

Participatory Algorithmic Management: Elicitation Methods for Worker Well-Being Models

Min Kyung Lee
minkyung.lee@austin.utexas.edu
The University of Texas at Austin

Ishan Nigam
ishann@cs.utexas.edu
The University of Texas at Austin

Angie Zhang
angie.zhang@austin.utexas.edu
The University of Texas at Austin

Joel Afriyie
joel.afriyie@austin.utexas.edu
The University of Texas at Austin

Zhizhen Qin
zhizhenqin@ucsd.edu
University of California San Diego

Sicun Gao
sicung@ucsd.edu
University of California San Diego

ABSTRACT

Artificial intelligence is increasingly being used to manage the workforce. Algorithmic management promises organizational efficiency, but often undermines worker well-being. How can we computationally model worker well-being so that algorithmic management can be optimized for and assessed in terms of worker well-being? Toward this goal, we propose a participatory approach for worker well-being models. We first define worker well-being models: Work preference models—preferences about work and working conditions, and managerial fairness models—beliefs about fair resource allocation among multiple workers. We then propose elicitation methods to enable workers to build their own well-being models leveraging pairwise comparisons and ranking. As a case study, we evaluate our methods in the context of algorithmic work scheduling with 25 shift workers and 3 managers. The findings show that workers expressed idiosyncratic work preference models and more uniform managerial fairness models, and the elicitation methods helped workers discover their preferences and gave them a sense of empowerment. Our work provides a method and initial evidence for enabling participatory algorithmic management for worker well-being.

CCS CONCEPTS

• **Human-centered computing** → HCI theory, concepts and models; • **Social and professional topics** → *Employment issues*.

KEYWORDS

Worker well-being, preference elicitation, algorithmic management, algorithmic fairness, participatory design

ACM Reference Format:

Min Kyung Lee, Ishan Nigam, Angie Zhang, Joel Afriyie, Zhizhen Qin, and Sicun Gao. 2021. Participatory Algorithmic Management: Elicitation Methods for Worker Well-Being Models. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (AIES '21)*, May 19–21, 2021, Virtual Event, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3461702.3462628>

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.

AIES '21, May 19–21, 2021, Virtual Event, USA

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

ACM ISBN 978-1-4503-8473-5/21/05...\$15.00

<https://doi.org/10.1145/3461702.3462628>

1 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, work scheduling, team formation, and performance evaluation—are automated or assisted by algorithms in various workplaces: crowdsourcing [18] and on-demand work platforms [41, 63, 69], as well as in offline workplaces such as warehouses [47], call centers [57], restaurants and retail stores [50], and offices [46]. AI integration is based on the promise that algorithmic management will boost workplace efficiency and economic value [31].

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 [47]; Uber and Lyft drivers feel automated evaluation is unfair and distrust the system's opaque payment calculations [20, 38, 41, 69]; shift workers suffer from unpredictable schedules that destabilize work-life balance and disrupt their ability to plan ahead [59]. 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 [53].

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 [58] along with recent work that considers social media user well-being [62]. 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 method to model worker well-being, adopting a participatory algorithmic governance framework [40]. We first define worker well-being models: Work preference models—preferences about work and working conditions that impact workers' physical, psychological, and financial well-being, and managerial fairness models—beliefs about fair resource allocation among multiple workers. We then propose elicitation methods to enable workers to build their own well-being models leveraging pairwise comparisons and ranking. As a case study, we applied our method to algorithmic work scheduling and shift workers, which comprise between 26–38 million American adults, or about 25% of the American working population [43]. We conducted

formative interviews with 9 shift workers and 2 managers to inform shift work well-being model design. We then had 25 shift workers use our method to build their well-being models and interviewed 3 shift scheduling managers. The findings show that workers expressed idiosyncratic work preference models and more uniform managerial fairness models, and the elicitation methods helped workers discover their preference and gave them a sense of empowerment. 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.

2 RELATED WORK

2.1 Algorithmic Management of Workers

AI is transforming multiple functions in workplaces, changing communication, collaboration, training, and workforce management [41, 46, 53]. 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 [31]. 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 [66]. However, recent advances in AI are expanding the scale of these tools, the number 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 [28]. Recent research has highlighted how shifting schedules negatively impact workers' health [59]. 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 [69]. Similarly, crowd workers on crowdsourcing platforms have little control in their work arrangement, task assignment, and evaluation [18].

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 [6, 12]. Khovanskaya et al. proposed adapting union tactics for algorithmic platform workers, such as requesting data transparency or contesting wage decisions [32]. Anya proposed ways to design worker-centered crowdsourcing task and environment design, leveraging the literature on job design [5]. Unde et al. studied what scheduling norms are perceived as fair by nurses [65]. Most recently, scholars called for the inclusion of worker well-being as an explicit focus when algorithmic work is designed [53]. To our knowledge, no research has explored methods of computationally defining worker well-being, and our work addresses this gap.

2.2 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 [9, 30, 44]. In the past, preference elicitation was utilized to measure the economic value of goods [49, 67]. 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 [2, 15, 34]. Preferences can be explicitly stated or revealed from observed behaviors [25, 44]. 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 [48]. 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 [2]. In our work, we adopt the discrete choice, specifically pairwise comparisons, and ranking methods to elicit worker's preferences for their well-being.

3 ELICITATION METHODS FOR WORKER WELL-BEING MODELS

Well-being is an umbrella term that describes "what it means to be functioning as a healthy person across multiple domains" [55]. 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 at scale. In this section, we present well-being related preference constructs in work preference and managerial fairness models.

3.1 Work Preference Model

Work preference models capture workers' preferences in work and working conditions that impact their physical, psychological, and financial well-being. Working conditions refer to the environment and context that workers perform their job, and encompasses aspects related to well-being such as schedule predictability, physical and social risks, job security, work activities and autonomy, and work-life balance [45]. To create our constructs for work preference models, we considered the ways these forms of well-being may materialize in preferences.

Physical well-being means "the ability to perform physical activities and carry out social roles that are not hindered by physical limitations and experiences of bodily pain, and biological health




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.

indicators” [8]. 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. Work-life balance also impacts physical well-being [19, 53], as work can hinder workers’ abilities to carry out social roles outside of work. This is particularly relevant to work scheduling and 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” [23]. In the context of work, the following factors contribute to psychological well-being [51, 56]: Autonomy in work—how much control workers have over what tasks to carry out and how; meaning derived from work—whether the work helps others or reinforces workers’ identities; and enjoyment in work—whether workers find the work pleasant and feel that 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” [7]. 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 [59]. We expect financial well-being matters most to the majority of workers, which could mean less individual variations in preferences. Still, workers may prioritize other types of well-being depending on their context, for example, whether they depend on the job for primary or supplementary income, and how well they can tolerate income volatility.

3.2 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 [16, 22, 61]. Managerial preference models capture what managerial rules workers deem fair when allocating work to multiple workers. Three dominant allocation principles exist [11]: 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.

3.3 Worker Well-Being Model Elicitation

To elicit worker well-being models for algorithmic management, we propose two methods drawing from literature on preference elicitation and adopting a participatory framework for algorithmic governance [40].

3.3.1 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 worker well-being and working condition papers [45, 51], 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 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.

3.3.2 Feature weight- & pairwise comparison-based elicitation. This method adopts the individual belief modeling part of WeBuildAI, a participatory algorithmic governance framework [40]. In this framework, users answer pairwise comparisons of alternatives to

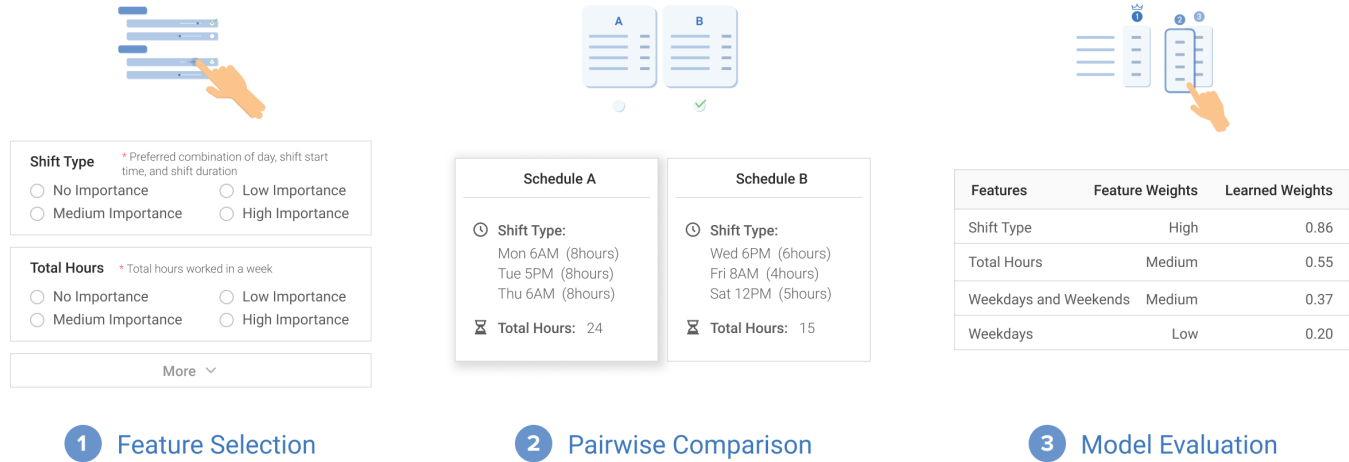


Figure 2: Feature weight- & pairwise comparison-based elicitation for schedule preference and managerial fairness models. 1) The worker chooses a set of relevant features. 2) The worker expresses preferences by choosing preferred options from a series of pairs of alternatives. 3) The worker evaluates the model learned from the pairwise comparison responses.

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).

4 CASE STUDY: ALGORITHMIC WORK SCHEDULING

Our case study applies the well-being elicitation method in the context of shift work scheduling. We chose shift work as recent research suggests worker well-being is compromised due to limited agency over schedules and on-demand dynamic scheduling [59]. 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.

4.1 Impact of Work Scheduling on Well-Being

Precarious work is a type of labor often characterized by uncertainty of working hours, lack of control by workers, and low wages

[26, 52]. A form of precarious work, shift work, employs 25% of workers in the United States [43]. 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 [14, 36, 37, 68]. Recently, researchers began to investigate the temporal dimension of shift work, namely "predictability and stability of work hours" [59]. Many companies turn to staffing strategies that reflect on-demand practices which minimize costs by matching the number of shift workers with forecasted demand [1, 37, 59]. 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 [59]. Shift workers also experience high levels of inconsistency in the hours and days they work each week [21, 36]. As most shift workers are paid hourly, any unpredictability of how much work they receive is directly correlated with financial stability [17]. Additionally, the inconsistency of which days they work prevents workers from freely planning non-work commitments [21]. These conditions culminate in circumstances where shift workers have to defer control over their own scheduling and consequently, their well-being [1, 4, 36].

4.2 Shift Worker Well-Being Model Design

To model worker well-being for shift work, we need to understand what factors of work scheduling influence worker well-being. We derived a set of features that constitute shift work preference and managerial fairness models based on the literature review and formative interviews with shift workers and scheduling managers.

	Preference Feature	Explanation
Schedule Preferences	Shift Type	The 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.
Managerial Fairness Preferences	Same Days	Shifts assigned for the same number of days over weeks.
	Reliability	Worker who is very reliable, i.e., shows up on time to shifts, rarely cancels.
	Performance	Worker who is high-performing, i.e., is productive, completes tasks effectively, assists coworkers.
	Fewer Hours	Worker who received fewer hours than requested.
	Limited Availability	Worker who received fewer hours due to external circumstances (healthcare, childcare, etc.).
	Fewer Preferred Shifts	Worker who received fewer preferred shifts.
	Volunteering	Worker who volunteered last month for shifts considered undesirable by their coworkers.
Seniority	Worker who has high seniority (years at the company).	

Table 1: Shift worker well-being model features. Schedule preference features capture shift work and working condition characteristics that influence workers’ physical, psychological, and financial well-being. Managerial fairness features capture factors that could be used to determine which workers should get assigned work/shift.

4.2.1 Formative interviews with shift workers & managers. We conducted one hour interviews with nine shift workers and two scheduling managers¹ employed in the fast food, retail, and healthcare industries to inform our design of the default preference features. The interview questions focused on understanding their workplace scheduling practices, how the schedules support or hinder worker goals, and factors considered when scheduling multiple workers for a limited number of shifts. We also complemented our interview findings with surveys on managers’ scheduling practices [35] and the impact of scheduling on worker well-being [59] (Section 4.1).

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 scheduling managers manually created schedules; others used scheduling software. Relationships between managers and workers varied. In a smaller workplace, the scheduling manager directly interacted with workers on the floors. In a larger workplace, workers interacted with middle managers but had no interaction with the scheduling manager. The degree to which workers’ preferences were considered in schedule creation largely depended on the presence of interaction. Some managers directly asked workers to share their preferences. In other workplaces, there was no formal system for worker preference information. In case scheduling managers knew a few workers personally, they considered the workers’ preferences and needs, giving them better schedules than the others (i.e. a worker who is a single mom and may have greater financial needs). In workplaces where managers did not interact with workers, there was no way for workers to convey their preferences as the methods workers used to indicate their availability did not solicit preference information.

4.2.2 Task preference. In our interviews, shift workers performed an array of tasks and they had well-established opinions on how different tasks took a different physical and/or mental toll. Accounting for workers’ task-level preferences would improve their well-being. We thus apply a ranking method asking workers to order tasks they could be assigned by preference. We also assess how each task

impacts their well-being by asking questions derived from work design literature [29, 33, 51, 64]: 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.

4.2.3 Schedule preference. In line with prior work [21], we learned from workers that consistent schedules promoted predictability in both work schedule and income, which allowed for a healthy work-life balance. As a metric of schedule predictability, 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 for their financial well-being.

4.2.4 Managerial fairness preference. In the scheduling manager report [35] and our interviews, the majority of managers gave more hours to workers who were good at sales and were reliable, which informed the *performance* and *reliability* features to indicate reliable or high performing workers. While the report found that managers did not typically consider workers’ financial needs, in our interviews, both managers and workers wanted to assign more shifts to those who have received less hours or less desirable shifts than others in the past, which led to the *fewer-hours* and *fewer-preferred-shifts* features. We also created the *limited-availability* feature as

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

	Age	Race	Gender	Industry	Education	Income
Workers						
W1	44	Latinx/Hispanic	Female	Government	Bachelor’s degree or equivalent	Less than \$20,000
W2	22	White	Female	Restaurant	Postgraduate or professional degree	Less than \$20,000
W3	18	White	Female	Fast Food	High school incomplete (Currently in trade school)	Less than \$20,000
W4	20	Asian	Female	Retail	Postgraduate or professional degree	Less than \$20,000
W5	32	American Indian	Female	Fast Food	Some college, no degree	\$20,000 to \$34,999
W6	18	White	Male	Fast Food	High school incomplete (Current undergraduate)	Less than \$20,000
W7	19	Latinx/Hispanic	Female	Fast Food	Some college, no degree (Current undergraduate)	Less than \$20,000
W8	19	Latinx/Hispanic	Male	Fast Food	Some college, no degree (Current undergraduate)	Less than \$20,000
W9	18	Latinx/Hispanic	Male	Fast Food	High school graduate (Current undergraduate)	Less than \$20,000
W10	20	Latinx/Hispanic	Male	Retail	Prefer not to say (Currently in school, not specified)	Prefer not to say
W11	18	Other (Middle Eastern)	Male	Fast Food	High school graduate (Current undergraduate)	\$50,000 to \$74,999
W12	18	Latinx/Hispanic	Male	Fast Food	High school graduate (Current undergraduate)	Less than \$20,000
W13	33	White	Male	Manufacturing	Postgraduate or professional degree	Over \$100,000
W14	19	White	Female	Retail	Some college, no degree (Current undergraduate)	Less than \$20,000
W15	22	Latinx/Hispanic	Male	Fast Food	High school graduate (Current undergraduate)	\$35,000 to \$49,999
W16	45	Latinx/ Hispanic	Female	Fast Food	High school graduate	Less than \$20,000
W17	43	American Indian	Female	Fast Food	High school graduate	Less than \$20,000
W18	42	Black/African American	Female	Social Work	2-Year Associate’s Degree	\$50,000 to \$74,999
W19	40	White	Female	Retail	Postgraduate or professional degree	Over \$100,000
W20	36	White	Male	Retail	High school graduate	\$20,000 to \$34,999
W21	35	Asian	Female	Retail	Some college, no degree	Less than \$20,000
W22	31	Asian	Female	Healthcare	Postgraduate or professional degree	\$75,000 to \$99,999
W23	27	Asian	Female	Healthcare	Some college, no degree	Less than \$20,000
W24	48	White	Female	Retail	Some college, no degree	\$35,000 to \$49,999
W25	41	White	Female	Retail	High school graduate	\$20,000 to \$34,999
Managers						
M1	29	Asian	Male	Fast Food	Prefer not to say	\$20,000 to \$34,999
M2	41	Black/African American	Female	Fast Food	Associate Degree	\$35,000 to \$49,999
M3	22	White	Female	Fast Food	Some college, no degree (Current undergraduate)	Prefer not to say

Table 2: Participant demographic information.

prior work suggests that the lower a worker’s socioeconomic status is, their availability becomes more limited, for example, due to no economic means to handle unexpected child care duties [60]. Finally, our interviews showed worker seniority played a role in determining schedules in some workplaces, and managers wanted to reward those who volunteered to work on a short notice, which resulted in the *seniority* and *volunteering* features.

4.3 Shift Worker Well-Being Model Elicitation

We built a web tool where workers can use the ranking and pairwise comparison-based elicitation methods to construct their shift work preference and managerial preference models.

4.3.1 Task evaluation and ranking. The web tool asked workers to enter information about tasks that they can be assigned. Workers then evaluated each task by answering questions about the task’s impact on physical and psychological well-being [45, 51, 59]. After evaluating each task, workers ranked the tasks where the first task was their most preferred task (Figure 1).

4.3.2 Schedule & managerial fairness model learning. We modeled scheduling and managerial fairness preferences (Figure 2) using the features in Table 1 and following a participatory algorithm design framework [40]. On the web tool, workers first indicated importance levels for each feature. The features with low to high

importance were then used to generate pairwise comparisons of alternatives based upon fractional factorial design principles [24].² Over a series of pairwise comparisons, workers chose a schedule they preferred (schedule preference) and a worker that should be assigned a shift (managerial fairness). Using the workers’ choice data, we used a logistic regression model based on random utility theory [40, 48] to train schedule preference and managerial fairness models.³

4.3.3 Model summary & evaluation. After pairwise comparison, workers saw the learned preference model summary: The model accuracy as well as a set of five alternatives ranked by the model. The web tool also showed feature weights learned by the model and the importance levels specified by the workers for each of the features. Using the information, the workers chose the model features—learned weights or explicitly specified importance—that better represented their preference as their final model.

²The number of comparisons shown to a worker depended on the number of selected features and their importance - more features and higher importance resulted in more comparisons. On average workers answered 28 pairwise comparisons for schedule preference models and 34 pairwise comparisons for managerial fairness models.

³Model learning is initialized from a normal prior distribution, and maximum likelihood estimation is performed using the BFGS method.

4.4 Method

4.4.1 Participants. We recruited 25 shift workers and 3 shift scheduling managers (Table 2). To recruit workers, we posted a recruiting message with a link to a screening survey on shift work related Reddit threads and ran Facebook ads. The survey asked about their current employment as a shift worker, demographics, scheduling practices in their workplace, and an attention check question. To recruit scheduling managers, we ran Facebook ads with a link to a screening survey. The survey asked for their employer and whether they currently created employee schedules.

Our worker participants included 9 White, 8 Latinx, 4 Asian, 2 American Indian, 1 Black, and 1 Middle Eastern. Their average age was 29.12 years (SD=10.7; Min-Max: 18-48), and 16 of them were female. They spanned a variety of industries, with the fast food industry being the majority (11 workers). 24 of our 25 participants answered optional education and income questions. 10 participants were enrolled in college; the highest completed education level of the other 14 participants varied from some high school to receiving a college degree. The average hourly wage of our worker participants was \$14.07 per hour. Our scheduling managers were White, Black, and Asian. Two managers were female. All of them were employed in the fast food industry.

4.4.2 Elicitation method evaluation with shift workers. We asked the 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 conferencing so the researcher could observe interactions with the tool. They interacted with the tool at their own pace and were encouraged to think aloud or ask for clarifications. The researcher observed the interactions and asked questions about their thought process throughout the session to understand their personal experiences and preferences. For example, during the feature selection stage, the researcher asked participants to provide their reasoning for how they assigned importance levels to each preference feature. Upon completion, a semi-structured interview was conducted to understand their overall experience followed by an exit survey with optional demographic information.

4.4.3 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 a 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.

4.4.4 Analysis. All interviews were recorded and transcribed. The research team had weekly meetings to discuss observations and emerging themes. Three researchers read all transcripts and one researcher coded them within Dedoose following the thematic analysis method [54]. The team worked collectively to discuss and iteratively refine and consolidate themes around participant experiences and perceptions of the elicitation methods. We also analyzed well-being models that our participants 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 from pairwise comparisons to understand how their preferences evolved 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.

5 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.

5.1 Elicited Worker Well-Being Models

5.1.1 Task preferences. Worker 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., book-keeping and inventory). On average, each participant worked on 3.28 ± 1.04 tasks, often more than one per shift. We analyzed how they rated the tasks along the nine well-being metrics (Section 4.2.2), comparing their most and least preferred tasks. The most preferred tasks were less stressful, provided desirable social interaction, and afforded 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 workers. Cashier was the most common task, performed by 18 participants working in fast-food and retail. Different workers reported cashier as most-preferred (N=6), least-preferred (N=5), and neutral tasks (N=7). During interviews, they shared distinct reasons for their preferences: W25 said that solving problems and assisting customers made the task enjoyable while W21 found negative interactions with mean customers to be stressful. This diverse span in preferences for the same task demonstrates heterogeneity in task preferences and supports the potential for personalized task assignment that maximizes each workers' preferences.

5.1.2 Schedule preference and managerial fairness models. Workers 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 workers 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 workers chose the learned schedule preference model, and 21 out of 25 workers chose the learned managerial fairness model.⁴

The workers 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, some preferred sacrificing the total hours in order to work during their preferred day/time windows or have consistent

⁴We note that workers paid more attention to the preference feature weights rather than the model accuracy and judged how the weights were in line with their preferences.




Figure 3: Schedule preference model (Top) and managerial fairness model (Bottom). The heatmap shows how important each feature is to each worker. Feature weights range from 0 to 1 (min to max importance); features that workers did not select are denoted with a grey background color. Workers expressed idiosyncratic scheduling preferences resulting in almost all features being weighted of the highest importance by at least one worker. A more clear delineation appeared for managerial fairness features: Merit features (*performance, reliability*) were strongly favored by workers while auxiliary features (*volunteering, seniority*) were rated of low or no importance by most workers.

across-week schedules. Additionally, workers 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” -W10). The distinctive preferences are also captured in the schedule preference model that participants built (Figure 3). *Shift-type* was weighted highly by many workers. However, there was no clear consensus as multiple workers rated different features highly, resulting in almost all features being weighted as the top factor by at least one worker. On an aggregate level, we observe that the mean value of *shift-type* feature for part-time workers ($M=0.78$) is significantly higher than for full-time workers ($M=0.51$) indicating that preference may depend on employment types.

The managerial fairness features can be divided into three categories: Merit (*performance, reliability*), need (*limited-availability, fewer-preferred-shifts, fewer-hours*), and auxiliary (*seniority, volunteering*). During pairwise comparisons, workers 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 is shown in the trends displayed in Figure 3. Across all workers, *reliability* received the highest response followed by *performance*. *Seniority* was rated lowest and had significantly higher value for full-time workers ($M=0.30$) compared to part-time workers ($M=0.10$).

5.2 Effects of Participatory Elicitation

5.2.1 Preference discovery through pairwise comparisons. Pairwise comparisons were effective in drawing out workers’ preferences by allowing them to visualize realistic scenarios with combined features, as opposed to viewing each feature in isolation and out of context. Workers could recognize the compromises they were willing to make when all preferences could not be satisfied and estimate the level of importance they held for each feature in practice. Initially, W19 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. W19 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.” W9 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, workers appeared to trade off schedule consistency and preferred 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 workers’ preference models, other needs-based features became more important. W4 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*⁵ 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 workers, the 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

⁵We clarified the meaning of this feature to participants, but its unfamiliarity may have resulted in it not being prioritized much.

stable preferences with some features that workers did not hold strong preferences for. The managerial fairness feature *volunteering* had similar 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.” -W2). Interestingly, while the feature weights of *volunteering*, *fewer-hours*, and *fewer-preferred-shifts* were nearly unchanged, *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 began to mix need-based and auxiliary preferences with merit-based preferences.

5.2.2 Empowerment through participation. Worker empowerment refers to workplaces sharing power and information with employees to motivate and empower them [13]. In this work, we define empowerment as supporting workers’ agency in decision-making in the workplace. We observed that our preference elicitation methods could 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 workers suggested that these manual processes did not always use worker input or preferences and could result in frustrating errors. W7 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?’” W12 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 gained knowledge about their preferences to be their own advocates and having a degree of control over a workplace process affecting them. W19 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.” W14 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.

5.3 Organizational Context Considerations

5.3.1 Assisting managers with scheduling. Our manager participants appreciated the tool’s ability to give them insight into their workers’ task and schedule preferences so that they could incorporate those individual characteristics into their scheduling process. All three managers currently created schedules by hand.⁶ They also each recorded worker preferences manually, depending on observations and interactions to learn them: M1 updated a note in his phone, M2 kept a hand-written file for employees, and M3 asked workers to follow-up in personal conversations or with emails.

⁶A few of our worker participants told us their employers used software tools, however, we were unable to recruit managers who used those.

Upon seeing the tool, the managers told us that the well-being models provided information about 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 believed to have merit. They envisioned integrating worker preferences 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.”

5.3.2 Satisfying multiple workers’ preferences fairly. Workers and managers alike expressed that every worker’s desires cannot be satisfied to the fullest extent, and managers had various ideas for how to handle fairly meeting each worker’s preferences. Workers contemplated how popularity of certain scheduling features like *shift-type* may outweigh the supply. W7 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 unsure 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 had 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 over time.

5.3.3 Preserving worker-manager communication. Our objective is to build an assistive tool for workers to create well-being preference models and for managers to incorporate these preferences in making scheduling decisions instead of fully automating scheduling. 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 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.” -W23.

6 DISCUSSION

As AI changes work, workers are exposed to risks and harms that arise from the experimentation and implementation of new algorithmic management techniques [50]. In line with a recent call to address the impact of AI on power distribution [27], we seek to create worker well-being models as a vehicle to incorporate workers' priorities and voices in algorithmic management with the goal of creating worker-centered and procedurally-fair algorithmic workplaces [39]—that is, working conditions personalized for each worker's well-being and management practices based on workers' inputs. Our case study in the context of work scheduling provides initial evidence that our participatory methods could elicit worker well-being models, that workers had idiosyncratic work preferences and relatively more uniform managerial fairness models, and that well-being models could assist scheduling managers to better account for worker well-being.

In our study, the participatory process of building work preference and managerial fairness models gave workers a sense of empowerment. It is noteworthy that many participants discovered their schedule preferences while going over the pairwise comparisons, 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 (W8), or that they will tell their manager about their preferences when they go back to work (W19), point to a possibility that this participatory process could help increase awareness of scheduling possibilities for shift workers. These findings highlight that elicitation methods should be designed to facilitate the preference discovery of those who are marginalized [42, 59] and should be used repeatedly for them to update their models.

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 [3].

We see the potential to apply worker well-being models to other algorithmically managed workplaces such as gig work [41, 63]. 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 has begun to propose algorithmic approaches to define fairness for repeated decisions [10], 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 [40], 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.

7 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 algorithmic 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.

ACKNOWLEDGMENTS

We thank our participants for their valuable insights. Katherine Hill helped with literature reviews; Anqing Chen and Shruthi Krish built the web tool; Chang Jun Cho and Janet Dai created the illustrations; Hemank Lamba, Kyle Chin, Chang Jun Cho, Brent Hong, Nathan Jen, and Rachel Lee contributed to the initial web tool prototype. Our research was supported by National Science Foundation Awards IIS-1939606 and CNS-1952085, and Good Systems⁷, a UT Austin Grand Challenge to develop responsible AI technologies.

⁷<https://goodsystems.utexas.edu>

REFERENCES

- [1] Charlotte S Alexander and Elizabeth Tippet. 2017. The hacking of employment law. *Mo. L. Rev.* 82 (2017), 973.
- [2] Shehzad Ali and Sarah Ronaldson. 2012. Ordinal preference elicitation methods in health economics and health services research: using discrete choice experiments and ranking methods. *British Medical Bulletin* 103, 1 (2012), 21–44.
- [3] Ali Alkhatib and Michael Morstein. 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.
- [4] Elizabeth O Ananat and Anna Gassman-Pines. 2021. Work Schedule Unpredictability: Daily Occurrence and Effects on Working Parents' Well-Being. *Journal of Marriage and Family* 83, 1 (2021), 10–26.
- [5] Obinna Anya. 2015. Bridge the Gap! What Can Work Design in Crowdfork Learn from Work Design Theories?. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. 612–627.
- [6] Gro Bjerknes, Pelle Ehn, Morten Kyng, and Kristen Nygaard. 1987. *Computers and democracy: A Scandinavian challenge*. Gower Pub Co.
- [7] Elisabeth C Brüggen, Jens Högrev, Maria Holmlund, Sertan Kabadayi, and Martin Löfgren. 2017. Financial well-being: A conceptualization and research agenda. *Journal of Business Research* 79 (2017), 228–237.
- [8] Catherine M. Capio, Cindy H. P. Sit, and Bruce Abernethy. 2014. *Physical Well-Being*. Springer Netherlands, Dordrecht, 4805–4807. https://doi.org/10.1007/978-94-007-0753-5_2166
- [9] Li Chen and Pearl Pu. 2004. *Survey of preference elicitation methods*. Technical Report. Swiss Federal Institute of Technology in Lausanne.
- [10] Alexander D'Amour, Hansa Srinivasan, James Atwood, Pallavi Baljekar, D Sculley, and Yoni Halpern. 2020. Fairness is not static: deeper understanding of long term fairness via simulation studies. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 525–534.
- [11] Gwen Daverth, Catherine Cassell, and Paula Hyde. 2016. The subjectivity of fairness: managerial discretion and work–life balance. *Gender, Work & Organization* 23, 2 (2016), 89–107.
- [12] Pelle Ehn. 1988. *Work-oriented design of computer artifacts*. Ph.D. Dissertation. Arbetslivscentrum.
- [13] Sergio Fernandez and Tima Moldogaziev. 2013. Employee empowerment, employee attitudes, and performance: Testing a causal model. *Public Administration Review* 73, 3 (2013), 490–506.
- [14] Ryan Finnigan. 2018. Varying weekly work hours and earnings instability in the Great Recession. *Social science research* 74 (2018), 96–107.
- [15] A Myrick Freeman III, Joseph A Herriges, and Catherine L Kling. 2014. *The measurement of environmental and resource values: theory and methods*. Routledge.
- [16] Kaori Fujishiro. 2005. *Fairness at work: Its impacts on employee well-being*. Ph.D. Dissertation. The Ohio State University.
- [17] Lonnie Golden. 2015. Irregular work scheduling and its consequences. *Economic Policy Institute Briefing Paper* 394 (2015).
- [18] Mary L Gray and Siddharth Suri. 2019. *Ghost work: how to stop Silicon Valley from building a new global underclass*. Eamon Dolan Books.
- [19] David E Guest. 2002. Perspectives on the study of work-life balance. *Social Science Information* 41, 2 (2002), 255–279.
- [20] Andrew Hawkins. 2020. California labor commissioner sues Uber and Lyft for alleged wage theft. *The Verge* (2020).
- [21] Julia R Henly and Susan J Lambert. 2014. Unpredictable work timing in retail jobs: Implications for employee work–life conflict. *Ilr Review* 67, 3 (2014), 986–1016.
- [22] Annekatrin Hoppe, Catherine A Heaney, and Kaori Fujishiro. 2010. Stressors, resources, and well-being among Latino and White warehouse workers in the United States. *American Journal of Industrial Medicine* 53, 3 (2010), 252–263.
- [23] Felicia A Huppert. 2009. Psychological well-being: Evidence regarding its causes and consequences. *Applied Psychology: Health and Well-Being* 1, 2 (2009), 137–164.
- [24] F Reed Johnson, Emily Lancsar, Deborah Marshall, Vikram Kilambi, Axel Mühlbacher, Dean A Regier, Brian W Bresnahan, Barbara Kanninen, and John FP Bridges. 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 (2013), 3–13.
- [25] Robert J Johnston, Kevin J Boyle, Wiktor Adamowicz, Jeff Bennett, Roy Brouwer, Trudy Ann Cameron, W Michael Hanemann, Nick Hanley, Mandy Ryan, Riccardo Scarpa, et al. 2017. Contemporary guidance for stated preference studies. *Journal of the Association of Environmental and Resource Economists* 4, 2 (2017), 319–405.
- [26] Arne L Kalleberg. 2009. Precarious work, insecure workers: Employment relations in transition. *American sociological review* 74, 1 (2009), 1–22.
- [27] Pratyusha Kalluri. 2020. Don't ask if artificial intelligence is good or fair, ask how it shifts power. *Nature* 583, 7815 (2020), 169–169.
- [28] Jodi Kantor. 2014. Working Anything but 9 to 5. *The New York Times* (2014).
- [29] Katherine A Karl and Cynthia L Sutton. 1998. Job values in today's workforce: A comparison of public and private sector employees. *Public Personnel Management* 27, 4 (1998), 515–527.
- [30] Ralph L Keeney, Howard Raiffa, and Richard F Meyer. 1993. *Decisions with multiple objectives: preferences and value trade-offs*. Cambridge University Press.
- [31] Katherine C Kellogg, Melissa A Valentine, and Angele Christin. 2020. Algorithms at work: The new contested terrain of control. *Academy of Management Annals* 14, 1 (2020), 366–410.
- [32] Vera Khovanskaya, Lynn Dombrowski, Jeffrey Rzeszutarski, and Phoebe Sengers. 2019. The Tools of Management: Adapting Historical Union Tactics to Platform-Mediated Labor. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–22.
- [33] Kenneth A Kovach et al. 1987. What motivates employees? Workers and supervisors give different answers. *Business Horizons* 30, 5 (1987), 58–65.
- [34] Alan Krupnick, Anna Alberini, Maureen Cropper, Nathalie Simon, Bernie O'Brien, Ron Goeree, and Martin Heintzelman. 2002. Age, health and the willingness to pay for mortality risk reductions: a contingent valuation survey of Ontario residents. *Journal of Risk and Uncertainty* 24, 2 (2002), 161–186.
- [35] Susan Lambert. 2015. Managers' Strategies for Balancing Business Requirements with Employee Needs. *EPRN* (2015).
- [36] Susan J Lambert. 2008. Passing the buck: Labor flexibility practices that transfer risk onto hourly workers. *Human relations* 61, 9 (2008), 1203–1227.
- [37] Susan J Lambert, Julia R Henly, and Jaeseung Kim. 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 (2019), 218–257.
- [38] Min Kyung Lee. 2018. Understanding perception of algorithmic decisions: Fairness, trust, and emotion in response to algorithmic management. *Big Data & Society* 5, 1 (2018), 2053951718756684.
- [39] Min Kyung Lee, Anuraag Jain, Hea Jin Cha, Shashank Ojha, and Daniel Kusbit. 2019. Procedural justice in algorithmic fairness: Leveraging transparency and outcome control for fair algorithmic mediation. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–26.
- [40] Min Kyung Lee, Daniel Kusbit, Anson Kahng, Ji Tae Kim, Xinran Yuan, Allissa Chan, Daniel See, Ritesh Noothigattu, Siheon Lee, Alexandros Psomas, et al. 2019. WeBuildAI: Participatory framework for algorithmic governance. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–35.
- [41] Min Kyung Lee, Daniel Kusbit, Evan Metsky, and Laura Dabbish. 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.
- [42] Min Kyung Lee and Kate Rich. 2021. Who Is Included in Human Perceptions of AI?: Trust and Perceived Fairness around Healthcare AI and Cultural Mistrust. In *Proceedings of the 39th annual ACM conference on human factors in computing systems (CHI)*.
- [43] Harris R Lieberman, Sanjiv Agarwal, John A Caldwell, and Victor L Fulgoni III. 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 (2020), zsz240.
- [44] Jordan J Louviere, David A Hensher, and Joffre D Swait. 2000. *Stated choice methods: analysis and applications*. Cambridge University Press.
- [45] Nicole Maestas, Kathleen J Mullen, David Powell, Till Von Wachter, and Jeffrey B Wenger. 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 (2017).
- [46] Alexandra Mateescu and Aiha Nguyen. 2019. Algorithmic Management in the Workplace. *Data Society Report* (2019).
- [47] Mac McClelland. 2012. I was a warehouse wage slave. *Mother Jones* March/April (2012).
- [48] Daniel McFadden. 1986. The choice theory approach to market research. *Marketing science* 5, 4 (1986), 275–297.
- [49] Robert Cameron Mitchell, Richard T Carson, and Richard T Carson. 1989. *Using surveys to value public goods: the contingent valuation method*. Routledge.
- [50] Pegah Moradi and Karen Levy. 2020. The Future of Work in the Age of AI: Displacement or Risk-Shifting? (2020).
- [51] Frederick P Morgeson and Stephen E Humphrey. 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 (2006), 1321.
- [52] International Labour Organization. 2011. Policies and regulations to combat precarious employment.
- [53] PAI. 2020. Framework for Promoting Workforce Well-being in the AI-Integrated Workplace. *Partnership in AI* (2020).
- [54] Michael Quinn Patton. 2014. *Qualitative research & evaluation methods: Integrating theory and practice*. Sage publications.
- [55] Sarah D. Pressman, Tara Kraft, and Stephanie Bowlin. 2013. *Well-Being: Physical, Psychological, Social*. Springer New York, New York, NY, 2047–2052. https://doi.org/10.1007/978-1-4419-1005-9_75
- [56] Hackman J Richard and Greg Oldham. 1976. Motivation through the design of work: Test of a theory. *Organizational behavior and human performance* 16, 2 (1976), 250–279.

- [57] Kevin Roose. 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> (2019).
- [58] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. 2001. Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th international conference on World Wide Web*. 285–295.
- [59] Daniel Schneider and Kristen Harknett. 2019. Consequences of routine work-schedule instability for worker health and well-being. *American Sociological Review* 84, 1 (2019), 82–114.
- [60] D Schneider and K Harknett. 2019. It's about time: How work schedule instability matters for workers, families, and racial inequality. *The Shift Project Research Brief, October* (2019).
- [61] Jennifer L Sparr and Sabine Sonnentag. 2008. Fairness perceptions of supervisor feedback, LMX, and employee well-being at work. *European journal of work and organizational psychology* 17, 2 (2008), 198–225.
- [62] Jonathan Stray. 2020. Aligning AI Optimization to Community Well-Being. *International Journal of Community Well-Being* 3, 4 (2020), 443–463.
- [63] Will Sutherland, Mohammad Hossein Jarrahi, Michael Dunn, and Sarah Beth Nelson. 2020. Work precarity and gig literacies in online freelancing. *Work, Employment and Society* 34, 3 (2020), 457–475.
- [64] Ermanno C Tortia. 2008. Worker well-being and perceived fairness: Survey-based findings from Italy. *The Journal of Socio-Economics* 37, 5 (2008), 2080–2094.
- [65] Alarith Uhde, Nadine Schlicker, Dieter P. Wallach, and Marc Hassenzahl. 2020. Fairness and Decision-Making in Collaborative Shift Scheduling Systems (*CHI '20*). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376656>
- [66] Jorne Van den Bergh, Jeroen Beliën, Philippe De Bruecker, Erik Demeulemeester, and Liesje De Boeck. 2013. Personnel scheduling: A literature review. *European journal of operational research* 226, 3 (2013), 367–385.
- [67] Lingappan Venkatachalam. 2004. The contingent valuation method: a review. *Environmental Impact Assessment Review* 24, 1 (2004), 89–124.
- [68] Emerson M Wickwire, Jeanne Geiger-Brown, Steven M Scharf, and Christopher L Drake. 2017. Shift work and shift work sleep disorder: clinical and organizational perspectives. *Chest* 151, 5 (2017), 1156–1172.
- [69] Alex J Wood, Mark Graham, Vili Lehdonvirta, and Isis Hjorth. 2019. Good gig, bad gig: autonomy and algorithmic control in the global gig economy. *Work, Employment and Society* 33, 1 (2019), 56–75.