing Access While Protecting Foster Youth from Online Risks. In *Proceedings of the 2017 Conference on Interaction Design and Children*. ACM, 17–26.
- [9] Karla Badillo-Urquiola, Xinru Page, and Pamela Wisniewski. 2019. Risk vs. Restriction: The Digital Divide between Providing a Sense of Normalcy and Keeping Foster Teens Safe Online. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM.
- [10] Solon Barocas, Asia J Biega, Benjamin Fish, Jędrzej Niklas, and Luke Stark. 2020. When not to design, build, or deploy. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 695–695.
- [11] Solon Barocas and Andrew D Selbst. 2016. Big data’s disparate impact. *Calif. L. Rev.* 104 (2016), 671.
- [12] Eric PS Baumer. 2017. Toward human-centered algorithm design. *Big Data & Society* 4, 2 (2017), 2053951717718854.
- [13] Eric PS Baumer and M Six Silberman. 2011. When the implication is not to design (technology). In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2271–2274.
- [14] Richard Berk, Hoda Heidari, Shahin Jabbari, Michael Kearns, and Aaron Roth. 2018. Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods & Research* (2018), 0049124118782533.
- [15] Joan M Blakey, Sonya J Leathers, Michelle Lawler, Tyreasa Washington, Chiralaine Natschke, Tonya Strand, and Quenette Walton. 2012. A review of how states are addressing placement stability. *Children and Youth Services Review* 34, 2 (2012), 369–378.
- [16] Herman L Boschken. 1994. Organizational performance and multiple constituencies. *Public Administration Review* 54, 3 (1994), 308–312.
- [17] Mark Bovens 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 (2002), 174–184.
- [18] Evelyn Z Brodtkin. 2008. Accountability in street-level organizations. *Intl Journal of Public Administration* 31, 3 (2008), 317–336.
- [19] Urie Bronfenbrenner. 1975. Reality and research in the ecology of human development. *Proceedings of the American Philosophical Society* 119, 6 (1975), 439–469.
- [20] Anna Brown, Alexandra Chouldechova, Emily Putnam-Hornstein, Andrew Tobin, and Rhema Vaithianathan. 2019. Toward Algorithmic Accountability in Public Services: A Qualitative Study of Affected Community Perspectives on Algorithmic Decision-making in Child Welfare Services. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 41.
- [21] Justin B Bullock. 2019. Artificial intelligence, discretion, and bureaucracy. *The American Review of Public Administration* 49, 7 (2019), 751–761.
- [22] Justin B Bullock, Jesper Rosenberg Hansen, and David J Houston. 2018. Sector Differences in Employee’s Perceived Importance of Income and Job Security: Can These be Found Across the Contexts of Countries, Cultures, and Occupations? *International Public Management Journal* 21, 2 (2018), 243–271.
- [23] Peter André Busch and Helle Zinner Henriksen. 2018. Digital discretion: A systematic literature review of ICT and street-level discretion. *Information Polity* 23, 1 (2018), 3–28.
- [24] Sarah Carnochan, Megan Moore, and Michael J Austin. 2013. Achieving placement stability. *Journal of evidence-based social work* 10, 3 (2013), 235–253.
- [25] Eleanor Carter, Clare FitzGerald, Ruth Dixon, Christina Economy, Tanyah Hameed, and Mara Airoidi. 2018. Building the tools for public services to secure better outcomes: Collaboration, Prevention, Innovation. *Government Outcomes*

Lab, University of Oxford, Blavatnik School of Government (2018).

- [26] Spencer P Chainey. 2013. Examining the influence of cell size and bandwidth size on kernel density estimation crime hotspot maps for predicting spatial patterns of crime. *Bulletin of the Geographical Society of Liege* 60 (2013), 7–19.
- [27] Paul Chalekian. 2013. POSDCORB: Core patterns of administration. In *Proceedings of the 20th Conference on Pattern Languages of Programs*. The Hillside Group, 17.
- [28] Angèle Christin, Alex Rosenblat, and Danah Boyd. 2015. Courts and predictive algorithms. *Data & CivilRight* (2015).
- [29] Shelby L Clark, Ashley N Palmer, Becci A Akin, Stacy Dunkerley, and Jody Brook. 2020. Investigating the relationship between trauma symptoms and placement instability. *Child Abuse & Neglect* 108 (2020), 104660.
- [30] Victoria Clarke, Virginia Braun, and Nikki Hayfield. 2015. Thematic analysis. *Qualitative psychology: A practical guide to research methods* (2015), 222–248.
- [31] Roxanne Connelly, Christopher J Playford, Vernon Gayle, and Chris Dibben. 2016. The role of administrative data in the big data revolution in social science research. *Social science research* 59 (2016), 1–12.
- [32] Stephanie Cuccaro-Alamin, Regan Foust, Rhema Vaithianathan, and Emily Putnam-Hornstein. 2017. Risk assessment and decision making in child protective services: Predictive risk modeling in context. *Children and Youth Services Review* 79 (2017), 291–298.
- [33] Mark Davidson. 2020. Extreme municipal fiscal stress and austerity? A case study of fiscal reform after Chapter 9 bankruptcy. *Environment and Planning C: Politics and Space* 38, 3 (2020), 522–538.
- [34] Maria De-Arteaga, Riccardo Fogliato, and Alexandra Chouldechova. 2020. A Case for Humans-in-the-Loop: Decisions in the Presence of Erroneous Algorithmic Scores. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [35] James W Dean Jr and Mark P Sharfman. 1993. Procedural rationality in the strategic decision-making process. *Journal of management Studies* 30, 4 (1993), 587–610.
- [36] Ganesh Chandra Deka. 2014. Big data predictive and prescriptive analytics. In *Handbook of research on cloud infrastructures for Big Data analytics*. IGI Global, 370–391.
- [37] Benoit Desmarchelier, Faridah Djellal, and Faiz Gallouj. 2019. Innovation in public services in the light of public administration paradigms and service innovation perspectives. *European Review of Service Economics and Management* 8 (2019), 91–120.
- [38] Paul Dourish and Edgar Gómez Cruz. 2018. Datafication and data fiction: Narrating data and narrating with data. *Big Data & Society* 5, 2 (2018), 2053951718784083.
- [39] Grant Drawve, Joel Caplan, and Michael Ostermann. 2019. Utilizing a risk of crime measure for recidivism research: Moving environmental corrections forward. *Crime & Delinquency* 65, 5 (2019), 606–629.
- [40] Phoebe C Ellsworth. 2005. Legal Reasoning (The Cambridge Handbook of Thinking and Reasoning, KJ Holyoak and RG Morrison Jr.(eds). (New York.
- [41] K Anders Ericsson and Jacqui Smith. 1991. *Toward a general theory of expertise: Prospects and limits*. Cambridge University Press.
- [42] Virginia Eubanks. 2018. *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press.
- [43] Ali Farazmand. 2009. *Bureaucracy and administration*. CRC Press.
- [44] Ali Farazmand. 2017. Governance reforms: The good, the bad, and the ugly; and the sound: Examining the past and exploring the future of public organizations. *Public Organization Review* 17, 4 (2017), 595–617.
- [45] Patricia Fredericksen and Rosanne London. 2000. Disconnect in the hollow state: The pivotal role of organizational capacity in community-based development organizations. *Public Administration Review* 60, 3 (2000), 230–239.
- [46] H George Frederickson, Kevin B Smith, Christopher W Larimer, and Michael J Licari. 2015. *The public administration theory primer*. Westview Press.
- [47] Batya Friedman, Peter Kahn, and Alan Borning. 2002. Value sensitive design: Theory and methods. *University of Washington technical report* 2-12 (2002).
- [48] David J Fritzsche. 1991. A model of decision-making incorporating ethical values. *Journal of Business Ethics* 10, 11 (1991), 841–852.
- [49] Eileen Gambrill and Aron Shlonsky. 2000. Risk assessment in context.
- [50] Eileen Gambrill and Aron Shlonsky. 2001. The need for comprehensive risk management systems in child welfare. *Children and Youth Services Review* 23, 1 (2001), 79–107.
- [51] Sarah Giest and Nadine Raaphorst. 2018. Unraveling the hindering factors of digital public service delivery at street-level: the case of electronic health records. *Policy Design and Practice* 1, 2 (2018), 141–154.
- [52] Gerd Gigerenzer and Wolfgang Gaissmaier. 2011. Heuristic decision making. *Annual review of psychology* 62 (2011), 451–482.
- [53] Philip Gillingham. 2018. Decision-making about the adoption of information technology in social welfare agencies: Some key considerations. *European Journal of Social Work* 21, 4 (2018), 521–529.

- [54] Philip Gillingham. 2019. Decision Support Systems, Social Justice and Algorithmic Accountability in Social Work: A New Challenge. *Practice* 31, 4 (2019), 277–290.
- [55] Travis Greene, Galit Shmueli, Jan Fell, Ching-Fu Lin, Mark L Shope, and Han-Wei Liu. 2020. The Hidden Inconsistencies Introduced by Predictive Algorithms in Judicial Decision Making. *arXiv preprint arXiv:2012.00289* (2020).
- [56] Nina Grgić-Hlača, Christoph Engel, and Krishna P Gummadi. 2019. Human decision making with machine assistance: An experiment on bailing and jailing. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–25.
- [57] Kenneth R Hammond. 1996. *Human judgment and social policy: Irreducible uncertainty, inevitable error, unavoidable injustice*. Oxford University Press on Demand.
- [58] MD Romael Haque, Katherine Weathington, Joseph Chudzik, and Shion Guha. 2020. Understanding Law Enforcement and Common Peoples’ Perspectives on Designing Explainable Crime Mapping Algorithms. In *Conference Companion Publication of the 2020 on Computer Supported Cooperative Work and Social Computing*. 269–273.
- [59] MD Romael Haque, Katherine Weathington, and Shion Guha. 2019. Exploring the Impact of (Not) Changing Default Settings in Algorithmic Crime Mapping-A Case Study of Milwaukee, Wisconsin. In *Conference Companion Publication of the 2019 on Computer Supported Cooperative Work and Social Computing*. 206–210.
- [60] Alison Hendricks, Lisa Conradi, and Charles Wilson. 2011. Creating trauma-informed child welfare systems using a community assessment process. *Child Welfare* 90, 6 (2011), 187.
- [61] Joana Hois, Dimitra Theofanou-Fuelbier, and Alischa Janine Junk. 2019. How to Achieve Explainability and Transparency in Human AI Interaction. In *International Conference on Human-Computer Interaction*. Springer, 177–183.
- [62] Naja Holten Møller, Irina Shklovski, and Thomas T Hildebrandt. 2020. Shifting concepts of value: Designing algorithmic decision-support systems for public services. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society*. 1–12.
- [63] Steven J Jackson, Tarleton Gillespie, and Sandy Payette. 2014. The policy knot: re-integrating policy, practice and design in cscw studies of social computing. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*. 588–602.
- [64] Elizabeth K Hopper, Ellen L Bassuk, and Jeffrey Olivet. 2010. Shelter from the storm: Trauma-informed care in homelessness services settings. *The Open Health Services and Policy Journal* 3, 1 (2010).
- [65] Geraldine Kennett et al. 2013. The impact of training practices on individual, organisation, and industry skill development. *Australian Bulletin of Labour* 39, 1 (2013), 112.
- [66] Rob Kitchin and Tracey Lauriault. 2014. Towards critical data studies: Charting and unpacking data assemblages and their work. (2014).
- [67] Eun Koh and Mark F Testa. 2008. Propensity score matching of children in kinship and nonkinship foster care: Do permanency outcomes differ? *Social Work Research* 32, 2 (2008), 105–116.
- [68] Jan-Erik Lane. 2000. *New public management*. Taylor & Francis US.
- [69] Soo-Young Lee and Andrew B Whitford. 2013. Assessing the effects of organizational resources on public agency performance: Evidence from the US federal government. *Journal of Public Administration Research and Theory* 23, 3 (2013), 687–712.
- [70] Per LËgreid. 2017. *Transcending new public management: the transformation of public sector reforms*. Routledge.
- [71] Agnieszka Leszczynski. 2016. Speculative futures: Cities, data, and governance beyond smart urbanism. *Environment and Planning A: Economy and Space* 48, 9 (2016), 1691–1708.
- [72] Jill Levenson. 2017. Trauma-informed social work practice. *Social Work* 62, 2 (2017), 105–113.
- [73] Q Vera Liao, Daniel Gruen, and Sarah Miller. 2020. Questioning the AI: Informing Design Practices for Explainable AI User Experiences. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [74] Michael Lipsky. 2010. *Street-level bureaucracy: Dilemmas of the individual in public service*. Russell Sage Foundation.
- [75] Linda Lobao, Mia Gray, Kevin Cox, and Michael Kitson. 2018. The shrinking state? Understanding the assault on the public sector.
- [76] Thomas Lodato and Carl DiSalvo. 2018. Institutional constraints: the forms and limits of participatory design in the public realm. In *Proceedings of the 15th Participatory Design Conference: Full Papers-Volume 1*. 1–12.
- [77] Kim Loyens and Jeroen Maesschalck. 2010. Toward a theoretical framework for ethical decision making of street-level bureaucracy: Existing models reconsidered. *Administration & Society* 42, 1 (2010), 66–100.
- [78] John S Lyons. 2009. *Communimetrics: A communication theory of measurement in human service settings*. Springer Science & Business Media.
- [79] John S Lyons, Dana Aron Weiner, and Melanie Buddin Lyons. 2004. Measurement as communication in outcomes management: The child and adolescent needs and strengths (CANS). *The Use of Psychological Testing for Treatment Planning and Outcomes Assessment. Volume 2: Instruments for Children and Adolescents* (2004).
- [80] Thomas M. Vogl. 2020. Artificial intelligence and organizational memory in government: the experience of record duplication in the child welfare sector in Canada. In *The 21st Annual International Conference on Digital Government Research*. 223–231.

- [81] Mariusz Maciejewski. 2017. To do more, better, faster and more cheaply: Using big data in public administration. *International Review of Administrative Sciences* 83, 1_suppl (2017), 120–135.
- [82] Azhar Manzoor. 2014. A look at efficiency in public administration: Past and future. *Sage Open* 4, 4 (2014), 2158244014564936.
- [83] Steven Maynard-Moody and Michael Musheno. 2003. Cops, teachers, counselors: Narratives of street-level judgment. *Ann Arbor, MI: University of Michigan Press* (2003).
- [84] Gary B Melton. 2005. Mandated reporting: A policy without reason. *Child abuse & neglect* 29, 1 (2005), 9–18.
- [85] Michael E Milakovich and George J Gordon. 2013. *Public administration in America*. Cengage Learning.
- [86] Terry D Moore, Thomas P McDonald, and Kari Cronbaugh-Auld. 2016. Assessing risk of placement instability to aid foster care placement decision making. *Journal of Public Child Welfare* 10, 2 (2016), 117–131.
- [87] Shabnam Mousavi and Gerd Gigerenzer. 2014. Risk, uncertainty, and heuristics. *Journal of Business Research* 67, 8 (2014), 1671–1678.
- [88] Henrik Mucha, Sebastian Robert, Rüdiger Breitschwerdt, and Michael Fellmann. [n.d.]. Towards Participatory Design Spaces for Explainable AI Interfaces in Expert Domains. ([n. d.]).
- [89] US Department of Health, Human Services, et al. 2020. The AFCARS Report. Preliminary FY 2019 Estimates as of June 23, 2020 (27). Retrieved December 23 (2020), 2019.
- [90] Stephen P Osborne and Louise Brown. 2013. *Handbook of innovation in public services*. Edward Elgar Publishing.
- [91] Juho Pääkkönen, Matti Nelimarkka, Jesse Haapoja, and Airi Lampinen. 2020. Bureaucracy as a Lens for Analyzing and Designing Algorithmic Systems. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [92] Samir Passi and Steven Jackson. 2017. Data vision: Learning to see through algorithmic abstraction. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. 2436–2447.
- [93] Samir Passi and Steven J Jackson. 2018. Trust in data science: collaboration, translation, and accountability in corporate data science projects. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–28.
- [94] Samir Passi and Phoebe Sengers. 2020. Making data science systems work. *Big Data & Society* 7, 2 (2020), 2053951720939605.
- [95] John W Payne, James R Bettman, and Eric J Johnson. 1993. The use of multiple strategies in judgment and choice. (1993).
- [96] Rik Peeters. 2020. The agency of algorithms: Understanding human-algorithm interaction in administrative decision-making. *Information Polity Preprint* (2020), 1–16.
- [97] B Guy Peters. 2014. Accountability in public administration. *The Oxford handbook of public accountability* (2014), 211–225.
- [98] Hitendra Pillay and Andrea R McCrindle. 2005. Distributed and relative nature of professional expertise. *Studies in Continuing Education* 27, 1 (2005), 67–88.
- [99] Anthony T Pinter, Pamela J Wisniewski, Heng Xu, Mary Beth Rosson, and Jack M Carroll. 2017. Adolescent online safety: Moving beyond formative evaluations to designing solutions for the future. In *Proceedings of the 2017 Conference on Interaction Design and Children*. ACM, 352–357.
- [100] Anja Pors and Jannick Schou. 2020. Street-level morality at the digital frontlines: An ethnographic study of moral mediation in welfare work. *Administrative Theory & Praxis* (2020), 1–18.
- [101] Joanna Redden, Lina Dencik, and Harry Warne. 2020. Datafied child welfare services: unpacking politics, economics and power. *Policy Studies* (2020), 1–20.
- [102] Jeanne S Ringel, Dana Schultz, Joshua Mendelsohn, Stephanie Brooks Holliday, Katharine Sieck, Ifeanyi Edochie, and Lauren Davis. 2018. Improving child welfare outcomes: balancing investments in prevention and treatment. *Rand health quarterly* 7, 4 (2018).
- [103] Samantha Robertson, Tonya Nguyen, and Niloufar Salehi. 2021. Modeling Assumptions Clash with the Real World: Transparency, Equity, and Community Challenges for Student Assignment Algorithms. *arXiv preprint arXiv:2101.10367* (2021).
- [104] Samantha Robertson and Niloufar Salehi. 2020. What If I Don't Like Any Of The Choices? The Limits of Preference Elicitation for Participatory Algorithm Design. *arXiv preprint arXiv:2007.06718* (2020).
- [105] Jodi R Sandfort. 2000. Moving beyond discretion and outcomes: Examining public management from the front lines of the welfare system. *Journal of Public Administration Research and Theory* 10, 4 (2000), 729–756.
- [106] Devansh Saxena, Karla Badillo-Urquiola, Pamela Wisniewski, and Shion Guha. 2020. Child Welfare System: Interaction of Policy, Practice and Algorithms. In *Companion of the 2020 ACM International Conference on Supporting Group Work*. 119–122.
- [107] Devansh Saxena, Karla Badillo-Urquiola, Pamela J Wisniewski, and Shion Guha. 2020. A Human-Centered Review of Algorithms used within the US Child Welfare System. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–15.

- [108] Devansh Saxena, Erhardt Graeff, Shion Guha, EunJeong Cheon, Pedro Reynolds-Cuellar, Dawn Walker, Christoph Becker, and Kenneth R Fleischmann. 2020. Collective Organizing and Social Responsibility at CSCW. In *Conference Companion Publication of the 2020 on Computer Supported Cooperative Work and Social Computing*. 503–509.
- [109] Devansh Saxena and Shion Guha. 2020. Conducting Participatory Design to Improve Algorithms in Public Services: Lessons and Challenges. In *Conference Companion Publication of the 2020 on Computer Supported Cooperative Work and Social Computing*. 383–388.
- [110] Philipp Schmidt and Felix Biessmann. 2020. Calibrating Human-AI Collaboration: Impact of Risk, Ambiguity and Transparency on Algorithmic Bias. In *International Cross-Domain Conference for Machine Learning and Knowledge Extraction*. Springer, 431–449.
- [111] Craig Schwalbe. 2004. Re-visioning risk assessment for human service decision making. *Children and Youth Services Review* 26, 6 (2004), 561–576.
- [112] Nick Seaver. 2017. Algorithms as culture: Some tactics for the ethnography of algorithmic systems. *Big Data & Society* 4, 2 (2017), 2053951717738104.
- [113] Aron Shlonsky and Dennis Wagner. 2005. The next step: Integrating actuarial risk assessment and clinical judgment into an evidence-based practice framework in CPS case management. *Children and youth services review* 27, 4 (2005), 409–427.
- [114] Ravi Shroff. 2017. Predictive Analytics for City Agencies: Lessons from Children’s Services. *Big data* 5, 3 (2017), 189–196.
- [115] Elizabeth A Smith. 2001. The role of tacit and explicit knowledge in the workplace. *Journal of knowledge Management* (2001).
- [116] Michael R Sosin. 2010. Discretion in human service organizations. *Human services as complex organizations* (2010), 381–403.
- [117] Megan Stevenson. 2018. Assessing risk assessment in action. *Minn. L. Rev.* 103 (2018), 303.
- [118] Brian J Taylor and Karen McQuillan. 2014. Perspectives of foster parents and social workers on foster placement disruption. *Child Care in Practice* 20, 2 (2014), 232–249.
- [119] James Topitzes, Timothy Grove, Erika E Meyer, Stacey M Pangratz, and Caitlin M Sprague. 2019. Trauma-responsive child welfare services: A mixed methods study assessing safety, stability, and permanency. *Journal of Child Custody* 16, 3 (2019), 291–312.
- [120] Marcel F van Assen. 2021. Training, employee involvement and continuous improvement—the moderating effect of a common improvement method. *Production Planning & Control* 32, 2 (2021), 132–144.
- [121] Michael Veale, Max Van Kleek, and Reuben Binns. 2018. Fairness and accountability design needs for algorithmic support in high-stakes public sector decision-making. In *Proceedings of the 2018 chi conference on human factors in computing systems*. 1–14.
- [122] Eran Vigoda-Gadot and Itai Beeri. 2011. Change-oriented organizational citizenship behavior in public administration: The power of leadership and the cost of organizational politics. *Journal of Public Administration Research and Theory* 22, 3 (2011), 573–596.
- [123] Paul Voigt and Axel Von dem Bussche. 2017. The eu general data protection regulation (gdpr). *A Practical Guide, 1st Ed., Cham: Springer International Publishing* (2017).
- [124] Sandra K Welch. 2008. *A metasynthesis of the transition from novice to expert: Can instructional interventions shorten the process?* Ph.D. Dissertation. Capella University.
- [125] Melody Williams. 2014. Child welfare in Milwaukee, Wisconsin. (2014).
- [126] Ben Williamson. 2016. Digital education governance: data visualization, predictive analytics, and ‘real-time’ policy instruments. *Journal of Education Policy* 31, 2 (2016), 123–141.
- [127] C Wilson. 2014. Integrating safety, permanency, and well-being: A view from the field. *Integrating safety, permanency and well-being series* (2014).
- [128] Matthew M Young, Justin B Bullock, and Jesse D Lecy. 2019. Artificial discretion as a tool of governance: a framework for understanding the impact of artificial intelligence on public administration. *Perspectives on Public Management and Governance* 2, 4 (2019), 301–313.
- [129] Kristen M Zgoba and Nishi P Dayal. 2015. Recidivism. *The Encyclopedia of Crime and Punishment* (2015), 1–5.
- [130] Haiyi Zhu, Bowen Yu, Aaron Halfaker, and Loren Terveen. 2018. Value-sensitive algorithm design: Method, case study, and lessons. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–23.
- [131] Marc A Zimmerman. 2013. Resiliency theory: A strengths-based approach to research and practice for adolescent health.

A CODEBOOK FOR RQ1

Algorithms and Associated Outcomes	Illustrative Quotation
Child and Adolescent Needs and Strengths (CANS)	
Compensation Calculations (85%, n=17)	"CANS has become all about the rate with little attention paid to actually helping the child. A foster parent said to me.. this is more like a \$1600 kid"
Mental Heath Assessment (65%, n=13)	"CANS scores translate to actionable steps in form of mental-health services for the child"
Level of Foster Care (50%, n=10)	"It recommends the level of foster care for the child, but that decision comes down to the availability of placements"
Seven Essential Ingredients (7ei)	
Trauma-informed care (90%, n=18)	"7ei allows us to center every case in the trauma-responsive model and the scores show that it actually is leading to better placement stability and permanency outcomes"
Placement Stability and Permanency (85%, n=17)	
Anti Sex-Trafficking Response Tool	
Monitor case for sex-trafficking (65%, n=13)	"It helps caseworkers be actively aware of and look for sex-trafficking indicators that we wouldn't have thought of 10 years ago"
Report case to HART (65%, n=13)	
Legal Permanency Status Tool	
Quality of current placement (60%, n=12)	"It helps us be actively aware of the current status of the placement and systemic barriers, and what we need to do to achieve permanency"
Identify systemic barriers (50%, n=10)	

Table 2. RQ1: Associated Algorithms and Outcomes. Percentages (%) represent the proportion of interview participants who recognized and discussed the algorithmic outcome.

B PARTICIPANT PROFILES

We interviewed participants who were consistently present at team meetings or who had more experience working with algorithmic systems. While on-site, we interacted with approximately 120 employees and external consultants of the agency. We did not collect demographic information from all meeting participants as we did not want to disrupt their work as it was being carried out. We did, however, collect the gender, job title, education-level, years of experience of our interviewees, and the algorithms they used in daily work practices as shown in Table 3:

- **Child-Welfare Supervisors (n=8):** Supervise a case management teams consisting of six to eight case managers. All supervisors were former case managers who now provide individual supervision to case managers and ensure their professional development. They facilitate the planning meetings and supervise about 140 cases each.
- **Permanency Consultants (n=5):** Responsible for permanency planning, which includes the drafting and completion of documents for court and initiating permanency legal process, targeted recruitment for adoption, provision of post-guardianship services and/or post-adoption services for families. On average, they provide direct consultations on about 150 cases each.
- **Child Welfare Case Managers (n=2):** The front-line workers who directly interact with families and are the point of contact between parents, CWS, and the court system. They conduct home visits, support and monitor the safety and well-being of children, document safety assessments, facilitate improving parental protective capacities by creating goals with families etc. On average, they manage about 20 cases.
- **Data Specialists (n=2):** Track case-level data for families and are also responsible for performance benchmarks at the agency as mandated by federal legislation. Data specialists also analyze data about the algorithms being used at the agency and present findings to agency leadership.

- **Child Welfare Program Directors (n=1):** Supervise and manage the child-welfare supervisors. Program directors also manage the training and professional development programs, research projects, as well as policy development initiatives.
- **Permanency Consultation Supervisor (n=1):** Supervise and manage the permanency consultants and facilitate permanency consultations.
- **Clinical Therapist (n=1):** Licensed clinical social workers who conduct mental-health assessments for foster children.

Since CWS experiences a high turnover rate in the case manager position, we focused our attention more so towards supervisors and permanency consultants when conducting our interviews. Moreover, supervisors facilitate planning meetings and permanency consultants facilitate permanency consultations, which ensured that these staff members were consistently present at all the meetings where collaborative work mediated by policies, social work practice, and algorithms occurred.

Participant	Sex	Job Title	Education	Experience (years)	Algorithms
P1	F	Permanency Consultation Supervisor	MSW, APSW	22	7ei, AST, LPS
P2	F	Permanency Consultant	MSW	20	7ei, AST, LPS
P3	F	Permanency Consultant	MSW, APSW	9	7ei, AST, LPS
P4	F	Permanency Consultant	MSW, APSW	8	7ei, AST, LPS
P5	F	Permanency Consultant	MSW	12	7ei, AST, LPS
P6	F	Permanency Consultant	BSW	3	7ei, AST, LPS
P7	F	Child Welfare Program Director	MSW	20	7ei, CANS
P8	F	Child Welfare Supervisor	MSW, CAPSW	19	CANS, 7ei, AST
P9	F	Child Welfare Supervisor	MSW	13	CANS, 7ei, AST
P10	F	Child Welfare Supervisor	MSW, LPC, NCC	9	CANS, 7ei, AST
P11	F	Child Welfare Supervisor	MSW, APSW	9	CANS, 7ei, AST
P12	F	Child Welfare Supervisor	MSW	7	CANS, 7ei, AST
P13	M	Child Welfare Supervisor	MSW	12	CANS, 7ei, AST
P14	M	Child Welfare Supervisor	MSW	30	CANS, 7ei, AST
P15	F	Child Welfare Supervisor	MSW	9	CANS, 7ei, AST
P16	F	Clinical Therapist	LCSW	5	CANS, 7ei, AST
P17	F	Child Welfare Case Manager	BSW	8	CANS, 7ei, AST
P18	M	Child Welfare Case Manager	BSW	2	CANS, 7ei, AST
P19	F	Data Specialist (Program Director)	MSW	17	CANS, 7ei, AST
P20	M	Data Specialist	MSW	17	CANS, 7ei, AST

BSW: Bachelor of Social Work
LCSW: Licensed Clinical Social Worker
APSW: Advanced Practice Social Worker
NCC: National Certified Counselor

MSW: Master of Social Work
LPC: Licensed Professional Counselor
CAPSW: Certified Advanced Practice Social Worker

Table 3. Interview Study Participants

Received January 2021; revised April 2021; accepted May 2021