Participatory Algorithmic Management for Worker Well-Being

Min Kyung Lee,[1] Ishan Nigam,[2] Angie Zhang,[1] Joel Afriyie,[2] Zhizhen Qin,[3] Sicun Gao[3]

[1] School of Information, UT Austin
[2] Computer Science Department, UT Austin
[3] Computer Science and Engineering, UC San Diego
minkyung.lee@austin.utexas.edu, ishann@cs.utexas.edu, angie.zhang@austin.utexas.edu, joel.afriyie@austin.utexas.edu, zhq005@ucsd.edu, sicung@ucsd.edu

Abstract

Increasingly more workplaces are managed by algorithms that handle scheduling, task assignment, and matching functions. Algorithms promise efficient streamlined results, but emerging evidence suggests that algorithmic management often undermines worker well-being. How can we incorporate worker well-being into algorithmic management so that workplaces can also be optimized for workers? As a first step toward this goal, we propose a tool to enable workers to build their own well-being models. In this tool, workers specify their preferences about tasks and schedules in their work preference models and their beliefs about what rules their organization should use to make trade-off decisions when managing workers. As a case study, we evaluated the tool in the context of shift scheduling, which impacts 20% of working adults worldwide. We studied how 25 shift workers built their well-being models and interviewed three shift managers to understand their perspectives about the tool. The findings illuminate the opportunities and challenges in defining worker well-being for algorithmic management.

participatory design, worker well-being, preference elicitation, algorithmic work management

Introduction

Artificial Intelligence (AI) is increasingly being used to manage the workforce. A wide range of mid-level managerial decisions—such as task assignment and matching, team formation, and performance evaluation—are automated or assisted by algorithms in various workplaces: crowdsourcing (Gray and Suri 2019) and on-demand transportation platforms (Lee et al. 2015; Wood et al. 2019), as well as in offline workplaces such as warehouses (McClelland 2012), call centers (Roose 2019), and offices (Mateescu and Nguyen 2019). Dynamic scheduling based on predictive demand (Kantor 2014) significantly impacts the lives of shift workers (Schneider and Harknett 2019), which comprise 20% of the workforce worldwide and 25% of Americans (Eurofound 2012; Boivin and Boudreau 2014; Lieberman et al. 2020). AI integration is based on the promise that algorithmic management will boost workplace efficiency and economic value (Kellogg, Valentine, and Christin 2020).

However, alarming evidence suggests algorithmic management can undermine worker well-being. Numerous reports show that warehouse workers are under serious physical and psychological stress due to task assignment and tracking without appropriate break times (McClelland 2012); Uber and Lyft drivers feel automated evaluation is unfair and distrust the system's opaque payment calculations (Hawkins 2020; Lee et al. 2015; Lee 2018; Wood et al. 2019); shift workers suffer from unpredictable schedules that destabilize work-life balance and disrupt their ability to plan ahead (Schneider and Harknett 2019). There is growing recognition that worker well-being must be considered when designing a workplace that integrates AI, and guidelines for achieving this goal have been proposed (PAI 2020).

We argue it is critical to computationally model worker well-being and directly incorporate it into algorithmic workplace design. Extensive research and industry efforts have investigated individual modeling and personalization to model consumption preferences and tailor online environments (Sarwar et al. 2001). Worker well-being models will enable designers to use well-being as an optimization goal, personalize for individual workers, and monitor the effects of algorithmic management on worker well-being.

As a first step toward this goal, we propose a participatory method to model worker well-being, adopting a participatory algorithmic governance framework. We first define the foundational dimensions of worker well-being for two types of models: work preference models for work conditions that impact physical, psychological, and financial well-being, and managerial fairness models for organizational rules that manage workers and tradeoff decisions. We built a web-tool where workers could build their well-being models. As a case study, we applied our method to shift work scheduling. We conducted formative interviews with 9 shift workers and 2 managers to inform shift work well-being models. We then worked with 25 shift workers who used our web-tool to build their well-being models and shared their experience with us through interviews. Using the elicited worker well-being models, we also conducted interviews with 3 shift scheduling managers. Our work makes a contribution to a growing body of literature on algorithmic work and management by offering a method that centers on worker well-being.

Working Paper

Worker Well-Being in Algorithmic Management

Algorithmic Management of Human Workers

AI is transforming multiple functions in workplaces, changing communication, collaboration, training, and workforce management (PAI 2020; Mateescu and Nguyen 2019; Lee et al. 2015). A driving force of the adoption of AI is the projected economic value of streamlined coordination and task execution, and improved accuracy due to data-driven insights (Kellogg, Valentine, and Christin 2020). In this paper, we focus on the use of AI in workforce management tasks such as scheduling, performance evaluation, and work assignment. The concept of automated workforce management is not new, and features such as automated scheduling and team formation have long been part of modern workplaces (Van den Bergh et al. 2013). However, recent advances in AI are expanding the scale of these tools, the numbers of workers affected, and the nature of automation decisions. For example, algorithmic scheduling tools now utilize dynamic forecasting, such as predicted customer demand, which makes schedules more unpredictable from workers' perspectives (Kantor 2014). Recent research has highlighted how shifting schedules negatively impact workers' health (Schneider and Harknett 2019). Recent research on on-demand work platforms has also discovered how algorithms are used as a control mechanism, resulting in negative impacts on worker well-being. For example, algorithmic management in on-demand transportation platforms can result in financial uncertainty and insecurity, lack of autonomy, and sleep deprivation for drivers (Wood et al. 2019). Research on crowd workers also highlights how little control crowd workers have in the arrangement of their work, and that their work assignment and evaluation is done with little input or consideration of workers' preferences and contexts (Gray and Suri 2019).

Designing worker-centered workplaces has been a central topic of research in human-computer interaction. Early research on participatory design has sought to involve workers in designing work policies and environments, and giving them a say in management practices (Ehn 1988; Bjerknes et al. 1987). Khovanskaya et al. proposed adapting union tactics for algorithmic platform workers, such as requesting data transparency or contesting wage decisions (Khovanskaya et al. 2019). Anya proposed ways to design worker-centered crowdsourcing task and environment design, leveraging the literature on job design for crowd workers' well-being (Anya 2015). Unde et al. studied what scheduling norms are perceived as fair when scheduling nurses (Uhde et al. 2020). Most recently, scholars called for the inclusion of worker well-being as an explicit focus when algorithmic work is designed (PAI 2020). To our knowledge, no research has explored methods of computationally defining worker well-being. In this work, we propose a method for eliciting worker well-being preferences in algorithmic management, so that worker well-being can be used to optimize or evaluate algorithmic management.

Preference Elicitation

While eliciting workers' preferences for their well-being has not been the focus of prior work, a long line of research has investigated how to model individual preferences. Preference elicitation is a methodology for understanding individual valuations of goods or services. In the past, preference elicitation was utilized to measure the economic value of goods. In more recent years, preference elicitation has been used in subjective evaluations of goods and services, such as determining which environmental sustainability strategies people prefer or which services contribute to people's satisfaction and happiness. Preferences can be explicitly stated or revealed from observed behaviors (Louviere, Hensher, and Swait 2000; Johnston et al. 2017). In the case of worker well-being, there is little data to infer worker well-being preferences, so in this paper, we focus on elicitation methods for stated preferences. The most widely-used preference elicitation method is discrete choice experiments (McFadden 1986). In this method, people are asked to make choices between two or more discrete alternatives where at least one attribute of the alternative is systematically varied. Other methods include a ranking method that asks people to rank alternatives in the order of their preferences and a matching method where people are asked to state their "willingness to pay" to obtain a particular good (Ali and Ronaldson 2012). In our paper, we adopt the ranking and discrete choice methods to elicit worker's preferences for their well-being.

Worker Well-Being Preference Model

Well-being is an umbrella term that describes "what it means to be functioning as a healthy person across multiple domains" (Pressman, Kraft, and Bowlin 2013). The literature on well-being is vast and covers several different dimensions, such as subjective well-being and social well-being. In our work, we focus on constructs of work-related well-being. There are different requirements for working conditions and worker well-being such as minimum wage, maximum shift duration and required breaks in each state in the U.S. Every worker also has their individual preferences for their own well-being; the same tasks or shifts could influence well-being differently depending on the worker's aptitude, life situation and goals. For example, some workers might appreciate long work hours for financial reasons, whereas others might prefer shorter shifts in order to preserve work-life balance. In a smaller workplace, a manager might know every worker's individual preferences when assigning tasks and schedules. Our goal is to learn and codify workers' well-being preferences to personalize algorithmic management, even in larger organizations. In this section, we present well-being related preference constructs in work preference and managerial fairness models.

Work Preference Model

Work preference models should capture preferences that are related to the work's impact on physical, psychological, and financial well-being.

Physical well-being means "the ability to perform physical activities and carry out social roles that are not hin-

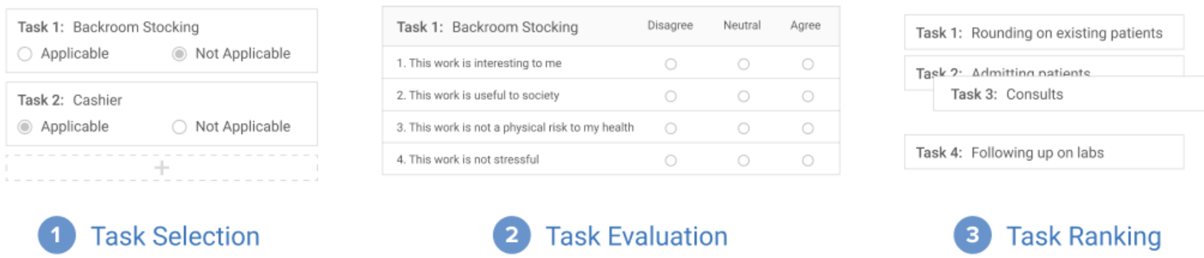

Figure 1: **Ranking-based elicitation for task preference model**. (1) The worker selects relevant tasks. (2) The worker provides inputs on their evaluations for each task. (3) The worker ranks the tasks according to their preferences.

dered by physical limitations and experiences of bodily pain, and biological health indicators" (Capio, Sit, and Abernethy 2014). The same tasks or work conditions can impact workers' physical well-being differently. For example, tasks that involve frequent social interaction or varying degrees of physical strength could be acceptable or even enjoyable for certain workers, but cause stress for workers who are introverted or physically weak. Another dimension of worker well-being is work-life balance (Guest 2002; PAI 2020), as work can hinder workers' abilities to carry out social roles outside of work. This is particularly relevant to shift scheduling and work hours. Shift work, particularly in the service industry, used to primarily employ younger adults with more schedule flexibility; however, the age range of shift workers has expanded in recent years, resulting in shift workers who need to handle family obligations.

Psychological well-being, also referred to as emotional or intellectual well-being, means "the combination of feeling good and functioning effectively" (Huppert 2009). In the context of work, the following factors contribute to psychological well-being (Morgeson and Humphrey 2006; Richard and Oldham 1976): autonomy in work, or how much control workers have over what tasks to carry out and how; meaning derived from work, or whether the work helps others or reinforces workers' identity; and enjoyment in work, or whether workers find the work pleasant and feel the tasks offer interesting challenges and learning opportunities. As with physical well-being, the same tasks or work conditions can impact workers' psychological well-being differently. For example, some workers prefer challenging tasks with more autonomy whereas other workers prefer well-defined, repetitive tasks.

Financial well-being refers to "the perception of being able to sustain current and anticipated desired living standards and financial freedom" (Brüggen et al. 2017). In the context of work, financial well-being depends on income adequacy—whether workers can cover their expenses and pay bills—and income volatility—whether workers earn regular incomes (Schneider and Harknett 2019). We expect this is the most important dimension for most workers, which means less individual level variations in their preferences. Still, sources of individual difference might include whether the worker depends on the job for primary or supplementary income, and how well they can tolerate income volatility.

Managerial Fairness Model

Another important dimension of worker well-being is the perceived fairness of their workplace managerial model. Worker well-being is greatly influenced by the perceived fairness of supervisors and management, with fair and ethical management highly correlated with worker well-being (Sparr and Sonnentag 2008; Fujishiro 2005; Hoppe, Heaney, and Fujishiro 2010). Managerial preference models capture what managerial rules workers deem fair when allocating work to multiple workers. Three dominant allocation principles exist (Daverth, Cassell, and Hyde 2016): the *equality* principle, which holds that "everyone should receive the same allocations regardless of performance or other contingencies"; the *equity* principle, which assumes that rewards and resources should be allocated based on merit, such as the workers' contributions to the organization; and the *need* principle, which argues that allocation should be based on individual circumstances, prioritizing those in the most need of a resource regardless of their input and output. Workers in different workplaces may perceive one or a combination of these principles as fair.

Worker Well-Being Model Elicitation

To elicit worker well-being preferences for algorithmic management, we first need to identify the workers and resources that are managed in the workplace, and the range of alternatives associated with those resources, such as types of tasks or shifts. Using the alternatives, we propose two methods that can elicit worker preferences, drawing from literature on preference elicitation and adopting a participatory framework for algorithmic governance (Lee et al. 2019b).

Ranking-based elicitation In this method, workers first evaluate each of the resource alternatives by answering several questions about the alternatives' impact on their well-being, and then rank them in the order of their preference (Figure 1).Based on the worker well-being and working condition papers (Morgeson and Humphrey 2006; Maestas et al. 2017), we derived nine questions that measure the tasks' impact on workers' physical and psychological well-being; furthermore, in our case study we describe these well-being

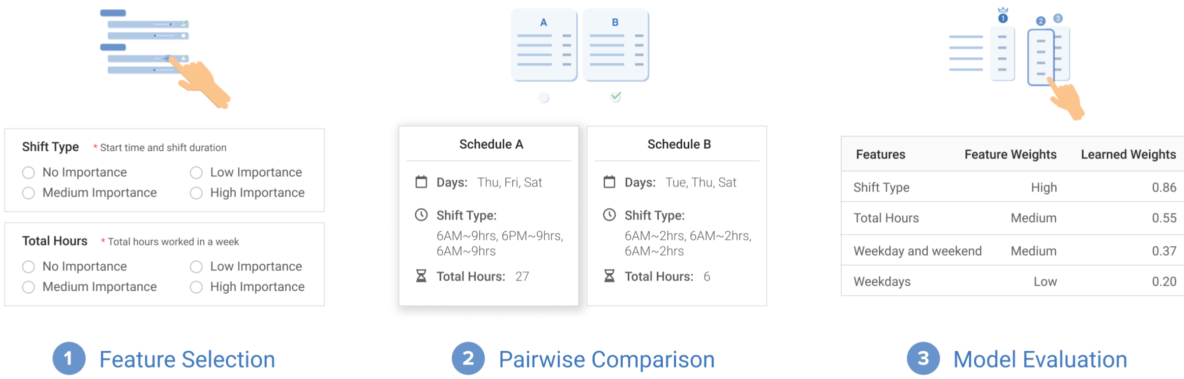

Figure 2: **Feature weight- & pairwise comparison-based elicitation for schedule preference and managerial fairness models**. (Left) The worker chooses a set of relevant features. (Center) The worker provides preferences by comparing pairs of hypothetical scenarios. (Right) The worker evaluates the AI model learned from the pairwise comparison responses.

questions by covering motivational characteristics, associated stress and physical risk, wage, and worker autonomy. These questions prepare the workers to evaluate the tasks holistically considering a range of impacts. This method is time-efficient and appropriate when workers already have well established preferences about the alternatives; thus, they can accurately express their preferences by ordering the alternatives. Additionally, answers to well-being questions can be used to measure and monitor the impact of workplace conditions on worker well-being. In our case study, we use this method to elicit workers' task preferences.

Feature weight- & pairwise comparison-based elicitation This method adopts the individual belief modeling part of WeBuildAI, a participatory algorithmic governance framework (Lee et al. 2019b). In this framework, users answer pairwise comparisons of alternatives to train an algorithm or explicitly specify weights for features (Figure 2). Workers first specify which characteristics (features) of the resources matter to them, and how important each feature is (explicit feature weights). Workers then answer a series of pairwise comparison questions about the alternatives. Their responses to the pairwise comparisons inform their individual feature weights (learned feature weights). In the final step, workers review the explicit and learned feature weights, the accuracy of the trained model, and an example list of ordered alternatives produced by the trained model. This method is suitable when it is difficult to select a manageably-sized set of representative alternatives for ranking. More importantly, pairwise comparisons are helpful when workers do not yet have well-formed, stable preferences and need to discover their preferences. In our case study, we use this method to elicit workers' preferences about different shift schedules (schedule preference model) and managerial rules that assign different shifts to workers (managerial fairness model).

Case Study: Shift Work Scheduling

Our case study applies the participatory well-being elicitation method in the context of shift work scheduling. We chose shift work because shift workers comprise 20% of working adults, and recent research suggests worker well-being is compromised due to limited agency over schedules and on-demand dynamic scheduling (Schneider and Harknett 2019). Our ultimate goal is to use shift workers' well-being models as part of optimization objectives in addition to organizational constraints such as cost and other resources, and use them to measure different schedules' impact on worker well-being.

Impact of Work Scheduling on Worker Well-Being

Precarious work is a type of labor often characterized by uncertainty of working hours, lack of control by workers, and low wages (Kalleberg 2009; Organization 2011). A form of precarious work, shift work, employs 25% of workers in the United States (Lieberman et al. 2020). Shift work has frequently been studied for its adverse effects on workers as related to low wages, unpredictability of hours, and the resulting effects on financial security (Lambert 2008; Finnigan 2018; Wickwire et al. 2017; Lambert, Henly, and Kim 2019). Recently, researchers have also been investigating the temporal dimension of shift work, namely "predictability and stability of work hours" (Schneider and Harknett 2019), in order to better understand its effects on worker well-being. Many companies turn to staffing strategies that reflect on-demand practices which minimize costs by matching the number of shift workers with forecasted demand (Alexander and Tippett 2017; Lambert, Henly, and Kim 2019; Schneider and Harknett 2019). However, this optimization practice for management efficiency often leads to schedule instability and compromised worker well-being. Workers commonly have their shifts extended, shortened, cancelled, or added—all with less than 24 hour notice—or are even assigned variable scheduling and work morning, day, and evening shifts within one week (Schneider and Harknett 2019). Shift workers also experience high levels of inconsistency in the hours and days they work each week (Lambert 2008; Henly and Lambert 2014). As most shift workers are paid hourly, any unpredictability of how much work they receive is directly correlated with financial stability (Golden 2015). Additionally, the inconsistency of which days they work prevents workers from freely planning non-work commitments (Henly and Lambert 2014). These conditions culminate in circumstances where shift workers have to defer control over their own scheduling and consequently, their well-being (Lambert 2008; Alexander and Tippett 2017; Ananat and Gassman-Pines 2021).

Shift Worker Well-Being Model Design

To model workers' well-being for shift work, we first need to know what factors of shift work schedules influence worker well-being. To gain insights on this, we conducted formative interviews with shift workers and scheduling managers.

Formative interviews with shift workers & managers We conducted one hour interviews with nine shift workers and two scheduling managers[1] employed in the fast food, retail, and healthcare industries to inform our design of the default preference features. The interview questions focused on understanding current scheduling practices, how current scheduling/schedules supports or hinders workers' goals, and factors that should be considered when scheduling multiple workers. We also complemented our interview findings with a report that surveyed shift managers' scheduling practice (Lambert 2015).

All workplaces of our interviewees used workers' availability, their cost budget —i.e. the total number of worker hours—, and the minimum number of workers to successfully complete each shift. Some manually created schedules; others used scheduling software. Relationships between scheduling managers and workers varied. In a smaller workplace, the scheduling manager directly interacted with workers on the floors. In a bigger workplace, workers interacted with middle managers without any interaction opportunities with the scheduling manager. The degree to which individual workers' preferences were considered in schedule creation largely depended on the presence of interaction. Some scheduling managers directly ask their workers to share their preferences. In other workplaces, there was no formal opportunity or system to solicit a worker's preferences. Workers who are personally close to their shift managers may directly convey their preferences. This may lead to these shift workers receiving better schedules, as some managers consider their workers preferences and needs that they learn through personal interaction (i.e, a worker who is a single mom and may have greater financial needs). In workplaces where shift managers do not interact with workers, there was often no way for workers to convey their preferences as the methods workers used to indicate their availability also do not solicit preference information.

Task preferences Worker interviews informed the creation of task preference input in our tool. From speaking to them, it was clear they performed an array of tasks and held different opinions about each one. We recognized that there was a physical and mental toll that these tasks uniquely took on the workers, and felt that accounting for task preferences could improve employee well-being if a balance in preferences could be satisfied. We thus extended our design to assess how each task impacted their well-being and asked them to order their tasks by preference. The questions included: Interest in task, perceived societal usefulness, desirable amounts of social interaction, perceived physical risk to health, stress associated with task, opportunities for high earnings, opportunities for career advancement, job security, and opportunities for independent decision making (Kovach et al. 1987; Karl and Sutton 1998; Morgeson and Humphrey 2006; Tortia 2008).

Schedule preference In line with previous work on precarious scheduling (Henly and Lambert 2014), we learned from workers that consistent schedules promoted predictability in both work schedule and income, which allowed for a healthy work-life balance. For this reason, we created two across week features, *same-days* each week and *same-number-of-days* each week. We also learned that not all available hours were equally favored by workers: they held strong preferences for morning vs. afternoon vs. evening shifts, which could vary depending on the day. Shifts assigned at less preferred times can negatively impact their sleep cycle and leisure-time. Thus, we implemented the *shift-type* feature—the collection of unique day, shift start time, and duration combinations—that workers input as their preferred time. Additionally, workers held well-formed preferences about which days they worked. For example, student workers often had less availability during the week and wished to maximize working on weekends. Alternatively, some preferred to maximize weekdays to spend weekends with family. We added working *weekdays* only and working *weekdays-and-weekends* to account for these situations. *Total-hours* was included as many workers expressed maximizing total hours was a crucial consideration.

Managerial fairness preference To create the managerial fairness features, we explored the types of practices that shift managers employ when they assign schedules. Lambert's scheduling study with shift managers (Lambert 2015) indicates that the majority of managers agreed that they gave more hours to workers who were good at sales and were reliable. This informed the creation of the *reliable* feature for reliable workers and the *performer* feature to indicate high performing workers. Though the report found that managers did not typically consider workers' financial needs when allocating hours, exploratory interviews with workers indicated that financial need formed an important aspect of their well-being - resulting in the design of three features to address this: (*fewer-hours, limited-availability, fewer-preferred-shifts*). We also learned from exploratory interviews that seniority played a large role in determining schedules at their workplace, and created a *seniority* feature accordingly. For our seventh feature, we incorporated *volunteering* after a scheduling manager explained that when she was unexpectedly shorthanded, she would first ask if anyone was willing to volunteer to work a shift.

[1]We clarify that there is no overlap between the formative interview participants and the evaluative study participants reported later in the paper.

Elicitation Tool for Shift Worker Well-Being Model

Task evaluation and ranking In order to model task preferences, the web-tool first asks workers to submit tasks that they can be assigned. Workers then evaluate each task by answering questions about the task's impact on physical and psychological well-being (Morgeson and Humphrey 2006; Maestas et al. 2017; Schneider and Harknett 2019). After evaluating each task, workers ranked the tasks where the first task was their most preferred task (Figure 1).

Schedule & managerial fairness model learning We modeled scheduling and managerial fairness preferences using the features described in Table 1. Using these features, we allow the workers to explicitly express their preferences by selecting No Importance, Low, Medium, High - for each of the features. The features chosen with No Importance level are removed from the analysis. We then proceed towards modeling the preferences through statistical learning.

We perform binary discrete choice experiments to model schedule preferences by asking workers to choose between hypothetical schedules. Each pair of hypothetical schedules is procedurally generated based upon fractional factorial design principles (Johnson et al. 2013). The manager fairness preferences are similarly modeled through binary discrete choice experiments where the participants are tasked with choosing between hypothetical workers. Each feature is represented as a binary variable. The shift worker's responses are used to learn a logistic regression model based on random utility theory (McFadden 1986). Model learning is initialized from a normal prior distribution and maximum likelihood estimation is performed using the `BFGS` method.

Model summary & evaluation. The worker is shown a summary of the learned preference model and asked to evaluate its performance. The overall model accuracy is shown followed by a comparison of each of their explicit preference responses with each of the preference weights learned by the model. The worker provides feedback on whether their explicit responses or the learned model weights better reflect their preferences. The worker also evaluates a set of five hypothetical scenarios ranked by the regression model. The worker provides feedback on how well the learned model captures their preferences. We conclude the evaluation by allowing the worker to provide open-ended feedback about the performance of the learned preference model as well as the process of creating the preference model.

Method

Participants We recruited 25 shift workers and 3 shift scheduling managers (Table 2). To recruit shift workers, We posted a recruiting message on shift work related Reddit threads and ran Facebook ads. Interested participants filled out a screening survey about their current employment as a shift worker, demographics, scheduling practices in their workplace, and an attention check question. To recruit shift work scheduling managers, we ran Facebook ads that directed participants to a screening survey. The survey asked for their employer and whether they currently create employee schedules, which was used as an eligibility criteria.

Our shift worker participants' average age was 29.12 years (SD=10.7; Min-Max: 18-48). 16 participants were female. Our participants included 9 white, 8 Latinx, 4 Asian, 2 American Indian/Alaska Native, 1 Black, and 1 Other [Middle Eastern]. Our participants spanned a variety of industries, with the fast food industry being the plurality (11 workers). 24 of our 25 participants provided responses to optional education and income bracket questions. 10 participants were enrolled in school; the highest education level completed of the other 14 respondents varied from some high school to receiving a college degree (associate or higher). The reported average hourly wage of our participants was $14.07 per hour.

Our shift scheduling managers were White, Black, and Asian. Two managers were female and one was male. All of them were employed in fast food industry. Full participant information is reported in (Tables 2 and 3) in the Appendix.

Elicitation tool evaluation with shift workers We asked our worker participants to use our web-tool to build their shift work well-being models. Participants visited our web-tool and screen-shared over video conference calls so the interviewer could observe interactions with the tool. Participants interacted with the tool at their own pace and were encouraged to think aloud or ask for clarifications. The interviewer observed the interactions and asked questions about the participant's thought process several times throughout the interview. For example, during the feature selection stage, interviewers asked participants to provide their reasoning for how they assigned importance levels to each preference feature. This not only provided insights into the participant's personal experiences and preferences but also allowed interviewers to clarify the meaning of any features that was not initially apparent. Upon completion, a semi-structure interview was conducted to understand their overall experience followed by an exit survey with optional demographic information.

Interviews with scheduling managers We conducted semi-structured interviews with scheduling managers via video-conferencing. The first set of questions focused on their scheduling practices such as walk-through of how they created the last period's schedule, how they currently learn and use worker preferences, and what fairness means to them in scheduling workers. We then solicited managers' perspectives on shift worker well-being models. To illustrate our concept, we showed screenshots of our web-tool as well as example schedule outputs optimized for worker well-being models given hard constraints such as the total number of employees and store hours.

Analysis All interviews were recorded and transcribed. The research team had a weekly meeting to discuss emerging themes and observations. Three researchers read all transcripts and one researcher coded them within Dedoose following the thematic analysis method (Patton 2014). The team worked collectively to discuss and iteratively refine and consolidate themes around participant experiences and perceptions of the scheduling preference elicitation web-tool. We also analyzed preference models that our partici-

	Preference Feature	Explanation
Schedule Preferences	Shift Type	Workers' preferred combination of day, shift start time, and shift duration.
	Total Hours	The total hours assigned in a week.
	Weekdays	Shifts assigned only on weekdays.
	Weekdays & Weekends	Shifts assigned on both weekdays and weekends.
	Same Number of Days	Shifts assigned on the same days over weeks.
	Same Days	Shifts assigned for the same number of days over weeks.
Manager Fairness Preferences	Seniority	Worker who has high seniority.
	Performer	High performing worker, i.e., is productive, completes tasks effectively, assists coworkers.
	Reliable	Worker who is very reliable. They show up on time to their shifts and they rarely cancel.
	Volunteer	Worker who volunteered last month for shifts considered undesirable by their coworkers.
	Fewer Hours	Received fewer hours than requested.
	Fewer Preferred Shifts	Received fewer preferred shifts.
	Limited Availability	Worker who received fewer hours due to external circumstances (healthcare, childcare, etc.).

Table 1: Features used for schedule preference and managerial fairness models. Participants choose as many features as they wish. The schedule preference and managerial fairness models are learned independently through discrete choice experiments.

pants built by examining task ratings and the feature weights of their final models. We examined how participants initial feature importance ratings compared with feature weights learned through pairwise comparisons to understand how the participants' preferences evolve through the process of building the models focusing on participants who chose the learned models as their final models. The explicit importance were treated as ordered categorical data and scaled to the same range as the learned preference feature weights.

Findings

In this section we describe the effectiveness of our participatory well-being elicitation method and its impact on workers. We also report managers' perspectives on the role of worker well-being models in organizational management.

Elicited Worker Well-Being Preference Models

We first describe the work preference model (task and scheduling preference) and managerial fairness model that participants built. We then illustrate how our elicitation method helped participants discover their preferences.

Task preferences Participants reported that they perform a wide spectrum of tasks in their workplaces ranging from interacting with customers (e.g., cashiers and servers) to providing operational support in the workplace (e.g., bookkeeping and inventory). On average, each participant worked on 3.28 ± 1.04 tasks, often more than one per shift. We analyzed how participants rated the tasks along the 9 well-being metrics, comparing their most and least preferred tasks. The most preferred tasks were less stressful and provided desirable social interaction and independent decision-making compared to the least preferred tasks. Perceived physical health risk did not differ between these tasks.

We also analyzed how the same task was ranked by different participants. Cashier was the most common task, performed by 18 participants working in fast-food and retail and was reported as most-preferred, least-preferred, and at neither of the two extremes by 6, 5, and 7 participants respectively. During interviews, participants reported distinct reasons for their preferences: P25 said that solving problems and assisting customers made the task enjoyable, while P21 found negative interactions with mean customers to be stressful. This diverse span in preferences for the same task demonstrates that workers have varying task preferences and supports the potential for personalized task assignment that maximizes each workers' preferences.

Schedule preference and managerial fairness model Participants selected 5.24 ± 0.81 features for their schedule preference models and 6.56 ± 0.85 features for their managerial fairness models. The average model cross-validation accuracy across all participants was $55.64 \pm 0.17\%$ for schedule preferences and $84.04 \pm 0.10\%$ for managerial fairness model. After reviewing their own explicit feature weights and learned weights, 15 out of 25 participants chose the learned schedule preference model, and 21 out of 25 participants preferred the learned managerial fairness model.[2]

In our study, the participants expressed idiosyncratic scheduling preferences-no feature was consistently used in the same manner by everyone. For example, rather than everyone attempting to maximize the total hours worked, people preferred either a limited set of part-time hours or maximized total hours. Additionally, participants held distinct preferences for the days that they worked: students often preferred weekends, those with childcare needs preferred weekdays, and some were indifferent ("I work basically any day and I don't really mind any of it" -P10). The distinctive preferences are also captured in the schedule preference model that participants built (Figure 3). *Shift-type* was weighted highly by participants on average. However, there was no clear consensus as multiple participants rated different features highly, resulting in almost all features being weighted as the top factor by at least one. On an aggregate level, we observe that the mean value of *shift-type* feature for part-time workers (0.78) is significantly higher than for full-time workers (0.51) indicating that preference importance may depend upon the circumstances of the participants.

[2]We note that participants paid more attention the preference weights rather than the accuracy judging how the final weights were in line with their preferences after pairwise comparisons.

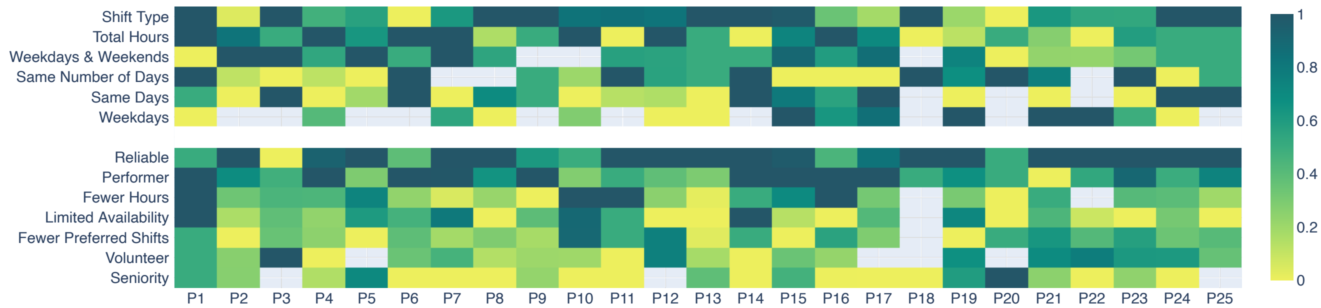

Figure 3: **Schedule preference and managerial fairness models**. (Left) Schedule preferences for each participant. (Right) Managerial fairness preferences for each participant. For both schedule preference model and managerial fairness model visualizations, we denote preferences not selected by participants with a default background color.

The managerial fairness features can be divided into three categories: merit (*performer*, *reliable*), need (*limited-availability*, *fewer-preferred-shifts*, *fewer-hours*), and auxiliary (*seniority*, *volunteering*). During pairwise comparisons, participants often made decisions based on merit first, followed by need, and employing auxiliary features when they felt the pairwise options were indistinguishable. This preference ordering for managerial fairness becomes apparent in the trends displayed in Figure 3. Across all demographics, *reliable* received the highest response from participants followed by *performer*. *Seniority* was rated lowest and had significantly higher value for full-time workers (0.30) compared to part-time workers (0.10).

Preference discovery through pairwise comparisons Pairwise comparisons were effective in drawing out users' preferences by allowing them to visualize realistic scenarios with combined features, as opposed to viewing each feature in isolation and out of context. Participants could recognize the compromises they were willing to make when all preferences could not be satisfied and were able to estimate the level of importance they held for each feature in practice. Initially, P19 cared most about *total-hours* and working the *same-number-of-days*. The model learned slightly different preferences though. Upon reviewing her explicit and learned weights side by side, she agreed that *total-hours* was actually the second least important factor while working *weekdays* only was the most important. P19 felt she understood her preferences better after the pairwise comparisons saying, "In the beginning, I said one thing was important. But going through the exercise, it became apparent that other things were actually important to me." P9 also remarked that pairwise comparisons allowed him to discover what preferences were more important to him: "After looking at the schedule that they [the AI] gave me, I'm thinking back like, oh, I think this is more important to me now."

In order to examine whether these patterns are also observed in the preference models themselves, we compared participants' explicit feature weights—feature weights that participants indicated before pairwise comparisons—and learned feature weights—feature weight learned through pairwise comparisons. In scheduling preference models, participants appeared to make a trade-off of schedule consistency and shift type for preferred choice of day: on average, weights of working *weekdays* only and working *weekdays-and-weekends* increased after pairwise comparisons, while weights of getting shifts for *same-number-of-days*, and *same-days*, and *shift-type* decreased. In managerial fairness models, weights of *reliability* and *performer* decreased. While *reliability* and *performer* still remained top features in participants' preference models, participants increased other needs-based features. P4 said "I used to think performance is really important but...I'm starting to see that number of hours [that workers have been assigned] is probably a little bit more important." Weights of *seniority* and *limited-availability*[3] also became less important overall. P8 said "[seniority is] one thing I'm willing to overlook...just because you have high seniority doesn't exactly mean you're a good performer or reliable."

We note that when participants had well-established preferences, feature weights did not change much between explicit and learned feature weights. For example, for many participants, scheduling feature *total-hours* hardly changed after pairwise comparisons: hours assigned are tied to workers' financial well-being so it stands to reason that many workers form firm preferences. We also saw stable preferences with some features that workers did not hold strong preferences for. The managerial fairness feature *volunteering* had close explicit and learned weights: only 3 of 25 workers assigned it high importance and correspondingly workers did not have strong preferences for it during pairwise comparisons ("It's nice to have but it's not that big of an importance to me." -P2). Interestingly, while *volunteering*, *fewer-hours*, and *fewer-preferred-shifts* had nearly unchanged learned and explicit feature weights *fewer-preferred-shifts* and *volunteering* increased in overall rankings after pairwise comparisons. This trend may be attributed to how, as mentioned above, the remaining features became less important in learned models and workers mixed need-based preferences and auxiliary preferences with merit-based preferences.

[3]We clarified the meaning of this feature to participants, but its unfamiliarity may have resulted in it not being prioritized much.

Empowerment through participation Worker empowerment refers to workplaces sharing power and information with employees to motivate and empower them (Fernandez and Moldogaziev 2013). We recognize empowerment as supporting workers' agency in decision-making in the workplace. We observed that our preference elicitation method can enable participation and empowerment of users. Speaking to managers, we found that tracking employee preferences was a manual process largely based on the manager's observations and/or explicit requests from the employees ("It's a book like a file, and I can write down what this person prefers over the next person... I find myself going based off of conversations that we've had." -M2). The feedback from shift workers echoed how these manual processes do not always use worker input or preferences and can result in frustrating errors. P7 said, "My schedule is usually just handed to me. It's just, 'these are the days you're going to work.' I don't have complete say and I'm not given options, like 'what shifts do you want to work?'" P12 commented that "everybody would probably be happier with [the AI] because the managers sometimes do forget some things you mention. The amount of times I've told her I wasn't available on this day and she still put me on there." In contrast, through the model creation process, workers gain knowledge about their preferences to be their own advocates and having a degree of control over a workplace process affecting them. P19 shared that she planned to use what she learned to communicate her preferences to her manager: "this sort of helped open my eyes to that. So I think that was a really good part of the exercise... That'll help me when I go to work later." P14 gave an anecdote that her boss continues to schedule one of her coworkers for the store's 5 AM shift despite that coworker's objections, but that a tool like this could take such preferences into account.

Potential Effects on Organizational Management

Assisting managers to create schedules that better account for worker well-being Our manager participants appreciated the tool's ability to give them insight into their worker's task and schedule preferences so that they could incorporate those individual characteristics into their scheduling process. All three managers currently create schedules by hand.[4] They also each record worker preferences manually, depending on observations and interactions to learn them: M1 updates a note in his phone, M2 keeps a handwritten file for employees, and M3 asks workers to follow-up in personal conversations or with emails. Upon seeing the tool, managers told us they would be able to use the information the tool provides as it utilizes attributes that they have either tried to collect but have not been able to so far or ones they had not thought of before but believe have merit.

The managers envisioned integrating preferences from the web-tool into scheduling. "I would definitely try to honor their feedback. I would take into high consideration how they feel and what they're saying because they're just as important as me, helping the store stay afloat and run smoothly." -M2. M3 also emphasized the importance of workers having a voice in the workplace and told us she felt the tool could empower junior workers for whom "this is their first job, so they may get nervous to speak to management." She added that if people could get their schedule preferences satisfied then "you're going into work a lot happier... it's going to be more of a positive vibe at work."

Fairness in workers' preference satisfaction We attempted to understand whether/how this tool might assist fairness in scheduling to improve overall worker well-being. A few participants mentioned a classic scheduling challenge that arises in the workplace, and managers shared their thoughts on solutions that they felt were most fair.

When asked what fairness means to them in scheduling, the managers had differing opinions. M1 felt fairness was achieved by making sure no one receives their preferred tasks more than twice a month so everyone works preferred and non-preferred tasks. M2 believed fairness meant honoring employee availability as much as possible so that workers were all happy with their scheduled hours. M3 explained that fairness meant distributing the number of hours amongst all workers as equally as possible.

Workers and managers alike expressed that every worker's desires cannot be satisfied to the fullest extent. P7 elaborated, "Let's say a lot of people want the morning shift, but they don't really need that many people. So how many people would they actually take into consideration to be at that shift? And how many people they would decide to not really care about their preferences?" Workers were not sure of the best way to resolve these cases, understandably, as they are not typically consulted on these situations by their supervisors. M3 brought up the same conundrum of fairly allocating shifts if all employees have similar preferences.

Managers shared their ideas for handling such situations. All of them agreed that a system for rotating preferences can balance whose preferences are met in each shift cycle. M1 suggested that the tool could also assist by tracking how everyone's task preferences were met each week to ensure fair distribution over time. Another theme that emerged was how managers attempted to meet everyone's preferences: all managers leaned towards schedules that emphasized equality over maximizing overall satisfaction of workers. They felt it would be fairer for everyone's preferences to be met to the same extent, as opposed to some people's preferences being met to a higher degree than others.

Interestingly, M2 did add that testing multiple schedules and asking employees for opinions would be the most fair. This feedback is useful as guidance for future tool iterations: these managerial strategies could be implemented and evaluated by workers in an extended fairness assessment.

Preserving worker communication Our overall objective with this study is to build an assistive tool for workers to create personalized preference models and for managers to incorporate these preferences in making scheduling decisions. Workers, particularly those who interact with their scheduling managers, suggested that a human should review and approve decisions made by an automated tool. One worker shared why she felt strongly that a human

[4] A few of our participants told us their employers used software tools, however, we were unable to recruit managers who used those.

should still be in charge of scheduling decisions: "Ideally, I would still prefer a human being. There's a level of empathy there. There's more complexity and nuance. I understand that it's more wasteful for the company's bottom line. But as a worker, of course, I would prefer to communicate with a good empathetic human being." -P23.

Discussion

In this paper, we proposed a participatory well-being elicitation method for algorithmic management with the goal of creating work conditions that are personalized for each worker's well-being. As a case study, we implemented our method in the context of shift work scheduling and evaluated it with 25 shift workers and three scheduling managers. Our findings provide initial evidence that our participatory method can elicit worker well-being models; moreover, we found that well-being models can assist scheduling managers to better account for worker well-being. Interviews with workers suggest that the shift schedule/work preference features capture their priorities, and our elicitation methods, particularly pairwise comparisons, helped workers discover their preferences in shift work.

The participatory process of building work and managerial fairness preference models also gave workers a sense of empowerment (Lee et al. 2019a). It is noteworthy that many of our participants discovered their schedule preferences, something that could have been overlooked if one assumes that workers have fully-formed preferences and do not require an elicitation process. One reason could be that in most of existing shift work, workers do not have much agency in scheduling other than providing their hard constraints such as availability; thus, workers might not have had a chance to form their preferences let alone realize the possibility. Our participants' comments that this process opened their eyes (P8), or that they will tell their manager about their preferences when they go back to work (P19), points to a possibility that this participatory process could help increase awareness of scheduling possibilities for shift workers.

Our study suggests that the key design decisions around worker well-being models—such as types of preference features, privacy and anonymity of preferences, fairness notions, and control—should be made in consideration of diverse organizational cultures and norms in which the system will be embedded. For example, one worker mentioned at the beginning that he was not planning to consider reliability or performance heavily because everyone at his workplace was a high performer. Some organizational structures allowed workers to interact with scheduling managers; others separated workers from scheduling managers, which prohibited informal exchanges for the managers to learn their workers' preferences. Our method could create new grounds where worker well-being preferences can be used to create and evaluate shift schedules; it can also risk compromising existing manager-worker interaction in some organizations. Careful consideration should be given in design so that both the strengths of computational worker well-being models and communication between managers and workers can be leveraged. For example, computational well-being models can play a role of boundary objects establishing common ground, while managers and workers can collectively update it and build consensus in response to changing situations (Alkhatib and Bernstein 2019).

We see the potential to apply worker well-being models to other algorithmically managed workplaces such as gig work. In gig work, there are no human managers; thus, developing relationships between workers and managers to learn preferences is not an option. Participatory well-being models will give workers a new opportunity to voice their well-being preferences in management, such as task assignment. For example, gig drivers' preferences about ride types and temporal assignment patterns could be used as information to break a tie when there are multiple drivers who are similarly distant from a ride requester.

Our work also points to future work on worker well-being modeling and AI fairness. In our study, managers had varying ideas on fairness over time. While emerging work began to propose algorithmic approaches to define fairness for repeated decisions, little work has investigated psychological and organizational perspectives on AI fairness in repeated allocation settings. This calls for expanding research on AI fairness to consider temporality. We also learned that scheduling preferences is more complex than we initially anticipated. While managerial fairness models achieved equivalent accuracy reported in prior work (Lee et al. 2019b), scheduling preference models' accuracy was lower. Future work should explore different modeling approaches to schedule preferences by using non-linear Bayesian methods and incorporating more elaborate fractional factorial design schemes that help model conditional features. Additionally, future work should investigate joint modeling of schedule preferences with task preferences.

We acknowledge the limitations of our study that readers should keep in mind. Our case study was conducted in the context of shift work with a small number of participants. Future studies should investigate our well-being elicitation method with a wider and representative sample of participants in different work contexts.

Conclusion

How can we center worker well-being as AI increasingly manages the workforce? As a first step toward this goal, we propose a participatory method for worker well-being models. We envision that such well-being models will enable management and working conditions to be optimized for worker well-being in addition to efficiency, and measure work's impact on worker well-being. Our case study in shift work scheduling suggests that our participatory method helps workers discover their preferences and build well-being models that they are satisfied with. Participation also provides workers a sense of empowerment. We hope our work will inspire further research that incorporates workers' voice and participation in AI integrated workplaces.

References

Alexander, C. S.; and Tippett, E. 2017. The hacking of employment law. *Mo. L. REv.* 82: 973.

Ali, S.; and Ronaldson, S. 2012. Ordinal preference elicitation methods in health economics and health services research: using discrete choice experiments and ranking methods. *British medical bulletin* 103(1): 21–44.

Alkhatib, A.; and Bernstein, M. 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.

Ananat, E. O.; and Gassman-Pines, A. 2021. Work Schedule Unpredictability: Daily Occurrence and Effects on Working Parents' Well-Being. *Journal of Marriage and Family* 83(1): 10–26.

Anya, O. 2015. Bridge the Gap! What Can Work Design in Crowdwork Learn from Work Design Theories? In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, 612–627.

Bjerknes, G.; Ehn, P.; Kyng, M.; and Nygaard, K. 1987. *Computers and democracy: A Scandinavian challenge*. Gower Pub Co.

Boivin, D.; and Boudreau, P. 2014. Impacts of shift work on sleep and circadian rhythms. *Pathologie Biologie* 62(5): 292–301.

Brüggen, E. C.; Hogreve, J.; Holmlund, M.; Kabadayi, S.; and Löfgren, M. 2017. Financial well-being: A conceptualization and research agenda. *Journal of Business Research* 79: 228–237.

Capio, C. M.; Sit, C. H. P.; and Abernethy, B. 2014. *Physical Well-Being*, 4805–4807. Dordrecht: Springer Netherlands. ISBN 978-94-007-0753-5. doi:10.1007/978-94-007-0753-5_2166. URL https://doi.org/10.1007/978-94-007-0753-5_2166.

Daverth, G.; Cassell, C.; and Hyde, P. 2016. The subjectivity of fairness: managerial discretion and work–life balance. *Gender, Work & Organization* 23(2): 89–107.

Ehn, P. 1988. *Work-oriented design of computer artifacts*. Ph.D. thesis, Arbetslivscentrum.

Eurofound. 2012. Fifth European working conditions survey.

Fernandez, S.; and Moldogaziev, T. 2013. Employee empowerment, employee attitudes, and performance: Testing a causal model. *Public Administration Review* 73(3): 490–506.

Finnigan, R. 2018. Varying weekly work hours and earnings instability in the Great Recession. *Social science research* 74: 96–107.

Fujishiro, K. 2005. *Fairness at work: Its impacts on employee well-being*. Ph.D. thesis, The Ohio State University.

Golden, L. 2015. Irregular work scheduling and its consequences. *Economic Policy Institute Briefing Paper* (394).

Gray, M. L.; and Suri, S. 2019. *Ghost work: how to stop Silicon Valley from building a new global underclass*. Eamon Dolan Books.

Guest, D. E. 2002. Perspectives on the study of work-life balance. *Social Science Information* 41(2): 255–279.

Hawkins, A. 2020. California labor commissioner sues Uber and Lyft for alleged wage theft. *The Verge* .

Henly, J. R.; and Lambert, S. J. 2014. Unpredictable work timing in retail jobs: Implications for employee work–life conflict. *Ilr Review* 67(3): 986–1016.

Hoppe, A.; Heaney, C. A.; and Fujishiro, K. 2010. Stressors, resources, and well-being among Latino and White warehouse workers in the United States. *American Journal of Industrial Medicine* 53(3): 252–263.

Huppert, F. A. 2009. Psychological well-being: Evidence regarding its causes and consequences. *Applied Psychology: Health and Well-Being* 1(2): 137–164.

Johnson, F. R.; Lancsar, E.; Marshall, D.; Kilambi, V.; Mühlbacher, A.; Regier, D. A.; Bresnahan, B. W.; Kanninen, B.; and Bridges, J. F. 2013. Constructing experimental designs for discrete-choice experiments: report of the ISPOR conjoint analysis experimental design good research practices task force. *Value in health* 16(1): 3–13.

Johnston, R. J.; Boyle, K. J.; Adamowicz, W.; Bennett, J.; Brouwer, R.; Cameron, T. A.; Hanemann, W. M.; Hanley, N.; Ryan, M.; Scarpa, R.; et al. 2017. Contemporary guidance for stated preference studies. *Journal of the Association of Environmental and Resource Economists* 4(2): 319–405.

Kalleberg, A. L. 2009. Precarious work, insecure workers: Employment relations in transition. *American sociological review* 74(1): 1–22.

Kantor, J. 2014. Working Anything but 9 to 5. *The New York Times* .

Karl, K. A.; and Sutton, C. L. 1998. Job values in today's workforce: A comparison of public and private sector employees. *Public Personnel Management* 27(4): 515–527.

Kellogg, K. C.; Valentine, M. A.; and Christin, A. 2020. Algorithms at work: The new contested terrain of control. *Academy of Management Annals* 14(1): 366–410.

Khovanskaya, V.; Dombrowski, L.; Rzeszotarski, J.; and Sengers, P. 2019. The Tools of Management: Adapting Historical Union Tactics to Platform-Mediated Labor. *Proceedings of the ACM on Human-Computer Interaction* 3(CSCW): 1–22.

Kovach, K. A.; et al. 1987. What motivates employees? Workers and supervisors give different answers. *Business Horizons* 30(5): 58–65.

Lambert, S. 2015. Managers' Strategies for Balancing Business Requirements with Employee Needs. *EPRN* .

Lambert, S. J. 2008. Passing the buck: Labor flexibility practices that transfer risk onto hourly workers. *Human relations* 61(9): 1203–1227.

Lambert, S. J.; Henly, J. R.; and Kim, J. 2019. Precarious work schedules as a source of economic insecurity and institutional distrust. *RSF: The Russell Sage Foundation Journal of the Social Sciences* 5(4): 218–257.

Lee, M. K. 2018. Understanding perception of algorithmic decisions: Fairness, trust, and emotion in response to algorithmic management. *Big Data & Society* 5(1): 2053951718756684.

Lee, M. K.; Jain, A.; Cha, H. J.; Ojha, S.; and Kusbit, D. 2019a. Procedural justice in algorithmic fairness: Leveraging transparency and outcome control for fair algorithmic mediation. *Proceedings of the ACM on Human-Computer Interaction* 3(CSCW): 1–26.

Lee, M. K.; Kusbit, D.; Kahng, A.; Kim, J. T.; Yuan, X.; Chan, A.; See, D.; Noothigattu, R.; Lee, S.; Psomas, A.; et al. 2019b. WeBuildAI: Participatory framework for algorithmic governance. *Proceedings of the ACM on Human-Computer Interaction* 3(CSCW): 1–35.

Lee, M. K.; Kusbit, D.; Metsky, E.; and Dabbish, L. 2015. Working with machines: The impact of algorithmic and data-driven management on human workers. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*, 1603–1612.

Lieberman, H. R.; Agarwal, S.; Caldwell, J. A.; and Fulgoni III, V. L. 2020. Demographics, sleep, and daily patterns of caffeine intake of shift workers in a nationally representative sample of the US adult population. *Sleep* 43(3): zsz240.

Louviere, J. J.; Hensher, D. A.; and Swait, J. D. 2000. *Stated choice methods: analysis and applications*. Cambridge university press.

Maestas, N.; Mullen, K. J.; Powell, D.; Von Wachter, T.; and Wenger, J. B. 2017. Working Conditions in the United States Results of the 2015 American Working Conditions Survey. *Maestas, Nicole, Kathleen J. Mullen, David Powell, Till von Wachter, and Jeffrey B. Wenger, Working Conditions in the United States: Results of the 2015 American Working Conditions Survey. Santa Monica, CA: RAND Corporation* .

Mateescu, A.; and Nguyen, A. 2019. Algorithmic Management in the Workplace. *Data Society Report* .

McClelland, M. 2012. I was a warehouse wage slave. *Mother Jones* (March/April).

McFadden, D. 1986. The choice theory approach to market research. *Marketing science* 5(4): 275–297.

Morgeson, F. P.; and Humphrey, S. E. 2006. The Work Design Questionnaire (WDQ): developing and validating a comprehensive measure for assessing job design and the nature of work. *Journal of applied psychology* 91(6): 1321.

Organization, I. L. 2011. Policies and regulations to combat precarious employment.

PAI. 2020. Framework for Promoting Workforce Well-being in the AI-Integrated Workplace. *Partnership in AI* .

Patton, M. Q. 2014. *Qualitative research & evaluation methods: Integrating theory and practice*. Sage publications.

Pressman, S. D.; Kraft, T.; and Bowlin, S. 2013. *Well-Being: Physical, Psychological, Social*, 2047–2052. New York, NY: Springer New York. ISBN 978-1-4419-1005-9. doi:10.1007/978-1-4419-1005-9_75. URL https://doi.org/10.1007/978-1-4419-1005-9_75.

Richard, H. J.; and Oldham, G. 1976. Motivation through the design of work: Test of a theory. *Organizational behavior and human performance* 16(2): 250–279.

Roose, K. 2019. A Machine May Not Take Your Job, but One Could Become Your Boss. *NY Times: The Shift (June 23, 2019), https://www. nytimes. com/2019/06/23/technology/artificial-intelligence-ai-workplace. html* .

Sarwar, B.; Karypis, G.; Konstan, J.; and Riedl, J. 2001. Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th international conference on World Wide Web*, 285–295.

Schneider, D.; and Harknett, K. 2019. Consequences of routine work-schedule instability for worker health and well-being. *American Sociological Review* 84(1): 82–114.

Sparr, J. L.; and Sonnentag, S. 2008. Fairness perceptions of supervisor feedback, LMX, and employee well-being at work. *European journal of work and organizational psychology* 17(2): 198–225.

Tortia, E. C. 2008. Worker well-being and perceived fairness: Survey-based findings from Italy. *The Journal of Socio-Economics* 37(5): 2080–2094.

Uhde, A.; Schlicker, N.; Wallach, D. P.; and Hassenzahl, M. 2020. Fairness and Decision-Making in Collaborative Shift Scheduling Systems. CHI '20, 1–13. New York, NY, USA: Association for Computing Machinery. ISBN 9781450367080. doi:10.1145/3313831.3376656. URL https://doi.org/10.1145/3313831.3376656.

Van den Bergh, J.; Beliën, J.; De Bruecker, P.; Demeulemeester, E.; and De Boeck, L. 2013. Personnel scheduling: A literature review. *European journal of operational research* 226(3): 367–385.

Wickwire, E. M.; Geiger-Brown, J.; Scharf, S. M.; and Drake, C. L. 2017. Shift work and shift work sleep disorder: clinical and organizational perspectives. *Chest* 151(5): 1156–1172.

Wood, A. J.; Graham, M.; Lehdonvirta, V.; and Hjorth, I. 2019. Good gig, bad gig: autonomy and algorithmic control in the global gig economy. *Work, Employment and Society* 33(1): 56–75.

Participant Demographics

	Age	Race	Gender	Industry	Education	Individual Income
P1	44	Latinx or Hispanic	Female	Government	Bachelor's degree or equivalent	Less than $ 20,000
P2	22	White	Female	Restaurant	Postgraduate or professional degree	Less than $ 20,000
P3	18	White	Female	Fast Food	High school incomplete (Currently in trade school)	Less than $ 20,000
P4	20	Asian	Female	Retail	Some postgraduate or professional (Now in grad school)	Less than $ 20,000
P5	32	American Indian or Alaskan	Female	Fast Food	Some college, no degree	$ 20,000 to $ 34,999
P6	18	White	Male	Fast Food	High school incomplete (Currently in undergrad)	Less than $ 20,000
P7	19	Latinx or Hispanic	Female	Fast Food	Some college, no degree (Currently in undergrad)	Less than $ 20,000
P8	19	Latinx or Hispanic	Male	Fast Food	Some college, no degree (Currently in undergrad)	Less than $ 20,000
P9	18	Latinx or Hispanic	Male	Fast Food	High school graduate (Currently in undergrad)	Less than $ 20,000
P10	20	Latinx or Hispanic	Male	Retail	Prefer not to say (Currently in school, not specified)	Prefer not to say
P11	18	Other (Middle Eastern)	Male	Fast Food	High school graduate (Currently in undergrad)	$ 50,000 to $ 74,999
P12	18	Latinx or Hispanic	Male	Fast Food	High school graduate (Currently in undergrad)	Less than $ 20,000
P13	33	White	Male	Manufacturing	Postgraduate or professional degree	Over $ 100,000
P14	19	White	Female	Retail	Some college, no degree (Currently in undergrad)	Less than $ 20,000
P15	22	Latinx or Hispanic	Male	Fast Food	High school graduate (Currently in undergrad)	$ 35,000 to $ 49,999
P16	45	Latinx or Hispanic	Female	Fast Food	High school graduate	Less than $ 20,000
P17	43	American Indian or Alaskan	Female	Fast Food	High school graduate	Less than $ 20,000
P18	42	Black or African American	Female	Healthcare/Social Work	2-Year Associate's Degree	$ 50,000 to $ 74,999
P19	40	White	Female	Retail	Postgraduate or professional degree	Over $ 100,000
P20	36	White	Male	Retail	High school graduate	$ 20,000 to $ 34,999
P21	35	Asian	Female	Retail	Some college, no degree	Less than $ 20,000
P22	31	Asian	Female	Healthcare/Social Work	Postgraduate or professional degree	$ 75,000 to $ 99,999
P23	27	Asian	Female	Healthcare/Social Work	Some college, no degree	Less than $ 20,000
P24	48	White	Female	Retail	Some college, no degree	$ 35,000 to $ 49,999
P25	41	White	Female	Retail	High school graduate	$ 20,000 to $ 34,999

Table 2: **Participant demographic information** Our study consists of 25 participants. We attempted to interview a wide range of participants based on their age, their racial history, their gender identity, the industry that they work in, their educational background, and their gross individual income.

	Age	Race	Gender	Industry	Education	Individual Income
M1	29	Asian	Male	Fast Food	Prefer not to say	$ 20,000 to $34,999
M2	41	Black or African American	Female	Fast Food	Associate Degree	$ 35,000 to $49,999
M3	22	White	Female	Fast Food	Some college, no degree (Currently in undergraduate college)	Prefer not to say

Table 3: **Manager demographic information** Our study consists of 3 managers who helped us understand current practices in their industry and get their inputs on the utility of our tool for managing shift scheduling their workplace.