Who Benefits from Salary Prediction Tools? A Simulation Study of Worker Power and Wage Inequality in AI-Mediated Salary Negotiations

SERAFINA KAMP, KATHLEEN CREEL, SINA FAZELPOUR

Abstract Pay transparency efforts, such as requiring firms to post salary ranges in job posts, aim to reduce information asymmetry in the labor market and correct historic pay inequalities between demographic groups, such as men and women. While some forms of pay transparency have been shown to improve wages for all workers, only some states require it and significant wage gaps persist between demographic groups. In response, online salary prediction tools powered by artificial intelligence have been developed to deliver pay transparency from new data sources, namely online crowd-sourced salary data. However, since these tools introduce uncertainty over the credibility and accuracy of the information, it is a priori unclear whether workers will benefit from these tools and whether they will reduce pay gaps. In this work, we model the extent to which salary prediction tools relying on crowd-sourced data reduce pay inequality and how their use influences whether workers or firms have more power in negotiations. To do so, we develop a labor market model where firms and workers are simultaneously learning to negotiate over time under different forms of pay transparency, including the presence of a salary prediction tool. Then, we use simulations to evaluate how firm and worker use of this tool affect worker wages as well as wage gaps between workers. We find that salary prediction tools can introduce more diverse information in the market by encouraging workers to negotiate more often and this improves worker wages overall, especially when firms directly post salary ranges. However, depending on the market conditions, the salary prediction tool can lead to wider wage gaps in expectation between worker types (corresponding to demographic differences) compared to firms posting salary ranges in a market without such a tool.

Working paper, please do not distribute.

1 Introduction

When some job seekers have better information about wages than others, opportunities for pay inequality or even pay discrimination abound [30, 33]. Better informed workers negotiate better, asking for higher wages when appropriate and avoiding making salary asks so high that the offer is withdrawn [22]. Pay transparency policies, such as requirements to include salary ranges in job advertisements, aim to address this: They equalize *information access* as a way to decrease wage inequality, particularly between demographic groups such as men and women [25, 27, 38].

However, since not all regions have pay transparency policies, some have posited that job seekers might instead rely on online *salary prediction tools* to gain information about what to ask for in salary negotiations [31, 32, 48]. On the other side of the market, firms typically perform *salary benchmarking* to gain information for salary negotiations through benchmarks curated by third-party sources [6, 20, 50]. Firms also rely on predictive tools to provide benchmark or pay recommendations [43, 51, 52]. In the best case, these predictive tools serve as valuable information sources, allowing firms and workers to pool information in order to make more accurate salary estimates, thereby leading to a market with widespread competitive wages. In the worst case, these tools can lead workers to develop incorrect wage expectations [31, 49] or can aid firms in artificially suppressing wages via tacit (or explicit) collusion [34, 46, 47].

Such predictive tools often rely on *crowd-sourced data* either from workers sharing their employment information directly (e.g., on Glassdoor [1] or Salary.com) or from firms sharing their compensation data with third-party benchmarks (e.g., Quarterly Census for Employment and Wages or through companies like Payscale). These data sources provide opportunities to collect large-scale granular salary information, but the accuracy of the information is uncertain and depends on factors such as whether information is honestly reported, the population that shares data is skewed towards a particular demographic, or firms and workers continue to share their information over time [9, 35, 40]. So, studying how firms and workers share and use data from such a tool over time is key to evaluating whether these tools are effective at providing pay transparency.

In this work, we develop a model to evaluate the effectiveness of predictive tools for both enabling competitive wages for workers and reducing wage gaps between demographic groups. In our model, we evaluate the particular property of the use of *crowd-sourced data* of the salary prediction tools: We abstract away the details of the predictive algorithms [2], and assume firms and workers can access wage percentiles from data pools comprised of wage data from workers and firms who chose to share the data. Notably, firms and workers have access to *different* data pools and we compare wage outcomes under several *information settings*, including markets with and without salary prediction tools and another form of non-algorithmic pay transparency. Finally, we model firms and workers from different demographic groups simultaneously learning data sharing and negotiation strategies over time in order to capture the dynamics of how asymmetric access to information about the underlying wage distribution may be resolved over time.

We assume firms always have access to at least two salary benchmarks, including one that simulates a third-party AI-recommended benchmarks. Then, we evaluate the effect of adding a salary prediction tool for workers (e.g., Glassdoor) by comparing four possible

[1] For example, see https://help.glassdoor.com/s/article/What-Salary-Information-is-on-Glassdoor?language=en_US.
[2] Salary prediction tools used in practice are proprietary and cannot be evaluated directly, e.g., Glassdoor's algorithm: https://media.glassdoor.com/knowyourworth/know-your-worth-employer-guide.pdf

Author's Contact Information: Serafina Kamp, Kathleen Creel, Sina Fazelpour.

information settings: workers see 1) their current wage only (hereafter `Neither`), 2) firm-posted salary ranges (hereafter `Firm-posted`), 3) firm-specific salary range estimates from a salary prediction tool (hereafter `Predicted`), and 4) firm-posted salary ranges which could be chosen from estimates from the salary prediction tool (hereafter `Both`). `Neither` represents a naive baseline of workers making offers based on no information besides their current wage and `Firm-posted` gives the baseline of the popular pay transparency policy of firms posting salary ranges with a job opening. [3] The last two settings consider the intervention of a salary prediction tool where workers either completely rely on the tool's estimates for market information (`Predicted`) or they can see if a firm chooses to follow the estimate or not in their posted range (`Both`). Thus, our analysis provides insight into whether salary prediction tools convey useful enough information on their own to workers of different demographic groups and whether they provide benefits over and above firms posting salary ranges based on their private information alone.

For each information setting, we empirically investigate potential long-term wage outcomes with simulation experiments. In our experiments, we consider firms learning to choose salary benchmarks and acceptance thresholds and, simultaneously, workers learning whether to share data, which offers to make, and to which firms. Agents learning strategies simultaneously creates a non-stationary market environment that is notably difficult to analyze theoretically, and our simulation approach is in line with previous work studying wage and price setting strategies [23, 36]. We aim to answer the following research questions with this work.

- (RQ 1) What is the effect of a salary prediction tool on *average worker wages* in markets with and without firm-posted salary ranges?
- (RQ 2) What is the effect of a salary prediction tool on *wage gaps between workers who differ along demographic lines* in markets with and without firm-posted salary ranges?
- (RQ 3) What characteristics in worker offer behavior and firm benchmarking behavior could contribute to differences in wage outcomes between information settings?

The rest of this paper is organized as follows: Section 2 reviews related work, Section 3 introduces our labor market and learning models, Section 4 describes the experimental design for our simulations, Section 5 provides the analysis of the simulation results, and, finally, Section 6 discusses the implications of our results and future work.

2 Related Work

Imperfect Information in Labor Markets. Standard competitive labor market models assume firms and workers have perfect information (i.e., equal and complete information access) about market conditions, [4] though in real labor markets neither firms nor workers have perfect information [2]. Firms conduct costly salary benchmarking to learn about the wage distribution for their job postings and workers engage in costly search to learn about the market directly or approximate it with noisy data from online sources. Although both parties have imperfect information, firms generally have the power to reduce their uncertainty more than workers can [2, 45]. Further, some workers are able to reduce uncertainty more than other workers of comparable skill sets (e.g., those who face lower job search frictions). Both factors contribute to the observed wage dispersion between workers and the wage disparities between, e.g., men and women [12, 33]. However, recent research on (non-algorithmic) pay transparency has shown that different forms of transparency are not equal in terms of the power they grant firms and workers [1, 12, 20–22, 29, 33]. Predicting the effect of particular pay transparency policies on worker pay, therefore, requires a better understanding of the effects of differential firm and worker access to information about the market, especially given that firms can leverage information to increase their market power and suppress wages [3, 42]. Our work addresses the effects of firm and worker access to *salary prediction tools.*

Evaluating Pay Transparency Policies. There is a growing body of economic research on the impact of pay transparency policies that highlights both positive and negative outcomes of different policies. For example, policies related to firms posting salary ranges, using the same third-party benchmark, or other public disclosures of salary information generally positively impacts workers [18, 27]: Studies have found increased wages along with new posted ranges, reduced gender wage gaps, wage compression and increased wages for low-wage workers after firms had access to the same benchmark, and wage compression of salaries publicly perceived to be excessive [1, 4, 20, 41]. However, online posted ranges have been found to be less informative than traditional benchmarks [5], and increased information availability could, theoretically, facilitate collusion among firms [27, 42]. On the other hand, policies aimed at allowing pay disclosure among colleagues within a single firm have been shown to be less effective since they not only can lead to decreased wages, negotiation attempts, and employee morale [8, 11, 21], but they can also be inhibited salary disclosure taboos [10, 19]. We add to this literature by analyzing the effectiveness of information from a *crowd-sourced salary prediction tool over time* and compare it to the current best practice of firms posting salary ranges with job descriptions.

[3] Several states have enacted pay transparency laws that require salary ranges in job postings, see https://www.paycor.com/resource-center/articles/pay-transparency-laws-by-state/ for an overview of pay transparency laws by state.
[4] See [39] for an introduction to standard models and terms used in economic analysis of competitive labor markets.

Labor and Product Market Simulations. There is recent work in simulating labor and product markets with reinforcement learning (RL) and large-language model (LLM) agents [14–16, 23, 28, 36]. In the labor market, there is work on analyzing wage choices by firms in online labor platforms [23] as well as search and matching models to analyze unemployment and wage outcomes [15, 16]. In the product market, there is work in analyzing how firms learn to set prices over time, and supra-competitive pricing outcomes have been found in simulations with both RL and LLM agents [14, 28, 36]. Our work differs from these studies in that we focus on the *information sources that both firms and workers have access to when learning negotiation strategies in a labor market*, rather than assuming workers always choose the firm with the highest posted wage or fixing the bargaining power of firms or workers.

Evaluating Effectiveness of AI Tools for Salary Negotiation. Recent studies have explored the effectiveness of AI tools for salary negotiation [9, 17, 24, 31, 35]. In particular, several studies have shown that LLM agents are not more effective than traditional handbooks at giving negotiation advice, and they have even been shown to give biased advice for opening offers if the prompt reveals gender or age [17, 24, 31]. Additionally, recent work has analyzed the effectiveness of online salary platforms from two angles: Brown et al. [9] investigate the reasons employees share information on such a platform (in a single-shot decision) and find it is based on 1) how useful information appears to be on the website, 2) how uncertain they are about the unfairness of their current pay, and 3) how restrictive their current employer is about pay disclosure whereas Karabarbounis and Pinto [35] analyze wage information from Glassdoor and find it to be largely in line with national wage statistics, indicating that the platform could be a credible source of wage information. In our work, we model the effectiveness of a salary prediction tool, much like Glassdoor's, over time and assume firms and workers are motivated to share information insofar as the tool is useful to them in wage negotiation.

3 Model

3.1 Labor Market Model

Agents. We consider a labor market with a finite number of firms, N_f, and a finite number of workers, N_w, where firms and workers are negotiating over wages. For the firms, we assume each firm will hire at most N_w workers, each worker up to this point brings in the same amount of additional revenue to the firm (called their *marginal revenue product (of labor)*, or MRPL [39]), and the firm will not pay a wage above the MRPL value. For the workers, we assume no job differentiation beyond wages, equal skills, and workers only switch employers if their wage does not decrease. Workers can be employed or unemployed, and we will refer to their employer by an index $j \in [N_f]$ or $\emptyset$ to indicate unemployed. The set of employees at firm j at time t is denoted $E_j^{(t)}$.

Worker Types. In our experiments, workers can be one of two types: type H with high exploration ability and high initial wages and type L with low exploration ability and low initial wages. Here, "exploration ability" corresponds to features that are correlated with demographic differences that have been shown to contribute to observed pay inequalities, e.g., gender differences in search costs and size of geographic search area for new jobs [7, 13, 26, 44] which exist alongside enduring pay inequalities [7, 37]. We operationalize exploration ability as a difference in *learning process*, described in Section 3.2. Keeping in line with our analogy to wage inequality between demographic groups, we will assume firms can identify and condition their negotiation strategy on worker type.

Market Parameters. First, the set of possible worker wages is finite, denoted $\mathcal{W}$. We assume $\mathcal{W}$ uniformly divides the interval $[0, 1]$, and we say the wages are *k-discretized* when there are k even intervals, i.e., $\mathcal{W} = \left\{0, \frac{1}{k}, \ldots, \frac{k-1}{k}, 1\right\}$. Under this representation, we take each firm's MRPL to be equal to 1 and a worker's leisure wage (i.e., utility value without employment) to be equal to 0. When a worker receives a wage equal to the MRPL value, we say they are receiving a *competitive wage* [39]. We refer to the wage of worker i at time t as $w_i^{(t)} \in \mathcal{W}$. Firms and workers will have access to market wage distribution information in the form of *wage percentiles* from different data sources, and we use the notation $\eta_c(W)$ to refer to the c-th percentile wage of a set of wages W with $c \in [0, 100]$.

Next, there are two parameters that we vary over in each experiment that we will refer to as the *market conditions*: Market riskiness and worker type distribution. First, attempting to negotiate for higher wages involves risk for workers, e.g., making offers that are too high may result in firms rescinding offers or otherwise ending negotiations [22], and we will call this condition *market riskiness*. Similar to Cullen et al. [22], we model market riskiness as the probability $r \in [0, 1]$ of a worker's wage being reset to leisure wage of 0 and becoming unemployed (i.e., their employer is $\emptyset$) if their offer is rejected. Finally, when there is more than one type of worker, the proportion of type H workers is given by $\tau_H \in [0, 1]$ and τ_L is $1 - \tau_H$.

Salary Prediction Tool. The salary prediction tool maintains a data pool of wages from workers who chose to share with the tool at time t, along with the employer they are from, and the set of wages shared from employees at firm j is denoted $D_j^{(t)}$. The set of wages shared across the market at time t is denoted $D^{(t)} = \cup_{j \in [N_f]} D_j^{(t)}$. At time t, the tool "predicts" a salary range for each firm j by reporting wage percentiles from $D_j^{(t-1)}$, defined as $\hat{\ell}_j^{(t)} = \eta_{10}\left(D_j^{(t-1)}\right), \quad \hat{u}_j^{(t)} = \eta_{90}\left(D_j^{(t-1)}\right)$. Additionally, if a worker shares their wage data at time t, then they get access to a median wage prediction across the market from the pool of wages shared in the previous round[5], denoted $\hat{m}^{(t)} = \eta_{50}(D^{(t-1)})$.

Firm Salary Benchmarks. In our market, firms use salary benchmarks to set their negotiation strategy. The set of benchmarks firm j chooses from at time step t is denoted $\mathcal{B}_j^{(t)}$. In all information settings, firms have access to the following two data sources: An *independent* salary benchmark, denoted $I_j^{(t)}$ for firm j, and a *shared* salary benchmark, denoted $S^{(t)}$. The independent benchmark reports wage percentiles from a firm's current employees: $I_j^{(t)} = \left(\ell_j^{(t)}, m_j^{(t)}, u_j^{(t)}\right)$, where $\ell_j^{(t)} = \eta_{10}\left(\left\{w_i^{(t)} : i \in E_j^{(t)}\right\}\right)$, $m_j^{(t)} = \eta_{50}\left(\left\{w_i^{(t)} : i \in E_j^{(t)}\right\}\right)$,$u_j^{(t)} = \eta_{90}\left(\left\{w_i^{(t)} : i \in E_j^{(t)}\right\}\right)$. Let $E_s^{(t)} = \left\{w_i^{(t)} : i \in \cup_{j \text{ sharing}} E_j^{(t)}\right\}$ be the set of wages of workers whose employer's shared their data at time t, then the shared benchmark reports wage percentiles from the set $S^{(t)} = \left(\ell_s^{(t)}, m_s^{(t)}, u_s^{(t)}\right)$, where $\ell_s^{(t)} = \eta_{10}\left(E_s^{(t)}\right)$,$m_s^{(t)} = \eta_{50}\left(E_s^{(t)}\right)$, $u_s^{(t)} = \eta_{90}\left(E_s^{(t)}\right)$. Here, $S^{(t)}$ simulates a "third-party" salary benchmark that are available to firms if they purchase a subscription and share their HR data to in order to gain access to the wage data. [6] So, if a firm chooses $B = S^{(t)}$, then they agree to share their employee's wage data at time step $t+1$. Finally, in `Predicted` and `Both`, firms also have access to the salary prediction tool and, for simplicity, we assume they consider the salary range the tool predicts *for themselves only*. This is the *predicted* benchmark, defined as $P_j^{(t)} = \left(\hat{\ell}_j^{(t)}, m_j^{(t)}, \hat{u}_j^{(t)}\right)$, where $\hat{\ell}_j^{(t)}$ and $\hat{u}_j^{(t)}$ are the firm-specific estimates described above and, since the salary prediction tool does not produce firm-specific median estimates, the firm views $m_j^{(t)}$ as the median.

Market Phases. Each time step t in the market consists of two phases for both firms and workers: information gathering and wage negotiation. The information sources used in information gathering depend on the *information setting*. An overview of a market time step for one worker i and one firm j under each information setting is given in Figure 1, and we describe the details below.

First, each firm j gathers information by choosing a benchmark $B_j^{(t)} \in \mathcal{B}_j^{(t)}$, and we refer to the cth percentile from the chosen benchmark as $b_{j,c}^{(t)} \in B_j^{(t)}$. Then, they choose an acceptance threshold, denoted $a_j^{(t)} \in B_j^{(t)}$. If there is more than one type of worker, the firm chooses an acceptance threshold for each type, i.e., $a_j^{(t)} = (a_{j,H}^{(t)}, a_{j,L}^{(t)})$. Next, worker i gathers information by considering their current wage $w_i^{(t)}$ and the posted ranges, if any. First, if there is a salary prediction tool, workers choose to share $w_i^{(t)}$ or not in order to gain access to $\hat{m}^{(t)}$, and this action is denoted $s_i^{(t)} \in \{\texttt{share, no share}\}$. Then, worker i chooses a firm to attempt negotiations with, denoted $f_i^{(t)} \in [N_f]$, based on the posted salary ranges, if any. Finally, worker i chooses an offer $o_i^{(t)}$ from a set of offers, denoted $\{\underline{o}_i^{(t)}, \ldots, \overline{o}_i^{(t)}\} \subseteq \mathcal{W}$, where $\underline{o}_i^{(t)}$ and $\overline{o}_i^{(t)}$ are conditional on $w_i^{(t)}, s_i^{(t)}, \hat{m}^{(t)}$, and the posted range of $f_i^{(t)}$, if any, or they can choose `no offer`.

Wage negotiation is a simultaneous, one round bargaining game, and we choose the worker to always make an offer to the firm since we are primarily interested in how workers learn to react to interventions aimed at helping them negotiate for better wages (i.e., do they negotiate more often and have more accepted offers in one information setting over another). During wage negotiation, worker i offers $o_i^{(t)}$ to firm $f_i^{(t)}$ who accepts if and only if $o_i^{(t)} \leq a_{f_i^{(t)}}^{(t)}$. If accepted, worker i updates their wage to be $w_i^{(t+1)} = o_i^{(t)}$ and their employer to be $f_i^{(t)}$. However, recall making offers is *risky*, and, if their offer is rejected, then, with probability r, $w_i^{(t+1)} = 0$ and their employer becomes $\emptyset$, and, otherwise, $w_i^{(t+1)} = w_i^{(t)}$ and the worker remains at their current employer.

3.2 Learning Model

To model firms and workers simultaneously learning a strategy to in our market, we use *multi-agent reinforcement learning*. Specifically, we model each agent as part of a Markov decision process (MDP) with a non-stationary transition function and the assumption that agents decide which action to take based on the market information they directly observe each time step only. The states and actions of the agents in each information setting follow the information source and agent choice description, respectively, in Figure **??**, and we provide more detail on the states and actions as well as the reward functions in Appendix A.2.

The agents learn a policy to use in the market via standard Q-learning [53] with ϵ-greedy action exploration. In an MDP with a non-stationary transition function, theoretical convergence guarantees no longer hold for Q-learning algorithms, but we follow previous literature that empirically analyze simulations of agents using Q-learning in dynamic market environments to investigate possible wage or price outcomes [23, 36].

Q-learning Algorithm. Agents learn in our market using a simple Q-learning algorithm [53] with ϵ-greedy action exploration where agents choose an action uniformly at random with probability ϵ and a Q-value maximizing action otherwise. We use an ϵ decay schedule where $\epsilon^{(t)} = e^{-\beta \cdot t}$ for a given decay rate β. Type H workers (and firms) use the decay rate β_H such that $\epsilon^{(\frac{T}{2})} = 10^{-3}$ and

[5]This models examples such as Glassdoor (https://www.glassdoor.com/Salaries/index.htm) and Salary.com (https://www.salary.com/research/salary) where a worker can learn more market information by sharing personal salary information and/or buying it.
[6]E.g., Payscale's data services terms (https://www.payscale.com/about/documentation) and Ravio's contract terms (https://ravio.com/pricing).

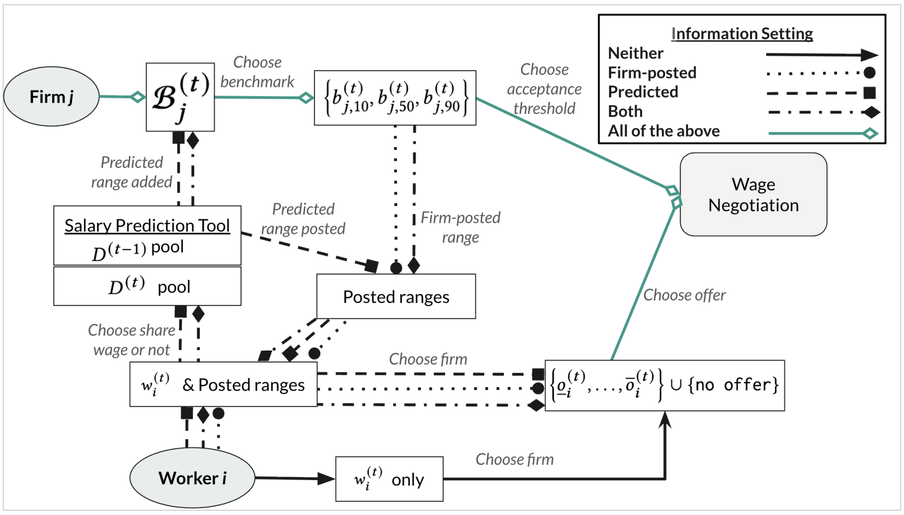

Fig. 1. Information flow and choices made at each time step t flow for worker i and firm j before negotiation under each information setting. The firms take the first actions by choosing their benchmark for time step t and posting ranges, if applicable, and choosing their private acceptance threshold. Then, the workers observe their current wage and the posted ranges, if any, before taking their actions of sharing their wage, if applicable, choosing a firm to negotiate with, and, finally, the offer they choose to make to that firm.

$\epsilon^{(T)} = 10^{-6}$ for total run time T, consistent with a similar (product) market simulation [36]. Type L workers use the (faster) decay rate β_L such that $\epsilon^{(\frac{T}{2})} = 10^{-5}$ and $\epsilon^{(T)} = 10^{-10}$.

The Q-values for each agent are initialized to 0 at the start, and the Q-learning update for agent $x \in \{w, f\}$ in state S after choosing action a, receiving reward $R_x(S, a)$, and transitioning to state S' (with corresponding action set A') is as follows.

$$Q_x^{(t+1)}(S,a) = (1-\alpha)\cdot Q_x^{(t)}(S,a) + \alpha \cdot \left(R_x(S,a) + \delta \cdot \max_{a'\in A'} Q_x^{(t)}(S',a')\right) \tag{1}$$

We choose a discount rate of $\delta = 0.9$ and a learning rate $\alpha = 0.3$, similar to a comparable (product) market simulation [36]. We provide further implementation details in Appendix A, including pseudocode for our market simulation.

4 Experimental Design

In this work, we run two main experiments: one where all workers are type H and one where there is a mixture of type L and type H workers. We will refer to the first experiment as *symmetric worker exploration* or Experiment S and the second as *asymmetric worker exploration* or Experiment A. Here, we give the details of each simulation run along with our methods for evaluating the simulations.

4.1 Simulation Details

In both experiments, we use $N_w = 100$, $N_f = 5$ [7], 5-discretized wage set $\mathcal{W}$, and $T = 20{,}000$ time steps. Workers are initially equally and randomly split among the firms, $D^{(0)}$ is initialized by workers randomly joining the pool IID with probability 0.5, and $S^{(0)}$ is initialized by firms randomly sharing their employee data IID with probability 0.5. The initial wages of workers are drawn from one of eight possible wage distributions over the set $\mathcal{W}$, ranging from left-skewed towards a wage of 0 to right-skewed towards a wage of 1, and details are given in Appendix A.1. Further, to implement the worker type difference in initial wages, we assume workers with wages in the bottom $1 - \tau_H$ percentile of the randomly drawn initial wages are type L and the rest are type H.

For Experiment S, we consider two market conditions: A market where negotiation is relatively riskier ($r = 0.5$) and one where it is relatively safer ($r = 0.25$). For Experiment A, we consider six market conditions: A riskier market with varying type distributions ($r = 0.5$ and $\tau_H = 0.25, 0.5, 0.75$) and a safer market with varying type distributions ($r = 0.25$ and $\tau_H = 0.25, 0.5, 0.75$). Each experiment consists of running the simulation T time steps per combination of information setting and market condition, repeating 15 times per initial wage distribution for a total of $N = 120$ runs.

4.2 Evaluation Methods

Simulation Output Evaluation. To compare the final wage outcomes of workers in each experiment, we consider two statistics: *average final wage of all workers* (in both experiments), denoted W_F, and *average final wage gap between worker types* (in experiment S only), denoted W_G. For each combination of information setting and market condition, we generate an empirical distribution of each statistic from the 120 simulation runs. First, W_F is defined as $W_F = \frac{1}{N_w}\sum_{i=1}^{N_w} w_i^{(T)}$, and note that $W_F \in [0, 1]$ where we take $W_F = 1$ to

[7] In Appendix C, we show our results extend qualitatively to more workers in the market $N_w \in \{250, 500\}$ and fewer firms $N_f \in \{2, 4\}$.

be the *fair* outcome where all workers are receiving a competitive wage. To evaluate this statistic, we compare the empirical CDFs $F_i(x)$ and $F_j(x)$ of W_F for information setting-market condition combinations i and j using the *Mann-Whitney U test* where the null hypothesis is $F_i = F_j$ and the alternative hypothesis is that F_i is stochastically greater than F_j. Thus, when the p-value of this test is significant ($p < 0.05$), we will say the final wages in combination j tend to be closer to the competitive wage of 1 than those in combination i.

Next, W_G is defined as $W_G = \frac{1}{N_H}\sum_{i=1}^{N_H} w_i^{(T)} - \frac{1}{N_L}\sum_{i=1}^{N_L} w_i^{(T)}$, where N_H is the number of type H workers and N_L is the number of type L workers. Note that $W_G \in [-1, 1]$ where we take $W_G = 0$ to be the *fair* outcome of no wage gap among worker types and, otherwise, $W_G > 0$ means type H workers are getting higher wages while $W_G < 0$ means type L workers are getting higher wages. To evaluate this statistic, we compare the empirical CDFs $F_i(x)$ and $F_j(x)$ of $|W_G|$ [8] for information setting-market condition combinations i and j again using the Mann-Whitney U test with the same null and alternative hypothesis as above. Here, when the p-value of this test is significant ($p < 0.05$), we will say wage gaps in combination i tend to be closer to 0 than wage gaps in combination j. We will also evaluate the mean, variance, and skewness of the empirical distributions of W_G to determine *which* worker type is likely to be advantaged when wage gaps exist.

Simulation Run Evaluation. We also track several statistics regarding characteristics of the firms' benchmark choices and workers' offer behaviors/outcomes throughout a single simulation run. For the *firms*, we track the proportion of timesteps where all possible benchmarks had identical ranges, at least two benchmark choices were not identical and the firm chose the one with the widest range, and at least two benchmark choices were not identical and the firm did not chose the one with the widest range. For the *workers*, we track the proportion of timesteps where the worker made no offer, an offer that was accepted, and an offer that was rejected. For each run, we average these proportions across all firms and workers, respectively, and then report the average value across all $N = 120$ runs.

5 Simulation Results and Analysis

For each experiment, we compare the following across settings: final worker wage outcomes, worker offer behavior, and firm salary benchmark choice. An overview of our findings for each research question can be found in Table 1.

5.1 Experiment S - Simulation Results and Analysis

In this experiment, the workers are all symmetric in the sense that they use the same $\epsilon^{(t)}$ decay rate in Algorithm 1. We compare information settings under the two market conditions considered, $r = 0.5$ and $r = 0.25$.

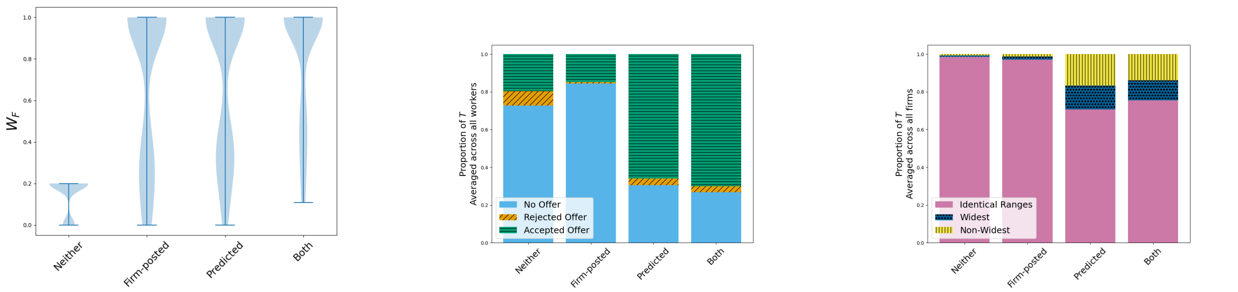

(a) W_F distribution comparison. $r = 0.5$ (b) Worker offer comparison. $r = 0.5$ (c) Firm benchmark comparison. $r = 0.5$

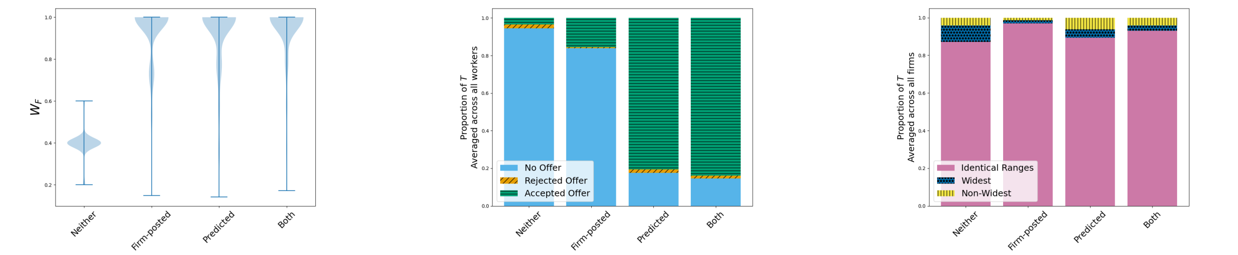

(d) W_F distribution comparison. $r = 0.25$ (e) Worker offer comparison. $r = 0.25$ (f) Firm benchmark comparison. $r = 0.25$

Fig. 2. Experiment S: W_F empirical distribution comparison, worker (average) offer behavior comparison, and firm (average) benchmark choice comparison across the information settings for $r = 0.5$ (top row) and $r = 0.25$ (bottom row).

[8] $|\cdot|$ is the absolute value function. This allows us to evaluate how likely wage gaps (in either direction) are in each setting.

RQ	Findings	Takeaway
RQ 1	The setting with firm-posted salary ranges *and* a salary prediction tool had significantly higher wages ($p < 0.05$) than all other settings in 3 out of 8 market conditions. The setting with *only* firm-posted salary ranges had significantly higher wages ($p < 0.05$) than the setting with *only* a salary prediction tool in 1 out of 8 market conditions. Settings *without* firm-posted salary ranges never had significantly higher wages ($p < 0.05$) than settings *with* them in 8 out of 8 market conditions.	Settings with firm-posted salary ranges are more likely to have *wages closer to the competitive wage* than those without.
RQ 2	The setting with *only* firm-posted salary ranges had significantly smaller gaps ($p < 0.05$) than the setting with *only* a salary prediction tool in 4 out of 6 market conditions. The setting with firm-posted salary ranges *and* a salary prediction tool had significantly smaller gaps ($p < 0.05$) than the setting with *only* a salary prediction tool in 3 out of 6 market conditions. The two settings with salary prediction tools resulted in wage gap distributions with larger skewness towards type H workers than the setting with *only* firm-posted salary ranges in 4 out of 6 market conditions (salary prediction tool *and* firm-posted salary ranges) and 6 out of 6 market conditions (*only* salary prediction tool), respectively.	Settings with firm-posted salary ranges *have smaller wage gaps* than the setting with only a salary prediction tool. Settings with salary prediction tools are more likely to have *wage gaps that favor workers with a historical wage advantage* than the setting with only firm-posted salary ranges.
RQ 3	The average worker made no offer 84-92% of the simulation run in the setting with *only* firm-posted salary ranges and 12-42% of the simulation in settings with salary prediction tools. The average firm saw non-identical salary benchmarks 97-99% of the simulation run in the setting with *only* firm-posted salary ranges and 66-96% of the simulation run in settings with salary prediction tools.	Settings with salary prediction tools are more likely to have *workers making more offers* and *firms choosing between non-identical salary benchmarks* than the setting with only firm-posted salary ranges.

Table 1. Summary of results for each research question.

Outcome comparison. In Figures 2a and 2d, the empirical distributions of W_F are given for the market conditions $r = 0.5$ and $r = 0.25$, respectively, as violin plots. Clearly, `Neither` leads to the worst W_F values since most outcomes are clustered around 0.2 ($r = 0.5$) or 0.4 ($r = 0.25$) while the other three settings have W_F values that skew towards 1, the competitive wage, though with larger variance. When $r = 0.5$, `Both` tends to have larger W_F values than `Neither` ($p < 0.05$), `Firm-posted` ($p < 0.05$), and `Predicted` ($p < 0.05$). However, when $r = 0.25$, `Both` only has significantly larger W_F values than `Neither` ($p < 0.05$). Thus, **introducing a salary prediction tool along with firm-posted ranges benefits workers when negotiations are relatively riskier whereas any kind of pay transparency benefits workers when negotiations are relatively safer.**

Worker strategy analysis. A worker's average offer characteristics are given in Figure 2b and Figure 2e for $r = 0.5$ and $r = 0.25$, respectively. These figures show that in both market conditions workers, on average, make more offers in `Predicted` and `Both` (70-73% of the simulation run for $r = 0.5$ and 82-85% for $r = 0.25$) than `Firm-posted` and `Neither` (15-28% of the simulation run for $r = 0.5$ and 5-16% for $r = 0.25$). In `Firm-posted`, `Predicted`, and `Both`, workers make more offers when $r = 0.25$ than when $r = 0.5$ whereas `Neither` shows the opposite trend, and we further examine this trend in the *firm strategy analysis* below. Additionally, in Figures 2b and 2e, `Both` shows slightly more favorable offers than `Predicted` where offers are accepted about 4% more of the simulation run in the former setting, suggesting one reason why W_F is more likely to be larger in `Both` than `Predicted`. Thus, **the presence of a salary prediction tool encourages all workers to negotiate more often than in settings where they do not have access to this tool**. Further, since `Both` led to significantly higher W_F values than `Firm-posted` when $r = 0.5$, **even in a risky market, encouraging more negotiations with a salary prediction tool can improve outcomes for workers beyond firms posting salary ranges alone.**

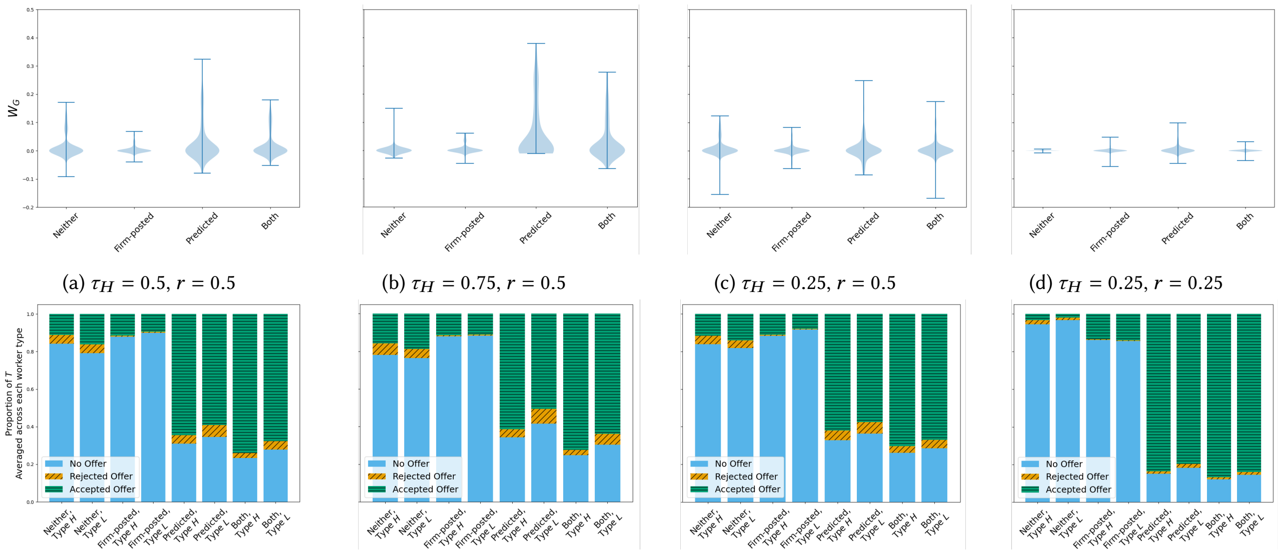

(a) $\tau_H = 0.5, r = 0.5$ (b) $\tau_H = 0.75, r = 0.5$ (c) $\tau_H = 0.25, r = 0.5$ (d) $\tau_H = 0.25, r = 0.25$

(e) $\tau_H = 0.5, r = 0.5$. (f) $\tau_H = 0.75, r = 0.5$ (g) $\tau_H = 0.25, r = 0.5$. (h) $\tau_H = 0.25, r = 0.25$

Fig. 3. Experiment A: W_G empirical distribution comparison (top row) and worker type (average) offer behavior comparison (bottom row) across information settings for ($\tau_H = 0.5, r = 0.5$), ($\tau_H = 0.75, r = 0.5$), ($\tau_H = 0.25, r = 0.5$), and ($\tau_H = 0.25, r = 0.25$).

Firm strategy analysis. A firm's average benchmark choice characteristics are given in Figure 2c and Figure 2f for $r = 0.5$ and $r = 0.25$, respectively. First, for both market conditions, in `Firm-posted`, firms almost always see identical benchmark ranges throughout the simulation (about 97% of the time) while in `Predicted` and `Both`, firms see identical benchmarks less often (about 70-93% of the time). This indicates that **the presence of the salary prediction tool encourages diverse salary benchmarks for firms**. However, when $r = 0.25$, we observe that firms see identical benchmarks even less often in `Neither` (about 87% of the time) than `Predicted` and `Both` (about 90-93% of the time). To investigate the effects of this trend, we next analyze *which* benchmark firms chose when benchmarks were non-identical, i.e., the one with the largest range or not.

In `Predicted` and `Both`, firms choose the non-widest range more often than the widest one (non-widest ranges are chosen 57-60% of the time when ranges are non-identical, see Figures 2c and 2f). This is in contrast with `Neither` when $r = 0.25$ where firms are more often choosing the widest range for their benchmark (non-widest ranges are chosen only 33% of the time when ranges are non-identical, see Figure 2f). This feature may explain why, counterintuitively, workers in `Neither` do not negotiate as often when $r = 0.25$ (Figure 2e) as they do when $r = 0.5$ (Figure 2b). In `Neither`, workers are essentially choosing firms to negotiate with randomly and, since firms are choosing wide ranges more often in this setting, then it is more difficult for workers to learn which offers to make. In fact, we observe that, when workers make offers, they are more likely to be rejected when $r = 0.25$ than when $r = 0.5$ (37% of the time compared to 28% of the time, see Figures 2e and 2b). Further, since `Both` and `Predicted` lead to significantly larger W_F values than `Neither`, **workers benefit from diverse salary benchmarks for firms when salary ranges are posted, either directly from the firms or estimates from the salary prediction tool.**

5.2 Experiment A - Asymmetric Worker Exploration

In this experiment, the workers are asymmetric in their *exploration ability*: Type H workers use a slower ϵ decay rate in Algorithm 1 while type L workers use a faster ϵ decay rate in Algorithm 1. We compare information settings under six market conditions in this experiment: A relatively *riskier* market ($r = 0.5$) and a relatively *safer* market, each under three proportions of worker types ($\tau_H = 0.25, \tau_H = 0.5, \tau_H = 0.75$).

Outcome comparison. We found similar empirical distributions of W_F in Experiment A across all market conditions in $r = 0.5$ and $r = 0.25$ to those in Figure 2a and 2d, respectively, and we provide the violin plots in Appendix B.1 for completeness. Further, similar statistical significance held compared to Experiment S: `Neither` has significantly lower wages compared to each of the other information settings ($p < 0.05$ for each), `Both` has significantly larger wages than `Predicted` in all $r = 0.5$ market conditions ($p < 0.05$), and `Both` loses its significance over other information settings in all $r = 0.25$ market conditions ($p > 0.05$). However, different from Experiment S, when $r = 0.5$, `Both` no longer has significantly larger wages than `Firm-posted` ($p > 0.05$), and `Firm-posted` now has significantly larger wages than `Predicted` in one market condition ($r = 0.5$ and $\tau_H = 0.25$, $p < 0.05$). This suggests that **requiring firm-posted salary ranges is more important than adding a salary prediction tool to a market with neither for worker wages to be closer to the competitive level.**

Next, the empirical distributions of W_G from 4 market conditions are given as violin plots in Figure 3, and the plots for the other 2 conditions and additional statistics for this experiment can be found in Appendix B.4. The analysis of the empirical distribution of $|W_G|$ reveals `Firm-posted` tends to have wage gaps closer to 0 than `Predicted` in 4/6 market conditions (Figures 3a to 3d, $p < 0.05$) and `Both` tends to have wage gaps closer to 0 than `Predicted` in 3/6 market conditions (Figures 3b to 3d, $p < 0.05$). Thus, **settings with firm-posted salary ranges are likely to have smaller wage gaps than settings with only a salary prediction tool**. Next, in `Predicted`, the empirical distributions of W_G have positive mean and skewness larger than that of `Firm-posted` in 6/6 market conditions (see Table 7 in Appendix B.4). Further, in `Both`, the empirical distributions of W_G have positive mean and skewness larger than that of `Firm-posted` in 4/6 market conditions. Figures 3a and 3b show two market conditions where both `Predicted` and `Both` are more positively skewed than `Firm-posted`. Figures 3c and 3d show the two market conditions where only `Predicted` is more positively skewed than `Firm-posted`. Thus, **settings with salary prediction tools are more likely to have wage gaps that favor type H workers than settings with firm-posted salary ranges only**.

Worker strategy analysis. Each worker type's average offer characteristics are given in Figures 3e to 3h. First, compared to experiment S, workers, on average, still make more offers in `Predicted` and `Both` (58-77% of the simulation run for $r = 0.5$ and 81-88% for $r = 0.25$) than in `Neither` and `Firm-posted` (8-24% of the simulation run for $r = 0.5$ and 3-17% for $r = 0.25$). Further, workers still make more offers when $r = 0.25$ than when $r = 0.5$ in `Firm-posted`, `Predicted`, and `Both` whereas the opposite trend still holds for `Neither` (see, e.g., Figures 3e to 3h).

Comparing the offer behavior of worker types, we observe that `Firm-posted` generally has more similar offer behavior between worker types (i.e., similar bar graph heights) across all market conditions than `Predicted` and `Both`. In particular, across all market conditions, the difference between proportion of time steps with no offers between the average type H worker and the average type L worker ranged between 0-3% in `Firm-posted` and ranged between 2-7% in `Predicted` and `Both`. Further, across all market conditions, the difference between the fraction of offers accepted over offers made between worker types ranged between 0-2% in `Firm-posted` and 1-7% in `Predicted` and `Both`. **These trends support the fact that `Firm-posted` is less likely to produce wage gaps than settings with a salary prediction tool across market conditions**. Additionally, in `Predicted` and `Both`, type L workers always negotiate less and have a smaller fraction of accepted offers over offers made compared to type H workers across all market conditions. **This trend possibly explains why settings with salary prediction tools are more likely to have W_G distributions skewed towards type H workers than `Firm-posted`**.

6 Conclusion

In this work, we set up a model to evaluate the role of salary prediction tools in the growing call for pay transparency, with a particular focus on how such tools may affect worker wages overall (W_F) and wage gaps between worker types associated with demographic differences (W_G).

Takeaways. Across all of our experiments, it appears that the main role of the salary prediction tool is to *encourage workers to make more offers in the market to introduce diverse information.* In theory, diverse information should encourage competition between firms and therefore boost worker wages. However, the presence of diverse information is not necessarily effective on its own: diverse information leads to more competitive wages *when communicated through a channel that both firms and workers find credible*, such as firm-posted salary ranges. Salary prediction tools on their own are less effective and can lead to wage gaps along demographic lines. Specifically, when workers learn to negotiate in different ways, i.e., one group can explore more than another, then *encouraging more negotiations seems to be more beneficial for the group that has the ability to explore more.* This suggests that policies that *equalize exploration costs* for groups that are known to typically have higher exploration costs can boost the effectiveness of salary prediction tools, e.g., providing salary negotiation coaching for women and other disadvantaged groups [22]. However, these kinds of interventions can be costly to develop and scale to different groups of workers, so, while it is useful to generate diverse market information and competition through encouraged negotiation, our results suggest it may be more fruitful to invest in policies that *set up channels of credible information communication between firms and workers.* Overall, our findings reinforce the growing policy calls for firms to post salary ranges along with their job descriptions. Even within a single job domain, workers face a large and costly space to explore, and removing the decision of *which information sources firms are likely to find credible in negotiations* by requiring firms to post the salary range they plan to use for the job makes it that much easier for all workers to prosper in the labor market.

Limitations and Future Work. In this work, we developed a simple model in order to isolate the effect of information source asymmetries as firms and workers are learning salary negotiation strategies, but there are a number of real-world complexities to consider for future work. In particular, our negotiation model is constrained: Only workers make offers, there are no counteroffers, workers cannot leverage offers against each other, and there is no job differentiation beyond wages. Next, firms often have more complicated strategies for compensation benchmarking, including designing benefits beyond base pay, so a more detailed study of these considerations is necessary to understand the incentives that lead to firms to react to pay transparency policies more or less competitively. Additionally, we only considered a model of firms and workers sharing information insofar as it helps them in negotiation later on. In reality, there are many reasons firms and workers may disclose information online, such as pay disclosure restrictiveness norms and fairness uncertainty salience as noted by previous research [9, 10]. Finally, as with any simulation study, our results are limited to the assumptions we make in our modeling choices, and real-world studies can further elucidate the role that salary prediction tools play in determining wage outcomes. We leave these questions to future work, and we view this simulation study as an important preliminary step to formulating and answering these questions.

References

[1] David Arnold, Simon Quach, and Bledi Taska. 2025. *The impact of pay transparency in job postings on the labor market.* Technical Report. National Bureau of Economic Research.

[2] Orley Ashenfelter, David Card, Henry Farber, and Michael R Ransom. 2022. Monopsony in the labor market: New empirical results and new public policies. *Journal of Human Resources* 57, S (2022), S1–S10.

[3] José Azar and Ioana Marinescu. 2024. Monopsony power in the labor market: From theory to policy. *Annual Review of Economics* 16, 1 (2024), 491–518.

[4] Michael Baker, Yosh Halberstam, Kory Kroft, Alexandre Mas, and Derek Messacar. 2023. Pay transparency and the gender gap. *American Economic Journal: Applied Economics* 15, 2 (2023), 157–183.

[5] Honey Batra, Amanda Michaud, and Simon Mongey. 2023. *Online job posts contain very little wage information.* Technical Report. National Bureau of Economic Research.

[6] Eric Blankmeyer, James P LeSage, JR Stutzman, Kris Joseph Knox, and R Kelley Pace. 2011. Peer-group dependence in salary benchmarking: a statistical model. *Managerial and Decision Economics* 32, 2 (2011), 91–104.

[7] Francine D Blau. 2025. Gender Inequality in the Labor Market: Continuing Progress? *ILR Review* 78, 2 (2025), 275–303.

[8] Emily Breza, Supreet Kaur, and Yogita Shamdasani. 2018. The morale effects of pay inequality. *The Quarterly Journal of Economics* 133, 2 (2018), 611–663.

[9] Michelle Brown, Peter Bamberger, Paul D Bliese, and John Shields. 2023. Fairness uncertainty and pay information exchange: Why and when employees disclose bonus pay to pay information websites. *Journal of organizational behavior* 44, 9 (2023), 1362–1379.

[10] Michelle Brown, Anthony J Nyberg, Ingo Weller, and Sam D Strizver. 2022. Pay information disclosure: Review and recommendations for research spanning the pay secrecy–pay transparency continuum. *Journal of Management* 48, 6 (2022), 1661–1694.

[11] Katharina Brütt and Huaiping Yuan. 2022. *Pitfalls of pay transparency: Evidence from the lab and the field.* Technical Report. Tinbergen Institute Discussion Paper.

[12] Sydnee Caldwell, Ingrid Haegele, and Jörg Heining. 2025. *Bargaining and inequality in the labor market.* Technical Report. National Bureau of Economic Research.

[13] Sydnee Caldwell and Nikolaj Harmon. 2019. Outside options, bargaining, and wages: Evidence from coworker networks. *Unpublished manuscript, Univ. Copenhagen* (2019), 203–207.

[14] Emilio Calvano, Giacomo Calzolari, Vincenzo Denicolo, and Sergio Pastorello. 2020. Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review* 110, 10 (2020), 3267–3297.

[15] Ruxin Chen. 2025. Deep Reinforcement Learning in a Search-Matching Model of Labor Market Fluctuations. *Economies* 13, 10 (2025), 302.

[16] Ruxin Chen and Zeqiang Zhang. 2025. Deep Reinforcement Learning in Labor Market Simulations. In *2025 IEEE Symposium on Computational Intelligence for Financial Engineering and Economics (CiFer).* IEEE, 1–7.

[17] Maxime C Cohen, Eddy Hage-Youssef, and Warut Khern-am nuai. 2025. When AI Sets Wages: Biases and Labor Discrimination in Generative Pricing. *Available at SSRN 5404966* (2025).

[18] Zoë Cullen. 2024. Is pay transparency good? *Journal of Economic Perspectives* 38, 1 (2024), 153–180.

[19] Zoë Cullen and Ricardo Perez-Truglia. 2023. The salary taboo privacy norms and the diffusion of information. *Journal of Public Economics* 222 (2023), 104890.

[20] Zoe B Cullen, Shengwu Li, and Ricardo Perez-Truglia. 2022. *What's my employee worth? The effects of salary benchmarking.* Technical Report. National Bureau of Economic Research.

[21] Zoë B Cullen and Bobak Pakzad-Hurson. 2023. Equilibrium effects of pay transparency. *Econometrica* 91, 3 (2023), 765–802.

[22] Zoë B Cullen, Bobak Pakzad-Hurson, and Ricardo Perez-Truglia. 2025. *Pushing the Envelope: The Effects of Salary Negotiations.* Technical Report. National Bureau of Economic Research.

[23] Herbert Dawid, Philipp Harting, and Michael Neugart. 2023. Implications of Algorithmic Wage Setting on Online Labor Platforms: a Simulation-Based Analysis. (2023).

[24] Veda Duddu, Jash Rajesh Parekh, Andy Mao, Hanyi Min, Ziang Xiao, Vedant Das Swain, and Koustuv Saha. 2025. Does AI Coaching Prepare us for Workplace Negotiations? *arXiv preprint arXiv:2509.22545* (2025).

[25] Deborah Thompson Eisenberg. 2011. Money, sex, and sunshine: A market-based approach to pay discrimination. *Ariz. St. LJ* 43 (2011), 951.

[26] Stefan Eriksson and Jonas Lagerström. 2012. The labor market consequences of gender differences in job search. *Journal of Labor Research* 33, 3 (2012), 303–327.

[27] Cynthia Estlund. 2014. Extending the case for workplace transparency to information about pay. *UC Irvine L. Rev.* 4 (2014), 781.

[28] Sara Fish, Yannai A Gonczarowski, and Ran I Shorrer. 2024. Algorithmic collusion by large language models. *arXiv preprint arXiv:2404.00806* 7 (2024).

[29] Christopher Flinn and Joseph Mullins. 2021. Firms' choices of wage-setting protocols. In *Technical Reports, Discussion paper.* New York University.

[30] Wolfgang Frimmel, Bernhard Schmidpeter, Rene Wiesinger, and Rudolf Winter-Ebmer. 2023. External pay transparency and the gender wage gap. (2023).

[31] R Stuart Geiger, Flynn O'Sullivan, Elsie Wang, and Jonathan Lo. 2025. Asking an AI for salary negotiation advice is a matter of concern: Controlled experimental perturbation of ChatGPT for protected and non-protected group discrimination on a contextual task with no clear ground truth answers. *PloS one* 20, 2 (2025), e0318500.

[32] Pijush Ghorai and Rupashri Barik. 2024. Salary Prediction Using Machine Learning Techniques. In *International Conference on Data & Information Sciences.* Springer, 201–213.

[33] Simon Jäger, Christopher Roth, Nina Roussille, and Benjamin Schoefer. 2024. Worker beliefs about outside options. *The Quarterly Journal of Economics* 139, 3 (2024), 1505–1556.

[34] Nathanael Jo, Kathleen Creel, Ashia Wilson, and Manish Raghavan. 2025. Homogeneous Algorithms Can Reduce Competition in Personalized Pricing.

[35] Marios Karabarbounis and Santiago M Pinto. 2018. What can we learn from online wage postings? Evidence from Glassdoor. *Evidence from Glassdoor* (2018).

[36] Timo Klein. 2021. Autonomous algorithmic collusion: Q-learning under sequential pricing. *The RAND Journal of Economics* 52, 3 (2021), 538–558.

[37] Rakesh Kochhar. 2023. *The Enduring Grip of the Gender Pay Gap.* Technical Report. Pew Reserach Center.

[38] Marianne DelPo Kulow. 2013. Beyond the paycheck fairness act: Mandatory wage disclosure laws-a necessary tool for closing the residual gender wage gap. *Harv. J. on Legis.* 50 (2013), 385.

[39] N Gregory Mankiw. 1998. *Principles of microeconomics.* Vol. 1. Elsevier.

[40] Dani M Marinova. 2016. On the use of crowdsourcing labor markets in research. *Perspectives on Politics* 14, 2 (2016), 422–431.

[41] Alexandre Mas. 2017. Does transparency lead to pay compression? *Journal of Political Economy* 125, 5 (2017), 1683–1721.

[42] Jonathan Masur and Eric A Posner. 2023. Horizontal collusion and parallel wage setting in labor markets. *U. Chi. L. Rev.* 90 (2023), 545.

[43] Qingxin Meng, Hengshu Zhu, Keli Xiao, and Hui Xiong. 2018. Intelligent salary benchmarking for talent recruitment: A holistic matrix factorization approach. In *2018 IEEE International Conference on Data Mining (ICDM)*. IEEE, 337–346.

[44] Armando Miano. 2023. Search costs, outside options, and on-the-job search. *Job market paper, Harvard University* (2023).

[45] Suresh Naidu and Arindrajit Dube. 2024. Monopsony power in labor markets. *NBER Reporter* 1 (2024), 7–12.

[46] Renato Nazzini and James Henderson. 2024. Overcoming the Current Knowledge Gap of Algorithmic" Collusion" and the Role of Computational Antitrust. *Stan. Computational Antitrust* 4 (2024), 1.

[47] U.S. Department of Justice and Federal Trade Commission. 2025. Antitrust Guidelines for Business Activities Affecting Workers. https://www.ftc.gov/system/files/ftc_gov/pdf/p251201antitrustguidelinesbusinessactivitiesaffectingworkers2025.pdf

[48] Glassdoor Team. 2016. *Introducing Know Your Worth by Glassdoor.* https://www.glassdoor.com/blog/introducing-know-worth-glassdoor/

[49] Payscale Team. 2025. *The pay confidence gap report.* https://www.payscale.com/featured-content/pay-confidence-gap

[50] Payscale Team. 2026. *2026 Compensation Best Practices Report.* https://www.payscale.com/featured-content/cbpr?utm_medium=cpc&utm_source=google&utm_campaign=all_prs_us-ca_srch-brd-ps_wv_evg&utm_content=dmo_brd_rqf_ps-brand-rad_ta_169369034340_716186191120_&utm_term=pay%20scale_b&gclid=Cj0KCQjw2MbPBhCSARIsAP3jP9wDj8PX9MzCda4EALUClrmrWA_Bfpxlw2lG_cTf-AW2C66tT9nOXPgaAnbEEALw_wcB#salary-data

[51] Payscale Team. 2026. *Payfactors Tool.* https://www.payscale.com/products/software/payfactors?utm_medium=cpc&utm_source=google&utm_campaign=all_prs_us-ca_srch-brd-ps_wv_evg&utm_content=dmo_brd_rqf_ps-brand-rad_ta_169369034340_716186191120_&utm_term=pay%20scale_b&gclid=Cj0KCQjw2MbPBhCSARIsAP3jP9wDj8PX9MzCda4EALUClrmrWA_Bfpxlw2lG_cTf-AW2C66tT9nOXPgaAnbEEALw_wcB

[52] Salary.com Team. 2026. *AI built for HR and compensation.* https://www.salary.com/solutions/ai

[53] Christopher John Cornish Hellaby Watkins et al. 1989. Learning from delayed rewards. (1989).

A Simulation Implementation Details

All experiments were run on a standard laptop (Apple Macbook air, M3 chip, 2024).

A.1 Initial Wage Distribution Generation

Each experiment generates wages from eight possible initial wage distributions to ensure diverse starting conditions. We describe the probability mass functions of these distributions over $\mathcal{W}$ and provide an example generation of each in Figure 4. In our experiments, we ensure at least one worker starts with wage equal to 1 so that there is always a chance for firms to encounter an offer of 1 immediately. Additionally, in Experiment A, formalizing our assumption about worker type differences, we set the type of workers with the bottom τ_L wages to be type L and the rest to be type H.

(1) Left-skewed

$$P(0) = 0.75, P(0.2) = 0.125, P(0.4) = 0.025, P(0.6) = P(0.8) = P(1.0) = 0.033$$

(2) Slightly left-skewed

$$P(0) = 0.5, P(0.2) = 0.125, P(0.4) = 0.075, P(0.6) = P(0.8) = P(1.0) = 0.1$$

(3) Uniform

$$P(0) = P(0.2) = P(0.4) = P(0.6) = P(0.8) = P(1.0) = 0.167$$

(4) Slightly right-skewed

$$P(0) = P(0.2) = P(0.4) = 0.1, P(0.6) = 0.5, P(0.8) = 0.125, P(1.0) = 0.075$$

(5) Right-skewed

$$P(0) = P(0.2) = P(0.4) = 0.033, P(0.6) = 0.025, P(0.8) = 0.125, P(1.0) = 0.75$$

(6) Bimodal, even

$$P(0) = 0.4, P(0.2) = P(0.4) = P(0.6) = P(0.8) = 0.05, P(1.0) = 0.4$$

(7) Bimodal, left-skewed

$$P(0) = 0.45, P(0.2) = P(0.4) = P(0.6) = P(0.8) = 0.05, P(1.0) = 0.35$$

(8) Bimodal, right-skewed

$$P(0) = 0.35, P(0.2) = P(0.4) = P(0.6) = P(0.8) = 0.05, P(1.0) = 0.45$$

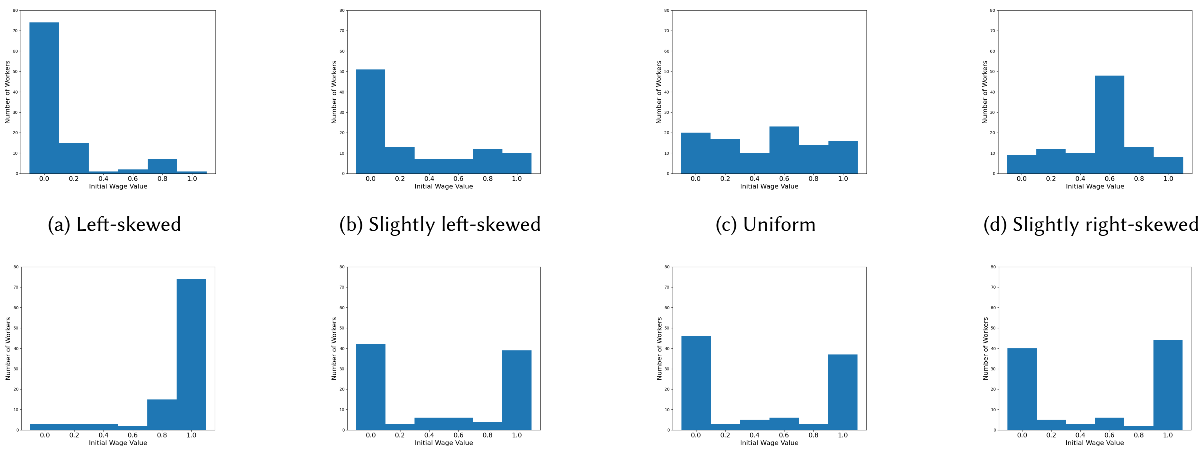

(a) Left-skewed (b) Slightly left-skewed (c) Uniform (d) Slightly right-skewed

(e) Right-skewed (f) Bimodal, even (g) Bimodal, left-skewed (h) Bimodal, right-skewed

Fig. 4. Example initial wage distributions over 5-discretized $\mathcal{W}$ for $N_w = 100$ workers.

A.2 MDP State, Action, and Reward Details

States + Actions. Workers share the same set of possible states and actions, denoted $\mathcal{S}_w$ and $\mathcal{A}_w$, respectively, as do firms, denoted $\mathcal{S}_f$ and $\mathcal{A}_f$, respectively. Let $\mathcal{A} = \prod_{i \in [N_w]} \mathcal{A}_w \times \prod_{j \in [N_f]} \mathcal{A}_f$ be the set of all possible combinations of actions by all agents, and let $A^{(t)} \in \mathcal{A}$ be the actions taken at time t. The state and action sets of each agent formalize the overviews given in Figures ?? and 1. We describe the details for each set of agents below.

For the workers, in `Neither` and `Firm-posted`, $\mathcal{S}_w$ is partitioned into two groups of states: `firm choice` states and `offer choice` states. In `Predicted` and `Both` there are also `sharing` states where workers decide whether to share $w_i^{(t)}$ with the salary prediction tool. First, $S \in \mathcal{S}_{w,\texttt{sharing}}$ is a tuple of $w_i^{(t)}$, the minimum of the posted lower bounds, and the maximum of the posted upper bounds. The later two values are defined as:

$$\underline{\ell}^{(t)} = \begin{cases} \min\{\hat{\ell}_j^{(t)} : j \in [N_f]\} & \texttt{Predicted} \\ \min\{b_{j,10}^{(t)} : j \in [N_f]\} & \texttt{Both} \end{cases}$$

and

$$\overline{u}^{(t)} = \begin{cases} \max\{\hat{u}_j^{(t)} : j \in [N_f]\} & \texttt{Predicted} \\ \max\{b_{j,90}^{(t)} : j \in [N_f]\} & \texttt{Both} \end{cases}$$

Then, $\mathcal{A}_{w,\texttt{sharing}} = \{\texttt{share}, \texttt{no share}\}$ where workers learn about $\hat{m}^{(t)}$ if they choose `share`. We assume all workers are optimistic about the salary prediction tool and use $\max\{\hat{m}^{(t)}, w_i^{(t)}\}$ as their *reservation wage*, i.e., the minimum wage they'll accept in negotiations. This is defined as:

$$w_{i,r}^{(t)} = \begin{cases} \max\{w_i^{(t)}, \hat{m}^{(t)}\} & \texttt{Predicted or Both and } s_i^{(t)} = \texttt{share} \\ w_i^{(t)} & s_i^{(t)} = \texttt{no share} \end{cases}$$

Next, $S \in \mathcal{S}_{w,\texttt{firm choice}}$ contains a single value relevant to the worker regarding the pay range of each firm. We make the assumption that workers prefer ranges where they *know they have a chance to increase their wage*, so if the firm does not post a range or if the firm's upper bound does not exceed the worker's reservation wage, we use the flag -1 as that firm's value. Otherwise, we use the difference between the posted η_{90} and η_{10} values of the firm. This value is defined as follows:

$$d_j^{(t)} = \begin{cases} \hat{u}_j^{(t)} - \hat{\ell}_j^{(t)} & \texttt{Predicted and } \hat{u}_j^{(t)} > w_{i,r}^{(t)} \\ b_{j,90}^{(t)} - b_{j,10}^{(t)} & \texttt{Firm-posted and Both and } b_{j,90}^{(t)} > w_{i,r}^{(t)} \\ -1 & \text{otherwise} \end{cases}$$

Then, each worker chooses a firm to negotiate with, i.e., $\mathcal{A}_{w,\texttt{firm choice}} = \{j : j \in [N_f]\}$.

Finally, $\mathcal{S}_{w,\texttt{offer choice}}$ contains the lower and upper bound of offers they can make, denoted $\underline{o}$ and $\overline{o}$, respectively. These values are set conditional on $w_{i,r}^{(t)}$ and the η_{90} value of the posted range (predicted or verified) of the chosen firm, as follows:

$$\underline{o}_i^{(t)} = w_{i,r}^{(t)}$$

and

$$\overline{o}_i^{(t)} = \begin{cases} \max\{\overline{b}_{x_i^{(t)},90}^{(t)}, \min\{1, w_{i,r}^{(t)} + \frac{1}{k}\}\} & \texttt{Firm-posted and Both and } d_{x_i^{(t)}}^{(t)} \neq -1 \\ \max\{\hat{u}_{x_i^{(t)}}^{(t)}, \min\{1, w_{i,r}^{(t)} + \frac{1}{k}\}\} & \texttt{Predicted and } d_{x_i^{(t)}}^{(t)} \neq -1 \\ \min\{1, w_{i,r}^{(t)} + \frac{1}{k}\} & d_{x_i^{(t)}}^{(t)} = -1 \end{cases}$$

Here, workers always use their reservation wage as the lower bound and consider offers at least one wage interval higher than their lower bound, up to a firm's posted η_{90} value. We also assume workers never make offers above the MRPL of 1. Thus, $\mathcal{A}_{w,\texttt{offer choice}} = \{\{w : w \geq \underline{o}, w \leq \overline{o}\} \cup \{\texttt{no offer}\} : (\underline{o}, \overline{o}) \in \mathcal{S}_{w,\texttt{offer choice}}\}$.

The full set of states and actions of the workers in each information setting are described in Tables 2 and 3, respectively.

	$\mathcal{S}_{w,\texttt{sharing}}$	$\mathcal{S}_{w,\texttt{firm choice}}$	$\mathcal{S}_{w,\texttt{offer choice}}$
`Neither`	-	$(d_j^{(t)} : j \in [N_f])$	$(\underline{o}_i^{(t)}, \overline{o}_i^{(t)})$
`Firm-posted`	-	$(d_j^{(t)} : j \in [N_f])$	$(\underline{o}_i^{(t)}, \overline{o}_i^{(t)})$
`Predicted`	$(w_i^{(t)}, \underline{\ell}^{(t)}, \overline{u}^{(t)})$	$(d_j^{(t)} : j \in [N_f])$	$(\underline{o}_i^{(t)}, \overline{o}_i^{(t)})$
`Both`	$(w_i^{(t)}, \underline{\ell}^{(t)}, \overline{u}^{(t)})$	$(d_j^{(t)} : j \in [N_f])$	$(\underline{o}_i^{(t)}, \overline{o}_i^{(t)})$

Table 2. Worker states in each information setting.

For the firms, in all information settings, $\mathcal{S}_f$ is partitioned into two groups of states: `benchmark choice` states and `acceptance threshold choice` states. First, $S \in \mathcal{S}_{f,\texttt{benchmark choice}}$ contains the 10th percentile wage and the 90th percentile wage for each benchmark in $\mathcal{B}_j^{(t)}$. Then, $\mathcal{A}_{f,\texttt{benchmark choice}} = \{\texttt{independent, shared, predicted}\}$, where `predicted` is only an option in information settings with a salary prediction tool. Next, $S \in \mathcal{S}_{f,\texttt{acceptance threshold choice}}$ is given as $\left(b_{j,10}^{(t)}, b_{j,50}^{(t)}, b_{j,90}^{(t)}\right)$. These values also constitute the action set $A \in \mathcal{A}_{f,\texttt{acceptance threshold choice}}$, and the firm chooses one (or two, if there are two worker types) of these values as their private acceptance threshold.

The states and actions of the firms in each information setting are described in Tables 4 and 5, respectively.

	$\mathcal{A}_{w,\texttt{sharing}}$	$\mathcal{A}_{w,\texttt{firm choice}}$	$\mathcal{A}_{w,\texttt{offer choice}}$
`Neither`	-	$\{j : j \in [N_f]\}$	$\{w : w \in \mathcal{W}, w \geq \underline{o}_i^{(t)}, w \leq \overline{o}_i^{(t)}\}$
`Firm-posted`	-	$\{j : j \in [N_f]\}$	$\{w : w \in \mathcal{W}, w \geq \underline{o}_i^{(t)}, w \leq \overline{o}_i^{(t)}\}$
`Predicted`	{`share`, `no share`}	$\{j : j \in [N_f]\}$	$\{w : w \in \mathcal{W}, w \geq \underline{o}_i^{(t)}, w \leq \overline{o}_i^{(t)}\}$
`Both`	{`share`, `no share`}	$\{j : j \in [N_f]\}$	$\{w : w \in \mathcal{W}, w \geq \underline{o}_i^{(t)}, w \leq \overline{o}_i^{(t)}\}$

Table 3. Worker actions in each information setting.

	$\mathcal{S}_{f,\texttt{benchmark choice}}$	$\mathcal{S}_{f,\texttt{acceptance threshold choice}}$
`Neither`	$\left((\ell_j^{(t)}, u_j^{(t)}), (\ell_s^{(t)}, u_s^{(t)})\right)$	$(b_{j,10}^{(t)}, b_{j,50}^{(t)}, b_{j,90}^{(t)})$
`Firm-posted`	$\left((\ell_j^{(t)}, u_j^{(t)}), (\ell_s^{(t)}, u_s^{(t)})\right)$	$(b_{j,10}^{(t)}, b_{j,50}^{(t)}, b_{j,90}^{(t)})$
`Predicted`	$\left((\ell_j^{(t)}, u_j^{(t)}), (\ell_s^{(t)}, u_s^{(t)}), (\hat{\ell}^{(t)}, \hat{u}^{(t)})\right)$	$(b_{j,10}^{(t)}, b_{j,50}^{(t)}, b_{j,90}^{(t)})$
`Both`	$\left((\ell_j^{(t)}, u_j^{(t)}), (\ell_s^{(t)}, u_s^{(t)}), (\hat{\ell}^{(t)}, \hat{u}^{(t)})\right)$	$(b_{j,10}^{(t)}, b_{j,50}^{(t)}, b_{j,90}^{(t)})$

Table 4. Firm states in each information setting.

	$\mathcal{A}_{f,\texttt{benchmark choice}}$	$\mathcal{A}_{f,\texttt{acceptance threshold choice}}$
`Neither`	{`independent`, `shared`}	$\{b_{j,10}^{(t)}, b_{j,50}^{(t)}, b_{j,90}^{(t)}\}$
`Firm-posted`	{`independent`, `shared`}	$\{b_{j,10}^{(t)}, b_{j,50}^{(t)}, b_{j,90}^{(t)}\}$
`Predicted`	{`independent`, `shared`, `predicted`}	$\{b_{j,10}^{(t)}, b_{j,50}^{(t)}, b_{j,90}^{(t)}\}$
`Both`	{`independent`, `shared`, `predicted`}	$\{b_{j,10}^{(t)}, b_{j,50}^{(t)}, b_{j,90}^{(t)}\}$

Table 5. Firm actions in each information setting.

Rewards. First, $\mathcal{R}_w$ is defined as follows for each worker i for $S_s \in \mathcal{S}_{w,\texttt{sharing}}$, $S_f \in \mathcal{S}_{w,\texttt{firm choice}}$, and $S_o \in \mathcal{S}_{w,\texttt{offer choice}}$.

$$R_w(S_s, A^{(t)}) = \begin{cases} 0 & s_i^{(t)} = \texttt{no share} \\ \epsilon_s & s_i^{(t)} = \texttt{share} \end{cases} \qquad R_w(S_f, A^{(t)}) = \begin{cases} b_{f_i^{(t)},90}^{(t)} - w_i^{(t)} & \text{range posted, } \texttt{Firm-posted} \text{ or } \texttt{Both} \\ \hat{u}^{(t)} - w_i^{(t)} & \text{range posted, } \texttt{Predicted} \\ 0 & \text{range not posted} \end{cases}$$

$$R_w(S_o, A^{(t)}) = \begin{cases} o_i^{(t)} - w_i^{(t)} & \text{negotiation succeeded} \\ -1 & \text{negotiation failed} \\ 0 & o_i^{(t)} = \texttt{no offer} \end{cases}$$

We assume agents prefer to take actions that increase the information they learn about the market, so $\epsilon_s > 0$ is the reward that workers get for sharing some information to gain more information, and we use $\epsilon_s = 10^{-2}$ throughout.

Next, $R_f : \mathcal{S}_f \times \mathcal{A}$ is defined as follows for each firm j for $S_b \in \mathcal{S}_{f,\texttt{benchmark choice}}$ and $S_a \in \mathcal{S}_{f,\texttt{acceptance threshold choice}}$. Let $O_j^{(t)}$ be the set of offers firm j receives at time t.

$$R_f(S_b, A^{(t)}) = \begin{cases} \epsilon_s & B_j^{(t)} = S^{(t)} \\ 0 & B_j^{(t)} \neq S^{(t)} \end{cases} \qquad R_f(S_a, A^{(t)}) = \begin{cases} \frac{1}{|O_j^{(t)}|} \left(\sum_{o \in O_j^{(t)}} \left[\mathbf{1}\{o \leq a_j^{(t)}\} \cdot (1 - o) - \mathbf{1}\{o > a_j^{(t)}\} \cdot 1 \right] \right) & |O_j^{(t)}| > 0 \\ 0 & |O_j^{(t)}| = 0 \end{cases}$$

In the firm's negotiation reward, each accepted offer results in the firm gaining the worker's MRPL of 1 and losing the wage they pay to that worker, o, each rejected offer is internalized by the firm as losing the potential MRPL of 1 from that worker, and the sum is normalized by number of offers received.

A.3 Labor Market Simulation Pseudocode

Our simulation pseudocode is provided in Algorithm 1 for one worker type only and similar logic is used for two types of workers. The subroutines ActionChoiceFirm and ActionChoiceWorker depend on the current value $\epsilon^{(t)}$ and perform ϵ-greedy action exploration, i.e., with probability $\epsilon^{(t)}$, a (valid) action is chosen uniformly at random and with probability $1 - \epsilon^{(t)}$, a (valid) action is chosen which maximizes the current Q-value estimate for the firm or worker, i.e., $Q^{(t)}_{f,j}$ or $Q^{(t)}_{w,i}$, respectively. Note that to ensure firms only set acceptance thresholds and workers only make offers from the set $\mathcal{W}$, we assume any percentile calculations (e.g., during firm salary benchmark and `salary prediction tool updates`) map back to the closest (in terms of ℓ-1 distance, ties broken by choosing the smaller value) value in $\mathcal{W}$ to use as the percentile value, i.e., $\arg\min\{\|w - \eta_c(W)\|_1 : w \in \mathcal{W}\}$. Additionally, if at any time a data pool used to produce wage percentile estimates is empty (e.g., for the salary prediction tool or any of the firm benchmarks), then the last available estimate is used, if available, or the flag -1 to indicate not enough data has been collected.

Algorithm 1 Labor Market Simulation via ϵ-Greedy Q-learning

Require: $\alpha, \delta, \beta, T, N_w, N_f, \mathcal{W}, p$
  $Q^{(0)}_{w,i} \leftarrow \mathbf{0} \forall i \in [N_w], Q^{(0)}_{f,j} \leftarrow \mathbf{0} \forall j \in [N_f]$
  $w^{(0)}_i$ drawn from distribution p over $\mathcal{W}$ $\forall i \in [N_w]$
  $D^{(0)}$ initialized with 50% of workers randomly
  $S^{(0)}$ initialized with 50% of firms randomly
  for $t = 0 \ldots T - 1$ **do**
    $\epsilon^{(t)} = e^{-\beta \cdot t}$
    Update all benchmarks in $\mathcal{B}^{(t)}_j$ $\forall j \in [N_f]$
    $B^{(t)}_j \leftarrow$ ActionChoiceFirm($\epsilon^{(t)}$) and $Q^{(t)}_{f,j}$ update via equation 1 $\forall j \in [N_f]$ ▷ Firm Benchmark Choice
    if Ranges posted **then**
      $\{d^{(t)}_j : j \in [N_f]\}$ provided to workers
    end if
    if `salary prediction tool` **then**
      $s^{(t)}_i \leftarrow$ ActionChoiceWorker($\epsilon^{(t)}$) and $Q^{(t)}_{w,i}$ update via equation 1 $\forall i \in [N_w]$ ▷ Worker Sharing Choice
      Update `salary prediction tool` with $D^{(t)}$
    end if
    $a^{(t)}_j \leftarrow$ ActionChoiceFirm($\epsilon^{(t)}$) $\forall j \in [N_f]$ ▷ Firm Acceptance Threshold Choice
    $f^{(t)}_i \leftarrow$ ActionChoiceWorker($\epsilon^{(t)}$) and $Q^{(t)}_{w,i}$ update via equation 1 $\forall i \in [N_w]$ ▷ Worker Firm Choice
    $o^{(t)}_i \leftarrow$ ActionChoiceWorker($\epsilon^{(t)}$) $\forall i \in [N_w]$ ▷ Worker Offer Choice
    WageNegotiation($\{a^{(t)}_j : j \in [N_f]\}, \{(f^{(t)}_i, o^{(t)}_i) : i \in [N_w]\}$)
    $Q^{(t)}_{w,i}$ update via equation 1 and $Q^{(t)}_{f,j}$ update via equation 1 $\forall j \in [N_f]$
  end for

B Additional Simulation Results

B.1 Experiment A: W_F Violin Plots

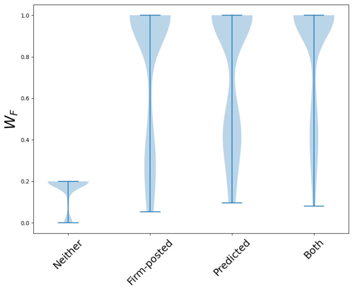

(a) $\tau_H = 0.75$, $r = 0.5$

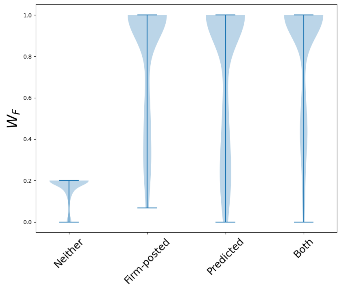

(b) $\tau_H = 0.5$, $r = 0.5$.

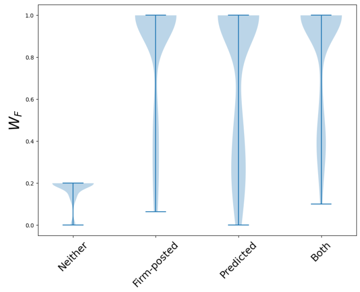

(c) $\tau_H = 0.25$, $r = 0.5$

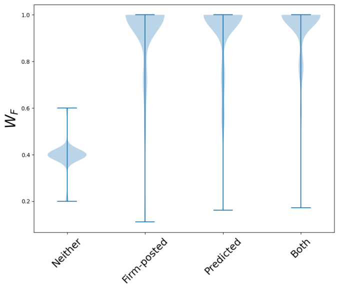

(d) $\tau_H = 0.75$, $r = 0.25$

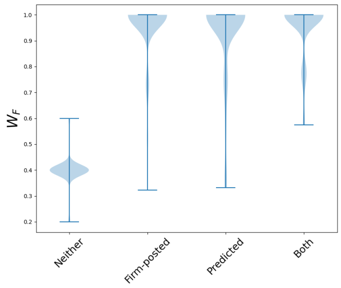

(e) $\tau_H = 0.5$, $r = 0.25$

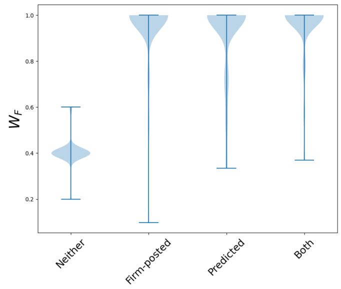

(f) $\tau_H = 0.25$, $r = 0.25$

Fig. 5. Experiment A simulation results: W_F empirical distribution across the information settings and all market conditions.

B.2 Experiment A: Benchmark Comparison Bar Charts

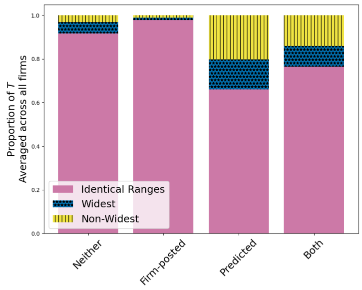

(a) $\tau_H = 0.75, r = 0.5$

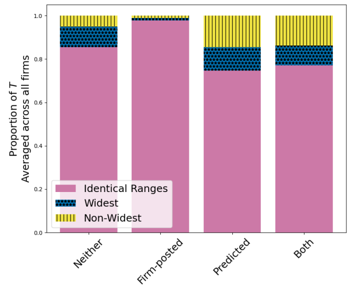

(b) $\tau_H = 0.5, r = 0.5$

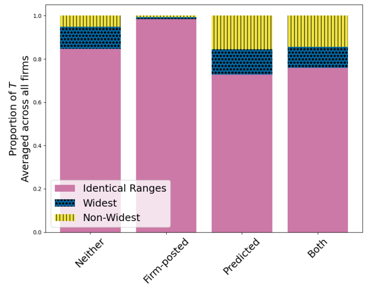

(c) $\tau_H = 0.25, r = 0.5$

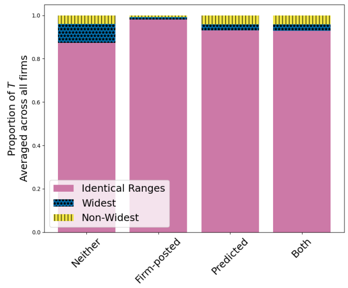

(d) $\tau_H = 0.75, r = 0.25$

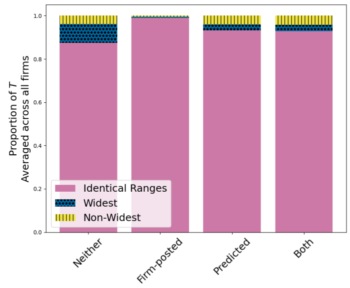

(e) $\tau_H = 0.5, r = 0.25$

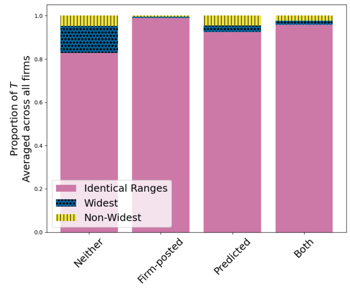

(f) $\tau_H = 0.25, r = 0.25$

Fig. 6. Experiment A simulation results: Firm benchmark choice analysis of a single simulation run across information settings and all market conditions.

B.3 Experiment A: Acceptance Threshold Gap Empirical Distributions

In Experiment A, we track the average gap between the firms' final worker type acceptance thresholds,

$$AT_G = \frac{1}{N_f} \sum_{j=1}^{N_f} \left(a_{j,H}^{(T)} - a_{j,L}^{(T)} \right).$$

The empirical distribution across $N = 120$ runs is given as a violin plot in Figure 7, and the mean, variance, and skewness of the empirical distribution is given in Table 6.

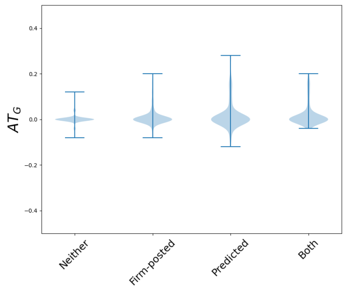

(a) $\tau_H = 0.5$, $r = 0.5$

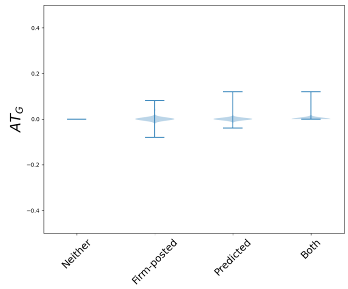

(b) $\tau_H = 0.5$, $r = 0.25$

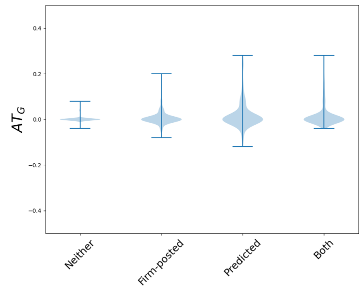

(c) $\tau_H = 0.75$, $r = 0.5$

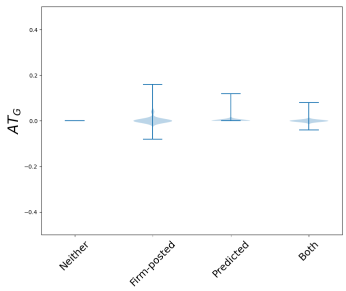

(d) $\tau_H = 0.75$, $r = 0.25$

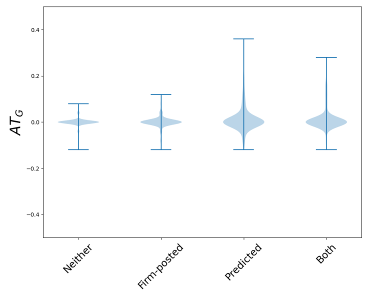

(e) $\tau_H = 0.25$, $r = 0.5$

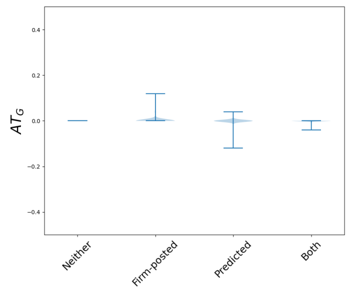

(f) $\tau_H = 0.25$, $r = 0.25$

Fig. 7. Experiment A simulation results: AT_G empirical distribution across all market conditions.

	Neither AT_G Mean (Variance) *Skewness*	Firm-posted AT_G Mean (Variance) *Skewness*	Predicted AT_G Mean (Variance) *Skewness*	Both AT_G Mean (Variance) *Skewness*
$\tau_H = 0.25$ $r = 0.25$	**0.0000 (0.0000)** *N/A*	0.0043 (0.0004) *4.8282*	-0.0003 (0.0001) *-7.4325*	-0.0007 (**0.0000**) *-7.5510*
$\tau_H = 0.25$ $r = 0.5$	**0.0000 (0.0003)** *-1.6771*	-0.0007 (0.0008) *-0.6841*	0.0143 (0.0033) *2.6300*	0.0143 (0.0026) *2.4317*
$\tau_H = 0.5$ $r = 0.25$	**0.0000 (0.0000)** *N/A*	0.0010 (0.0003) *0.6276*	0.0020 (0.0002) *5.9220*	0.0027 (0.0003) *6.6413*
$\tau_H = 0.5$ $r = 0.5$	**0.0007 (0.0003)** *1.7553*	0.0103 (0.0015) *2.500*	0.0187 (0.0038) *2.0777*	0.0207 (0.0025) *2.2140*
$\tau_H = 0.75$ $r = 0.25$	**0.0000 (0.0000)** *N/A*	0.0037 (0.0006) *2.8728*	0.0023 (0.0002) *7.0680*	0.0017 (0.0002) *5.1222*
$\tau_H = 0.75$ $r = 0.5$	**0.0013 (0.0001)** *3.1842*	0.0083 (0.0012) *2.9070*	0.0190 (0.0036) *2.140*	0.0167 (0.0022) *3.0428*

Table 6. Mean, variance, and skewness of the empirical distribution of AT_G across market conditions and information settings, rounded to 4 decimal places. The lowest (absolute) mean and lowest variance(s) are **bold** in each row. N/A indicates no skewness calculated due to all AT_G observations being the same number.

B.4 Experiment A: Additional Results

The figures and tables here document the results for market conditions and additional statistics not shown in the main paper.

	Neither W_G Mean (Variance) *Skewness*	Firm-posted W_G Mean (Variance) *Skewness*	Predicted W_G Mean (Variance) *Skewness*	Both W_G Mean (Variance) *Skewness*
$\tau_H = 0.25$ $r = 0.25$	0.0002 (**0.0000**) *-2.3978*	0.0002 (**0.0000**) *-0.8062*	0.0056 (0.0004) *2.7521*	**0.0000** (**0.0000**) *-0.5292*
$\tau_H = 0.25$ $r = 0.5$	0.0100 (0.0010) *0.4367*	**0.0016** (**0.0004**) *0.7314*	0.0058 (0.0017) *2.3968*	0.0057 (0.0012) *0.4981*
$\tau_H = 0.5$ $r = 0.25$	**0.0001** (**0.0000**) *1.4968*	-0.0014 (**0.0000**) *-3.3019*	0.0057 (0.0003) *3.4914*	0.0022 (0.0001) *3.8472*
$\tau_H = 0.5$ $r = 0.5$	0.0146 (0.0016) *1.7026*	**0.0037** (**0.0002**) *1.8777*	0.0279 (0.0046) *2.0701*	0.0195 (0.0019) *1.9430*
$\tau_H = 0.75$ $r = 0.25$	**0.0002** (**0.0000**) *3.4926*	**-0.0002** (0.0002) *0.3097*	0.0085 (0.0007) *3.2902*	0.0023 (0.0002) *4.5805*
$\tau_H = 0.75$ $r = 0.5$	0.0058 (0.0006) *4.5394*	**0.0053** (**0.0003**) *1.5235*	0.0596 (0.0096) *1.5514*	0.0356 (0.0050) *1.8546*

Table 7. Mean, variance, and skewness of the empirical distribution of W_G across market conditions and information settings, rounded to 4 decimal places. The lowest (absolute) mean(s) and lowest variance(s) are **bold** in each row.

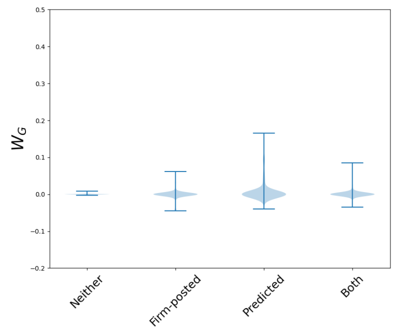

(a) $\tau_H = 0.75$, $r = 0.25$

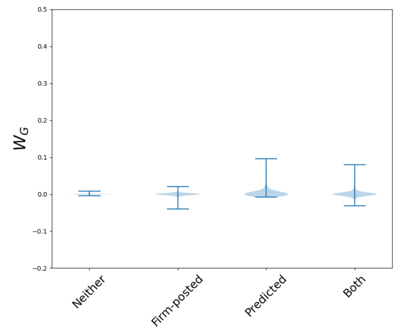

(b) $\tau_H = 0.5$, $r = 0.25$

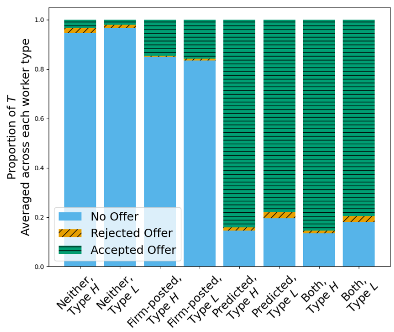

(c) $\tau_H = 0.75$, $r = 0.25$

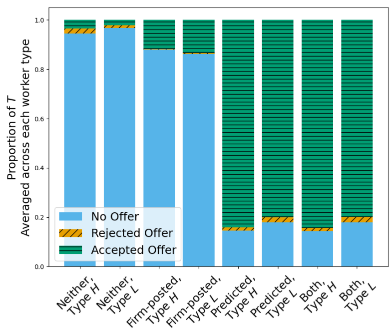

(d) $\tau_H = 0.5$, $r = 0.25$

Fig. 8. Experiment A simulation results: $\tau_H = 0.75$, $r = 0.25$ vs. $\tau_H = 0.5$, $r = 0.25$.

C Robustness Results

To evaluate the robustness of our results, we varied the number of workers (N_w) and firms (N_f) in our market. For each robustness test, we kept $\alpha = 0.3$, $\delta = 0.9$, $k = 5$-discretized wages, and $T = 20,000$. We tested the market conditions $\tau_H = 0.5$ and $r = 0.5$ in order to evaluate whether our claims about W_F and W_G between information settings holds (qualitatively) under these varied market sizes. For each experiment, we ran 2 simulation runs for each of the eight initial wage distributions for a total of $N = 16$ runs per experiment. For all of the experiments, we find that our qualitative claims still hold: 1) Workers negotiate more often in `Predicted` and `Both` than in `Neither` or `Firm-posted`, 2) Workers have the largest wages in `Firm-posted` or `Both`, 3) Firms see more diverse benchmarks in `Predicted` and `Both`, and, finally, 4) workers are more likely to have larger wage gaps in `Predicted` or `Both` than in `Firm-posted`.

C.1 More Workers

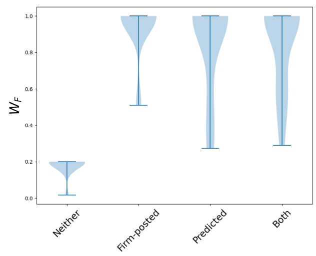

(a) W_F empirical distribution. $\tau_H = 0.5$, $r = 0.5$

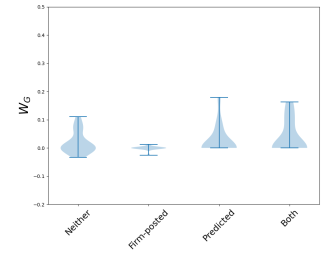

(b) W_G empirical distribution. $\tau_H = 0.5$, $r = 0.5$.

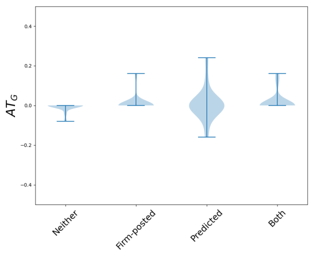

(c) AT_G empirical distribution. $\tau_H = 0.5$, $r = 0.5$

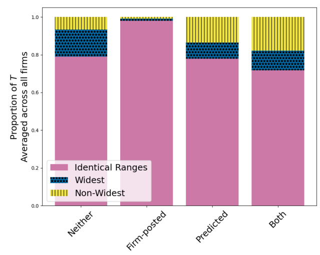

(d) Average firm benchmark choice comparison. $\tau_H = 0.5$, $r = 0.5$

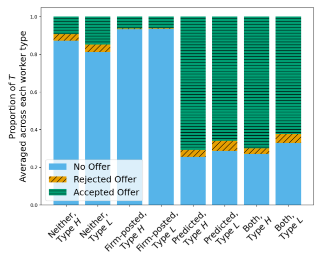

(e) Average worker offer choice comparison. $\tau_H = 0.5$, $r = 0.5$

Fig. 9. Robustness result: $N_w = 250$.

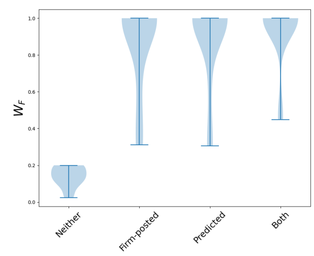

(a) W_F empirical distribution. $\tau_H = 0.5$, $r = 0.5$

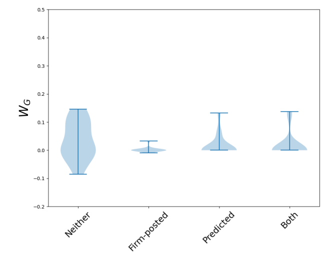

(b) W_G empirical distribution. $\tau_H = 0.5$, $r = 0.5$.

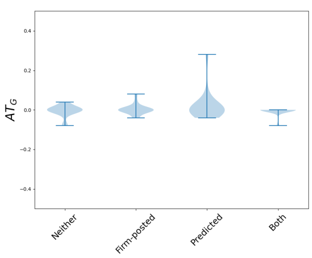

(c) AT_G empirical distribution. $\tau_H = 0.5$, $r = 0.5$

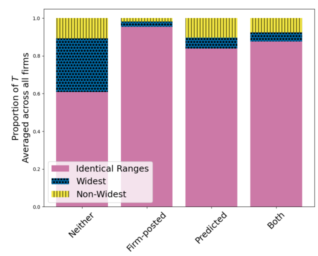

(d) Average firm benchmark choice comparison. $\tau_H = 0.5$, $r = 0.5$

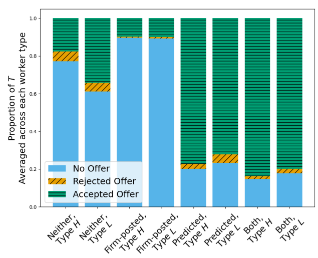

(e) Average worker offer choice comparison. $\tau_H = 0.5$, $r = 0.5$

Fig. 10. Robustness result: $N_w = 500$.

C.2 Fewer Firms

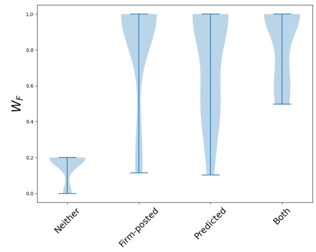

(a) W_F empirical distribution. $\tau_H = 0.5$, $r = 0.5$

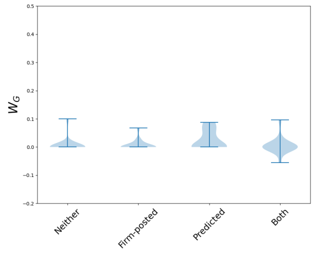

(b) W_G empirical distribution. $\tau_H = 0.5$, $r = 0.5$.

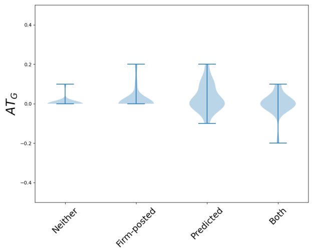

(c) AT_G empirical distribution. $\tau_H = 0.5$, $r = 0.5$

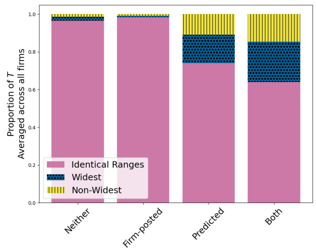

(d) Average firm benchmark choice comparison. $\tau_H = 0.5$, $r = 0.5$

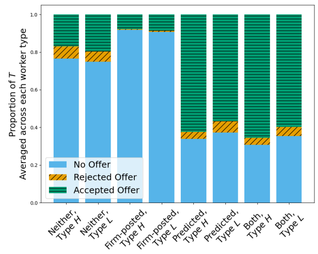

(e) Average worker offer choice comparison. $\tau_H = 0.5$, $r = 0.5$

Fig. 11. Robustness result: $N_f = 2$.

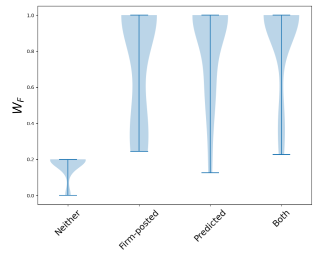

(a) W_F empirical distribution. $\tau_H = 0.5$, $r = 0.5$

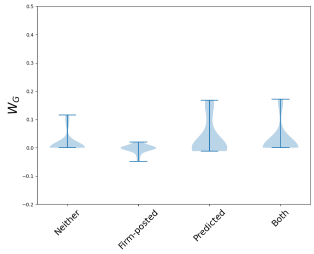

(b) W_G empirical distribution. $\tau_H = 0.5$, $r = 0.5$.

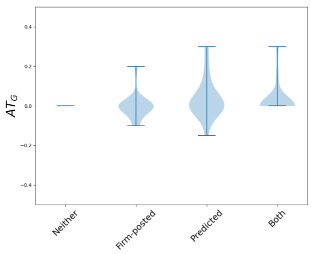

(c) AT_G empirical distribution. $\tau_H = 0.5$, $r = 0.5$

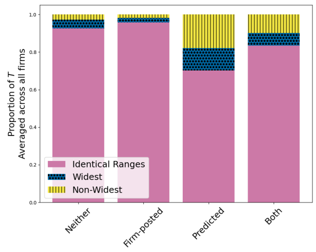

(d) Average firm benchmark choice comparison. $\tau_H = 0.5$, $r = 0.5$

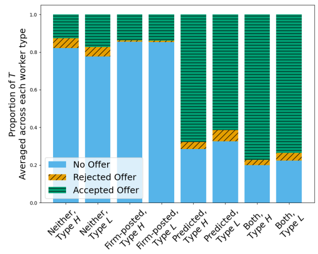

(e) Average worker offer choice comparison. $\tau_H = 0.5$, $r = 0.5$

Fig. 12. Robustness result: $N_f = 4$.