Equilibrium Effects of Pay Transparency in a Simple Labor Market*

Zoë B Cullen† and Bobak Pakzad-Hurson‡

April 2020

The public discourse around pay transparency has focused on the direct effect: how workers seek to rectify newly disclosed pay inequities through renegotiations. The question of how wage-setting, bargaining, and hiring practices change with higher transparency has received less attention. To study these outcomes, we combine a dynamic wage-bargaining model with data from online labor markets for low-skill, temporary jobs. We exploit naturally occurring variation in pay transparency as well as experimentally-induced variation. Wages are more equal, but lower under transparency. An increase in transparency raises the hiring rate when workers have sufficient bargaining power. Employer profits rise with transparency, increasing 27% in a field experiment. A key insight is that increasing transparency decreases the bargaining power of workers, as employers credibly refuse to pay high wages to any one worker to avoid costly renegotiations with others. We discuss implications for the gender wage gap and employers' endogenous transparency choices.

Keywords: Pay Transparency, Negotiation, Online Labor Market, Privacy, Wage Gap
JEL Classification Codes: C78, D82, D83, J31, L81, M52

*We are very grateful for guidance from Ran Abramitzky, Susan Athey, Nick Bloom, Matt Jackson, Fuhito Kojima, Ed Lazear, Luigi Pistaferri, Al Roth, and Gavin Wright. We are indebted to James Flynn for technical support. We also thank Mohammad Akbarpour, Lint Barrage, Jose Maria Barrero, Doug Bernheim, Ben Brooks, Eric Budish, Gabriel Carrol, Arun Chandrasekhar, Kalyan Chatterjee, Isa Chaves, Ken Chay, Bo Cowgill, Piotr Dworczak, Jack Fanning, Chiara Farronato, Laura Gee, Matt Gentzkow, Avi Goldfarb, Bob Hall, Gregor Jarosch, Scott Kominers, Maciej Kotowski, Vijay Krishna, Jon Levin, Shengwu Li, Erik Madsen, Davide Malacrino, Alejandro Martinez, Paul Milgrom, Muriel Niederle, Kareen Rozen, Ilya Segal, Isaac Sorkin, Jesse Shapiro, Chris Stanton, Bryce Millet Steinberg, Takuo Sugaya, Neil Thakral, Catherine Tucker, Christopher Snyder, Emmanuel Vespa, and Alistair Wilson and seminar attendees for helpful comments and suggestions. We appreciate TaskRabbit for providing data access. This research was supported by the Center for Comparative Studies in Race and Ethnicity and the B.F. Haley and E.S. Shaw Fellowship through SIEPR.

†Cullen: Harvard Business School, Rock Center, Soldiers Field, Boston, MA 02163, zcullen@hbs.edu
‡Pakzad-Hurson: Brown University, 64 Waterman Street, Providence RI 02912, bph@brown.edu

I. Introduction

Most pay transparency initiatives are based on the narrative that transparency gives workers more bargaining power. Pay transparency laws aim to increase workers' knowledge of the pay of their peers to ensure "victims of pay discrimination can effectively challenge unequal pay," equipping them for successful negotiations by revealing their employer's willingness to pay for labor (Phillips, 2009).[1] But the use of salary disclosure to remedy unequal pay for equal work is only half of the story; when salary transparency is anticipated by the employer and employees, optimal wage-setting, bargaining, and hiring practices also adjust. Despite a lack of evidence on the indirect effects of pay transparency,[2] thirteen U.S. states and ten EU countries have passed laws to increase pay transparency.[3]

Our paper studies both the direct and indirect impacts of pay transparency policies on wage negotiations, hiring, and profits in the context of on-demand labor markets. The short duration of on-demand jobs allows us to observe the effects of transparency propagating through the market and allow detailed visibility of wage negotiations between thousands of employers and workers at a level of detail that goes unrecorded in other settings. We combine a dynamic wage-bargaining model with administrative data from TaskRabbit for Business, an online platform where businesses source employees for in-person contract work, and a field experiment involving online labor market participants. We find that increasing transparency decreases workers' de facto bargaining power. Because of this, under higher transparency wages are more equal but lower, employer profits are higher, and hiring increases provided that workers start with sufficient bargaining power.

To introduce the mechanisms we study in this paper, we present two scenarios. First, suppose that a worker learns that a colleague with the same job is earning significantly more than she is. She reasons that her employer must be willing to pay her a higher wage as well. She confronts her boss and successfully negotiates a raise to bring her wage on par with her coworker's, perhaps by invoking Equal Pay legislation.[4] This direct effect of wage disclosure has both increased the worker's pay and reduced inequality. But this scenario

[1] Employers' private information can affect bargaining and lead to high levels of inequality in imperfectly competitive markets (Manning, 2005).

[2] In a recent paper, Mas (2017, page 1718) states, "More work could also be done to investigate other effects of pay disclosure...and whether transparency changes the relative bargaining of workers and employers in wage setting."

[3] The policies in the U.S. range from punishments for employers that prevent employees from discussing salaries (Siniscalco et al., 2017) to requirements that employers inform prospective employees of the range of salaries currently paid to workers (Pender, 2017). EU policies range from mandating full disclosure by companies over a threshold number of employees, like in Germany, to publication of salary statistics such as the mean, median and gap between men and women, as in the case of the U.K. and Denmark (European Commission; International Labour Organization).

[4] Caldwell and Harmon (2019) present empirical evidence from Denmark that workers renegotiate their wages on the job after learning new wage information.

does not account for equilibrium effects in an environment of pay transparency.

Now consider the initial wage negotiations of the same worker at a firm with full pay transparency. Even before agreeing to an initial wage, she observes the salary of the highest-paid worker. As before, the worker negotiates to no less than this observed wage. But she is not able to negotiate the wage any higher– her boss credibly states that he cannot afford to pay her more because her peers would immediately demand commensurate raises. If wage information were not public, the worker might be skeptical of her employer's claim and might bargain more aggressively. But in this case, the worker grasps the ramifications of asking for more than the current maximum wage earned by others. Full transparency leads to an unintended side effect: if workers all get the same wage and cannot negotiate this wage upward, then the firm gets the power to set the wage. In order to maximize its profit, the firm acts like a monopsonist and sets a relatively low wage. Thus transparency, rather than benefiting the worker, increases the de facto bargaining power of the employer, becoming the enforcement mechanism for a low wage.

Transparency, in this way, resembles best-price guarantees (also known as most-favored-customer agreements). Best-price models argue that agreements to rebate existing customers if prices fall in the future allows the firm to maintain higher prices (Butz, 1990; Cooper and Fries, 1991). The firm can commit to refusing marginal buyers because of the cost of rebates to previous consumers,[5] and empirical evidence supports these findings (Scott Morton, 1997a,b). Our study contains this effect, but an important point of departure is that our setting is one of (two-sided) incomplete information. Both sides of the market alter their bargaining strategies in response to transparency, and the interaction of these responses drives our theoretical predictions.

We present a baseline model of continuous-time wage negotiation in the presence of transparency, in which a continuum of workers individually bargain for wages with a firm. We extend this model in many ways, including a setting in which there are many firms and workers search among them for jobs. As our key findings are preserved in these extensions, we investigate our simple model to shed light on the forces at play.

Each worker has a heterogeneous outside option and the firm has a value for labor, both of which are private information. Bargaining is similar to double auction models, first studied by Chatterjee and Samuelson (1983): each worker specifies a minimum wage she is willing to accept, and the firm sets a maximum wage it is willing to accept. If the worker's minimum acceptable wage is below the firm's maximum acceptable wage, she is hired and the flow wage she receives is a convex combination of these two values.[6] Otherwise, the worker is

[5]See Bagwell and Wolinsky (2002) for a series of pricing models that demonstrate how the future observability of durable decisions can lead to increased commitment power.

[6]As argued by Larsen and Zhang (2018), double auctions serve as useful empirical approximations to real-world bargaining in settings with private information. For example, Larsen (2020) empirically estimates

permanently unmatched with the firm and receives her outside option. The convex weight on the firm's maximum wage measures its bargaining power. The higher the weight, the larger its role in determining the wage resulting from a negotiation.

Over time, workers stochastically learn the wages of their peers. An independent arrival process reveals the entire profile of wages at the firm to a particular worker, and the arrival rate characterizes the level of transparency: with higher levels of transparency, wage information arrives more quickly. Under full transparency, workers learn the wages of their peers immediately; under full privacy, they never do. Workers renegotiate their wage when information arrives by specifying a new minimum acceptable wage.

Transparency causes an information externality, as one worker's wage can affect (re)negotiations of others. In equilibrium, when a worker receives peer wage information, she will renegotiate her wage to match the highest wage she observes. This alters the de facto bargaining power through two equilibrium effects: a *demand effect* and a *supply effect*. As transparency rises, the firm's maximum willingness to pay for labor falls because information about one worker's pay raise spreads more quickly to others, who use that information to renegotiate (demand effect). At the same time, workers make lower initial wage offers to increase their chances of getting hired (supply effect). Because workers expect to quickly learn the wages of others and renegotiate with higher transparency, they are less concerned with securing a high initial wage.

Dynamic games with incomplete information frequently contain analogues of one of these effects, but not, to our knowledge, both. In the well-known chain store game, Kreps and Wilson (1982) and Milgrom and Roberts (1982) show that costly, predatory behavior against early competitors may be optimal in order to create a reputation favorable for later negotiations (demand effect).[7] Kuhn and Gu (1998, 1999) show that unions optimally delay making contract offers to employers so that they can freeride on information gathered from the negotiations of others (supply effect).[8] Our setting includes both of these supply and demand effects, which cause simultaneous adjustments of bargaining strategies by workers and the firm in response to changes in transparency. We prove several novel equilibrium results.

We show that increasing transparency has the same effect as increasing firm bargaining power. Formally, we show that the set of equilibria of this game under high transparency and low firm bargaining power is identical to the set of equilibria with lower transparency and higher firm bargaining power. Regardless of a firm's nominal bargaining power, full trans-

a double auction model using data from post-auction bargaining over used cars. Our empirical setting provides a rare opportunity to study details of bargaining for wages, rather that goods, in settings with private information.

[7] Fudenberg and Levine (1992) show a weaker, yet similar, effect in settings with imperfect observations, which is similar to the response to partial transparency in our model.

[8] Gu and Kuhn (1998) show that an analogue of the demand effect may be present in this setting, although they do not consider both effects simultaneously.

parency grants the firm full de facto bargaining power, as the firm commits to, and workers immediately observe, a maximum wage; and when a firm has full nominal bargaining power, all workers receive the same take-it-or-leave-it offer, implying that all workers immediately know the wages of their peers, whether or not they directly observe these wages through another channel.

Earnings are more equal with higher transparency, as workers renegotiate to a common wage. Because workers have heterogeneous outside options, this implies worker surplus–the difference between earnings and outside option–is more dispersed with higher transparency.

Pay transparency also shifts surplus from the workers to the firm, leading to lower wages and higher profit. Under full transparency, the firm will pick a maximum wage equal to the wage it would have selected as a monopsonistic firm posting a wage. Each worker will either accept the transparent maximum wage, or will consume her outside option.[9] Because of this, full transparency maximizes firm expected profits (Williams, 1987) and minimizes worker surplus and wages.

The amount of firm bargaining power has a non-monotonic effect on employment because employment is maximized when bargaining power is shared between the workers and the firm. When bargaining power is highly skewed, either the worker acts like a monopolist, making high wage demands that are often rejected; or the firm acts like a monopsonist, committing to low wages that deter high-outside-option workers from considering work at the firm. Granting either the firm or the workers all of the bargaining power minimizes expected hiring. Therefore increasing transparency, which effectively increases firm bargaining power, increases the expected hiring rate only if workers initially have sufficient bargaining power.[10]

The implications of increasing transparency are not the same for every worker and every firm type. The results described thus far are in expectation, given the distribution of firm and worker types. Workers with low outside options stand to gain relatively more from an increase in transparency because they benefit more than others from renegotiation. A firm with a low value for labor benefits relatively more from transparency. The reason is that, under pay privacy, workers demand a larger premium over their outside options (supply effect), and these demands exceed the value for labor of low firm types. Hence, low-value firm types are unable to hire when pay is private. When wages are transparent, the supply

[9]Because of this equilibrium equivalence between full transparency and posted wages, our theory unifies previous results from a variety of models (and our empirical evidence corroborates these findings). Michelacci and Suarez (2006) show that bargaining leads to more dispersed wages than posted wages (our Theorem 1); Ellingsen and Rosén (2003) find wage posting is more effective than bargaining when reservation wages are low (an implication of our Theorem 3); Brenzel et al. (2014) suggest that bargaining may lead to higher average worker wages than wage posting (our Theorem 2).

[10]The potential positive effect of pay transparency on hiring is perhaps surprising given results of other models of bargaining with private information (Bergemann and Hörner, 2018; Brancaccio et al., 2020; Hörner and Vieille, 2009; Kaya and Liu, 2015). All find that transparency decreases the number of (or prevents) transactions.

effect causes workers to reduce their initial demands and some will ask for less than the firm's value of labor, leading to hiring and positive profit. A high-value firm is harmed by too much transparency; due to its high value, it optimally sets a high maximum wage to hire many workers. Increased transparency means it must quickly give large raises to workers who initially demand a low wage.

Because a low-value firm prefers a higher level of transparency than a high-value firm, endogenously chosen transparency signals to workers the value of the firm. When the firm chooses a level of transparency after observing its value for labor, there is a unique equilibrium in which all firm types pool on full transparency. In this simple environment, the choice of transparency is a verifiable signal of the firm's value.[11] Therefore, the market unravels (Milgrom, 1981) toward the unique equilibrium.

Back-end data from TaskRabbit for Business allows us to study and verify these equilibrium effects in an environment uncomplicated by non-pecuniary benefits or career concerns. Most jobs on the platform are low-skill, standardized tasks, and jobs are typically carried out in person. We observe all transactions on the platform between 2010 and 2014, as well as job postings, worker bids, on-the-job bonuses, employer ratings of workers, and worker and employer demographics. Employers using the platform range from older and more established firms (approximately 10% of all Fortune 500 companies are in our sample) to household employers (which comprise roughly one quarter of our sample).

In many jobs, the worker has full nominal bargaining power and sends an initial wage offer to the employer. Wage offers are submitted privately online, and are observable only by the employer. Each accepted wage offer serves as a price floor, but the employer can raise the worker's wage before the contract is completed, allowing for the possibility of wage renegotiations on the job.

We exploit variation in the private submission of bids, and the downstream opportunity to communicate with co-workers about pay on the job. Workers participating in these markets expect to discuss their wages while on the job in 40% of circumstances, according to our survey of 5,000 workers. Communication channels on the job can vary substantially, however. For example it may be necessary for two workers to distribute marketing materials in the same location or different locations depending on where clients are during a particular shift. In our preferred specification, we isolate variation in the ability of particular workers to

[11]The verifiability of information spread by transparency, and the risk of negotiations breaking down if workers' initial offers are too high are key to this finding. Menzio (2007) studies a directed search model where firms can send non-binding messages to workers, and shows that there can exist a partial separating equilibrium which allows workers to direct their search. In Menzio's model, workers cannot extract the full surplus from firms, because information is non-verifiable and firms partially pool. Since high initial worker offers do not lead to an immediate breakdown of negotiations in equilibrium, the relatively lower firm types that signal high value are still able to hire workers. Our result extends with these key features even in a search environment, and we discuss support data for our theory from TaskRabbit for Business.

communicate about pay by comparing the same job description posted by the same employer, but the co-location of these workers and the length of time they overlap differs.

Our model predicts, first, workers should receive the same final wage as a higher-earning co-located worker *if* they renegotiate after the job begins. Second, the act of renegotiating should be uncorrelated with the size of the initial wage gap between coworkers. Note that this latter prediction helps differentiate our model from models in which fairness concerns are the primary driver of pay equalization, which we discuss later.

Our data verify these claims. We find that the bids workers submit privately to the platform before the job begins will causally affect the renegotiation outcomes of their coworkers if and only if workers are co-located. When workers are co-located and employers adjust wages above the initial wages, final wages are nearly completely equal. The degree of inequality between workers' initial wages does not predict whether a worker renegotiates at all.

We consider, and find little evidence for, alternative mechanisms for the wage equalization patterns we observe, such as employer preference for equal wages and productivity spillovers between workers equalizing observable output. We also do not find evidence that employers pre-emptively equalize wages to prevent worker morale and effort from falling in the face of known wage inequality, as we later discuss.

TaskRabbit for Business data also allow us to study the endogenous choice of employers to select full transparency. Employers can choose to be fully transparent by publicly posting a wage along with the job description, or by mentioning a wage in the text of a job description. Our model predicts first that employers with a low value for labor benefit most from choosing transparency; second, that wages are lower in transparent jobs; and third, that the market unravels over time toward greater transparency.

Our findings bear out all three predictions. We show that employers with lower household income[12]–a proxy for employers with a low value for labor–are the most likely to choose wage transparency. We show that employers who select full transparency transact at lower wages. The very same workers bid hourly wages 7.8% lower for work in the same job category when the job description mentions expected wages up front. We also observe the market unraveling towards higher levels of pay transparency across all employers. We observe a striking linear progression toward the use of transparent, posted wages month over month. For every month on the platform, the fraction of jobs using a transparent posted price increases by 1%. This trend is not explained by the changing composition or number of jobs on the platform.

We run a field experiment to directly measure the division of surplus between the employer and worker. Across all treatments, we hire 365 managers and 964 workers from an online labor market, Amazon Mechanical Turk, who are tasked with negotiating wages for, and carrying out, transcription services that take half a work day to complete on average. We

[12]Household income data on clients is estimated by SalesForce for marketing purposes.

elicit outside options from workers and assign a budget to managers, who are the residual claimants of that budget after paying wages for transcription. We randomize transparency by restricting wage negotiations to either a common chat room or a private chat room, where the only difference is whether coworkers can observe the negotiations with other workers. We document all interactions in these chat rooms, and we place no restrictions on the ways in which workers and managers can bargain.

Our field experiment results corroborate our theoretical findings and TaskRabbit for Business analyses: pay is equalized 100% of the time when workers negotiate in a common chat room, compared to 60% of the time in private chat rooms. In the common chat room, manager profits are 27% higher and wages 7.4% lower, and hiring (the share of workers who reach an agreement with the employer) rises by 10%. Worker surplus is lower and more dispersed, as predicted by our theory. Productivity levels, conditional on hire, are similar across the two groups.

Our paper builds on a rapidly-growing body of empirical literature on the effects of pay transparency. One strand of the literature focuses on individuals' behavioral responses to learning others' wages (Bracha et al., 2015; Breza et al., 2018; Card et al., 2012; Cohn et al., 2014; Dube et al., 2019; Perez-Truglia, 2019) . These papers generally find that workers decrease effort upon learning that others earn more, before wages are adjusted. Another strand studies equilibrium responses to pay transparency mandates (Baker et al., 2019; Bennedsen et al., 2019; Kim, 2015; Mas, 2017). Our evidence is broadly consistent with these studies on the effects of transparency and its tendency to compress wages and slow wage growth. Mas (2017) shows that top earners in municipal jobs experience a drop in nominal wages following the public revelation of wages. Baker et al. (2019) study the staggered introduction of mandated pay transparency among university faculty in Canada and show these transparency policies reduced the gender pay gap, possibly through the slowing of men's wage growth. Bennedsen et al. (2019) provide evidence that a Danish private-sector mandate to disclose relative average earnings of men and women to an employee representative had similar effects, reducing the gap and lowering wages overall. These results point to the first-order effect of transparency on bargaining power. Without an equilibrium response through bargaining, we would expect wages to rise after transparency is introduced, as transparency's direct effect allows low-wage workers to negotiate higher pay.

Our setting has the advantage of additional details about the negotiation process that allow us to model and test the microfoundations of bargaining under transparency, as well as extend the results both theoretically and empirically to include transparency's impact on hiring, worker surplus, profits and the endogenous choice of transparency. Our analysis is based on simpler labor markets than in these prior studies. While on-demand markets are not "traditional," it bears noting that they engage 7% of the US workforce, or 14 million

people (Robles and McGee, 2016). By an account based on proprietary nation-wide payroll data, 16% of workers inside firms are so-called gig workers today (Yildirmaz et al., 2020), and Farronato et al. (2019) find that a similar percentage of consumers hire contractors through online labor markets. Nevertheless, there are differences in some respects from some other labor markets.

We highlight four features about our setting that we believe are important to consider when determining where our results may generalize. First, the labor markets we study are short term and therefore workers have less time to get to know one other. Akerlof and Yellen (1990) propose that workers may lose morale and expend lower effort in the face of known wage inequality. The interpersonal comparisons at the core of these fairness concerns may be more limited in our setting than in other longer-term work arrangements studied in the literature. We show that a morale effect on effort does exist in our setting: workers reduce effort by 26% on average when there are persistent and observable wage inequalities. This figure is similar in magnitude to those found in longer work settings (Breza et al., 2018; Cullen and Perez-Truglia, 2019) and an even shorter-term work setting (Gächter and Thöni, 2010), suggesting that morale may play a similar role in many work settings. Theoretically, we model these morale concerns as in DellaVigna et al. (2016) and Breza et al. (2018) and find the bargaining channel is necessary to match our empirical findings of wage equalization.

Second, workers complete more standardized, low-skill tasks than in the general population, and their relative productivities may be more easily observed. We extend our model to allow for productivity differences between workers. If productivity differences are observable to workers, then the findings of our model hold. With unobserved productivity differences, the equilibrium outcome under full transparency is no longer identical to a single posted wage, but the other equilibrium impacts of transparency on hiring and profit share do not change. Wages are more equitable, in that workers of the same productivity receive more similar pay, and wages are more equal, in that the pay of workers with different productivities is compressed. Our empirical setting is not well suited to testing our model with unobserved productivity types, but Obloj and Zenger (2019) study the market for US academics where worker productivities vary. They find that transparency makes pay more equitable and more equal, matching our predictions.

Third, the bargaining protocol in TaskRabbit for Business, where workers submit a wage bid or employers post a wage, may differ elsewhere. We find similar directional effects of increased transparency in our field experiment, where there is an open format to negotiations and either party can start the process. Theoretically, the direction of our predictions holds as long as the workers have at least a modicum of bargaining power. Higher firm bargaining power affects the magnitude, but not the direction, of most consequences of increasing transparency. Therefore one prediction is that traditional labor markets with less

worker bargaining power will face similar but smaller effects from increases in transparency. We also extend our findings theoretically to various bargaining environments, and a search model with matching frictions and multiple firms.

Fourth, the observed level of transparency may be a less clear signal of a firm's value for labor in other settings. Firms might be compelled to disclose salaries by statute (for example, a California law mandates the disclosure of salary rages upon request), or to conform to the norms of the channels in which the job vacancy is advertised. When the choice of transparency is a weaker signal of the firms value of labor, we may expect unraveling toward transparency to be more muted. Nevertheless, we do observe high levels of transparency in many labor markets. By one careful account, over two-thirds of the U.S. workforce see transparent posted wages and nearly 75% of workers see a high degree of wage transparency at the time of hire (Hall and Krueger, 2012).[13] This share is higher among jobs with a high degree of standardization, such as entry-level jobs and jobs that require an occupational license. Other studies place this number at between 65% and 94% (Brenčič (2012); Brenzel et al. (2014); Caldwell and Harmon (2019); Niederle et al. (2006).[14]

Additional results emerge when there are other heterogenities across workers. In Appendix A.1, we study the differential effects of transparency on men and women. Empirically, we find that women have lower outside options, and our results imply that transparency will close the pay gap caused by differences in outside options. However, we also produce empirical evidence that suggests women communicate less often about wages than men. We layer this communication heterogeneity into our model and show that partial transparency can increase the gender pay gap because men are more likely than women to learn peer wages. This finding may give pause to those who advocate communication about pay to mitigate the gender pay gap. Men's communication advantage over women is predicted to disappear under full transparency, and our experimental data verifies this claim.

Our predictions *will* change for specialized jobs, or jobs at a superstar firm with a high value for labor exceeding all worker outside options. In Appendix A.8, we extend our model to such jobs and show theoretically and empirically how our predictions change. In particular, transparency benefits workers at the expense of the firm. This is supported by studies that find transparency of high-earning CEO and executive pay increases the wages of these workers (Faulkender and Yang, 2013; Mas, 2016; Schmidt, 2012).

[13]In their sample, two-thirds of employees faced transparent posted wages. Of the remaining third, 15% knew what their salary would be when they entered negotiations.

[14]In Denmark, Caldwell and Harmon (2019) find that 70% of workers see transparent posted wages, and those who do not are more likely to be in specialized roles. Brenzel et al. (2014) finds only 1/3rd of workers negotiate individually over wages at the time of hire. Brenčič (2012) show that across the US, the UK, and Slovenia, employers are more likely to post a wage offer when searching for low skilled workers. Niederle et al. (2006) consider entry-level jobs for gastroenterologists and find that 94% of employers pay common wages and "offers are not adjusted in response to outside offers and terms are not negotiable."

II. Model

II.A. Setup

Time is continuous, and is indexed by $t \in \mathbb{R}_+$. There is a single firm in the economy, and a unit mass of workers I. Each worker $i \in I$ has a private outside option $\theta_i \overset{iid}{\sim} G[0,1]$, which is the flow payment i receives when unemployed.[15] The firm has a constant-returns-to-scale production function. We assume for now that productivity of labor is common across all workers: $v \sim F[0,1]$, and is known only to the firm. (We discuss the case in which workers are heterogeneously productive in Section V.B.) All agents exponentially discount the future at rate δ, are risk neutral, and seek to maximize discounted expected flow payments. We assume that F and G are twice continuously differentiable with densities f and g, respectively. We also assume agents have strictly increasing virtual reservation values, i.e. $\theta + \frac{G(\theta)}{g(\theta)}$ is strictly increasing in θ and $v - \frac{1-F(v)}{f(v)}$ is strictly increasing in v.

Before any workers arrive, the firm selects a maximum wage it is willing to pay $\bar{w}(v) \in [0,1]$. $\bar{w}$ is not immediately observed by workers. An initial round of bargaining takes place at $t = 0$. Each worker i makes offer $w_{i,0}(\theta_i) \in [0,1]$. As in a double auction (Chatterjee and Samuelson, 1983), i is employed if and only if $w_{i,0} \leq \bar{w}$ and she receives a flow wage of $(1-k)w_{i,0} + k\bar{w}$. $k \in [0,1]$ is the known bargaining weight of the firm. If $w_{i,0} > \bar{w}$, then i is permanently unmatched from the firm, and she receives flow payments equal to her outside option θ_i.

We model transparency as the random arrival of information about current wages. At time $t \geq 0$ each matched worker observes the set of wages the firm pays to employed workers, W_t, according to an independent Poisson arrival process with (commonly known) rate $\lambda \in [0,\infty) \cup \{\infty\}$, where we take $\lambda = \infty$ to mean that the process arrives at every time $t \geq 0$. For convenience, we assume that $W_t \equiv \bar{w}$ if the firm does not have any currently employed workers.[16] Therefore, higher λ corresponds to more transparency.

When wage information arrives to worker i at time t, i and the firm renegotiate i's wage using the same bargaining protocol: i submits a new offer $w_{i,t}$ and she remains employed and receives a flow wage of $(1-k)w_{i,t} + k\bar{w}$ for all $t' \geq t$ if and only if $w_{i,t} \leq \bar{w}$. The interpretation

[15] There is a known measurability issue with the assumption of a continuum of i.i.d. random variables (Judd, 1985). A solution is to assume that worker outside options are drawn "almost" i.i.d. in the sense of Sun (2006). This solves the measureability issue and has the intuitive and intended property that the distribution of realized outside options is given by the same function $G(\cdot)$.

[16] Without this assumption, all workers under full transparency (and a measure zero set of workers for any $\lambda > 0$) face an openness issue of wanting to renegotiate wages at the earliest time $t > 0$. It is possible to deal with this issue as in Simon and Stinchcombe (1989): suppose workers can only renegotiate every $\frac{1}{N}$ periods, $N > 1$. Define a worker's payoff in continuous time as payoff of an agent at the limit point $N \to \infty$. Using this definition, even if a worker observes nothing at $t = 0$, her payoff under full transparency is equivalent to the case in which she receives a wage of $\bar{w}$ for all $t \geq 0$. For ease of notation, we continue with the simplifying, if unrealistic, assumption that $W_0 = \bar{w}$.

is that observing the exact wages of coworkers allows a worker to force the firm back to the bargaining table, perhaps due to the (unmodeled) threat of legal action that is only credible if the worker can provide direct proof that a peer is being paid more.

The timing of the stage game is as follows: First, at each time $t \geq 0$ each worker i learns W_t independently with arrival rate λ. Second, if either $t = 0$ or $t > 0$ and worker i observes W_t, then i bargains with the firm according to the protocol laid out above.

Of particular interest are the cases in which $k \in \{0, 1\}$. As in the TaskRabbit for Business data we study, $k = 0$ represents the case where workers make take-it-or-leave-it (TIOLI) offers in any wage negotiations. As a result, the transparency arrival process is particularly informative, since workers are not able to immediately deduce $\bar{w}$ from the results of the initial negotiation. When $k = 1$, the firm makes a TIOLI offer $\bar{w}$ in any negotiation. As $\bar{w}$ is time invariant, all employed workers receive wage $\bar{w}$ at each time t, and wages are never raised through renegotiation.

In Appendix A.7 we expand our model to allow workers to search for work across multiple firms, and show that many results are robust to this extension. We also consider extensions in which: the firm can accept or reject offers individually instead of picking a single $\bar{w}$, the firm can reject worker renegotiation offers at cost but without permanently unmatching, and the firm negotiates with workers under different bargaining protocols. Our main results extend to these settings, and further details are available upon request.

II.B. Equilibrium

We investigate pure strategy perfect Bayesian equilibria (PBE) of the game. Throughout, we write w_i^* to represent worker i's equilibrium wage offer at $t = 0$ assuming that she has not observed W_0. We restrict our attention to equilibria satisfying the following conditions:

A1 $0 \leq \bar{w} \leq v$ for all v. If $v \leq w_i^*$ for every worker i according to equilibrium strategies then $\bar{w} = v$.

A2 $\theta_i \leq w_i^* \leq 1$ for all i. If there is no v such that $\theta_i \leq \bar{w}$ according equilibrium strategies then $w_i^* = \theta_i$.

A3 For any λ, $\bar{w}$ and w_i^* are strictly increasing and absolutely continuous functions of v and θ_i, respectively.

A1 and **A2** restrict actions of agents who never match in equilibrium, because either the firm's value for labor is too low or the worker's outside option is too high. These assumptions rule out pathological equilibria in which, for example, $\bar{w} = 0$ for all v and all workers choose $w_i^* = 1$.

A3 assists in tractability. It also removes equilibria in which workers and the firm pool on a predetermined wage from consideration.[17]

As we will show, **A1-3** only play a role in our analysis for partial transparency levels $\lambda \in (0, \infty)$, or if $\lambda = 0$ and $k > 0$. If one is only interested in considering the impact of policies that implement full transparency, these assumptions are not necessary (with the exception of the transparency's effect on employment, which we state in Theorem 3).

There always exists an equilibrium of the game satisfying **A1-3**. If $k > 0$, all workers back out $\bar{w}$ immediately following initial negotiations, and will earn $\bar{w}$ upon renegotiating. Even if $k = 0$ and workers can not infer $\bar{w}$ from the outcome of initial negotiations, workers still receive $\bar{w}$ upon renegotiating; due to the continuum of workers entering the market at each time, in addition to our equilibrium selection criteria, workers trace out the set $[a, 1]$ for some $\bar{w} > a$ with their initial offers. Therefore, the highest wage paid by the firm (assuming it hires a positive measure of workers) is $\bar{w}$ for all $t \geq 0$.

PROPOSITION **1.** *The set of equilibria is non-empty. In any equilibrium, each worker receives $\bar{w}$ upon renegotiating.*

We sketch part of the argument for existence here, as it demonstrates an interesting connection between the rate of transparency and amount of firm bargaining power k.

Denote the *effective* level of transparency $\Lambda = \frac{\lambda}{\delta+\lambda}$ for all $\lambda \in [0, \infty)$ and $\Lambda \equiv 1$ for $\lambda = \infty$. We use Λ to represent transparency in the rest of the paper: a high rate of information arrival λ will be unimportant to workers if the discount rate δ is sufficiently higher than λ. $\lambda = \Lambda = 0$ corresponds to full privacy, while $\lambda = \infty$ and $\Lambda = 1$ correspond to full transparency.

An employed worker will receive w_i^* for the periods she is employed before learning the wages of her peers, and $\bar{w}$ thereafter. Letting $\bar{F}(x) = P(\bar{w} \leq x)$, and $\bar{G}(x) = P(w_i^* \leq x)$, we show in Appendix A.2 that each worker i and the firm respectively solve

$$w_i^* = \underset{w_i}{\operatorname{argmax}} \int_{w_i}^{1} \left((1 - \Omega) w_i + \Omega x - \theta_i \right) \bar{f}(x) dx \tag{1}$$

$$\bar{w} = \underset{w}{\operatorname{argmax}} \int_{0}^{w} \left(v - (1 - \Omega) y - \Omega w \right) \bar{g}(y) dy \tag{2}$$

where

[17] Leininger et al. (1989) suggest similarities between the set of continuous equilibria and a set of discontinuous equilibria in static double auctions, and so we do not believe this to be a conceptually limiting constraint. We discuss the connection of our game to static double auctions below.

$$\Omega = (1 - \Lambda)k + \Lambda \tag{3}$$

Equations 1 and 2 lead to two conclusions. First, these are the same objective functions as those in a static double auction between a single worker whose type is drawn according to G, and a single firm whose type is drawn according to F, with a bargaining weight of Ω on the firm's offer. Therefore, the set of equilibria of this static double auction corresponds to the set of equilibria of our game.

The first order conditions for workers and the firm are, respectively:

$$w_i^* - \theta_i = \overbrace{(1 - \Omega)}^{\text{direct effect}} \cdot \underbrace{\frac{1 - \bar{F}(w_i^*)}{\bar{f}(w_i^*)}}_{\text{indirect effect}} \tag{4}$$

$$v - \bar{w} = \overbrace{\Omega}^{\text{direct effect}} \cdot \underbrace{\frac{\bar{G}(\bar{w})}{\bar{g}(\bar{w})}}_{\text{indirect effect}} \tag{5}$$

The set of equilibria satisfying **A1-3** is characterized by solutions of these first order conditions, and in particular, this set of equilibria is non-empty (Satterthwaite and Williams, 1989, Theorem 3.1).

Second, increasing the level of transparency Λ has a similar effect as increasing the firm's nominal bargaining power k: both increase Ω, the *de facto bargaining power* of the firm. We see this from the fact that $\Omega = (1 - \Lambda)k + \Lambda$. There are two pairs (Λ, k) and (Λ', k') that yield any $\Omega \in (0, 1)$, where $\Lambda < \Lambda'$ if and only if $k > k'$.

Ω is submodular in Λ and k, implying that a fixed increase in transparency is more impactful the smaller k, the nominal bargaining power of the firm. When $\Lambda = 1$, $\Omega = 1$ implying that at full transparency, the nominal bargaining power k does not affect the equilibrium outcome. Similarly, when the firm has all of the nominal bargaining power, i.e. $k = 1$, the equilibrium outcome is constant in the level of transparency Λ. This matches our earlier, intuitive descriptions: under full transparency, all workers learn $\bar{w}$ immediately and secure this wage if it is higher than their outside option, regardless of k. The firm therefore "posts" $\bar{w}$ knowing that all employed workers will receive this wage. When $k = 1$ the firm makes an initial TIOLI offer $\bar{w}$ to each worker, and all workers with a lower outside option will be employed by the firm at this wage. Wages are transparent to workers, as all workers know the firm pays a common wage.

PROPOSITION **2.** *The function $\Omega : [0, 1]^2 \to [0, 1]$ which maps Λ and k into the firm's de facto bargaining power is increasing in both arguments, submodular, and surjective.*

We have thus far assumed that workers cannot rebargain until pay information arrives.

If workers can rebargain at will, any $k > 0$ will have the same equilibrium outcome as $\Omega = 1$ since all employed workers will back out $\bar{w}$ at $t = 0$ and immediately renegotiate by offering this amount (sidestepping the technical issue as in Footnote 16). When $k = 0$, workers do not learn $\bar{w}$ until the arrival of pay information. Interestingly, we show that even when $k = 0$, no employed worker will ever renegotiate her wage in equilibrium until the first moment information arrives. The key step in proving this result is showing that a worker does not learn exploitable information about $\bar{w}$ if her initial offer is accepted. Any worker strategy that says "offer w when initially hired at time 0 and offer $w' > w$ at time $t > 0$ if I have not learned the wages of my coworkers" is not optimal, because if offering w' at time t improves the expected utility of the worker, she would be even better off offering w' at time 0.[18] We formalize and prove the following result in the appendix.

PROPOSITION **3.** *Suppose each worker can elect to initiate wage renegotiations at any time t that she is employed at the firm, potentially renegotiating infinitely often. If $k = 0$ then no worker will ever renegotiate her wage in equilibrium until observing the wages of others, at which point she will offer and receive $\bar{w}$.*

This result is of particular interest to our empirical setting, where workers make TIOLI offers to their employer, i.e. $k = 0$. This further suggests that wage renegotiations we observe on TaskRabbit for Business are due to transparency.

II.C. Main Results

For any $k < 1$ the optimal bidding and wage-setting policies of workers and the firm, respectively, are interdependent for any $\Omega \in (0, 1)$. Workers decide how aggressively to bid depending on how the firm sets $\bar{w}$, while the firm sets $\bar{w}$ as a function of how aggressively the workers bid. While there exists a unique equilibrium when $\Omega \in \{0, 1\}$, Satterthwaite and Williams (1989) show that there exists a continuum of equilibria satisfying Equations 4 and 5 for $\Omega \in (0, 1)$. This set lacks natural ordering, limiting the possibility for general claims about the entire set of equilibria. However, experimental evidence in Radner and Schotter (1989) suggests that equilibria in which w_i^* and $\bar{w}$ are linear functions of θ_i and v, are focal and most likely to be played in practice. We similarly observe linear bidding strategies in our experiment.

To focus our attention on linear equilibria, we restrict attention to a two-parameter family of power law distributions of worker outside options and firm values. We show that this family–given in Equation 6–admits a unique linear equilibrium for any value of trans-

[18] This reasoning is shared in Lazear (1986) and Tirole (2016).

parency.[19] In this section we study the properties of the linear equilibrium, and analyze the effects of increasing transparency.

$$\begin{aligned} F(v) &= 1-(1-v)^r, \quad r>0 \\ G(\theta) &= \theta^s, \qquad\qquad\quad s>0 \end{aligned} \tag{6}$$

As r increases, v is on average lower and as s increases, θ is on average higher. Therefore, increasing r or s reduces the average surplus from employment. We define a linear equilibrium below and show that distributions of this type admit a unique linear equilibrium. This result extends work by Chatterjee and Samuelson (1983) who show existence of a linear equilibrium when F and G are uniform, corresponding to the case in which $r = s = 1$. We note again that these distributional restrictions are unnecessary for our analysis if one is only interested in comparing full transparency to full privacy.

Definition **1.** *A* linear equilibrium *is a pure strategy perfect Bayesian equilibrium satisfying* ***A1-3****, where $\bar{w}$ is a linear function of v whenever a positive mass of workers offers $w_i^* \leq v$, and where w_i^* is a linear function of θ_i whenever there is positive probability that $\theta_i \leq \bar{w}$.*

Proposition **4.** *For any pair of distributions within the family described in Equation 6 there exists a unique linear equilibrium.*

In what follows, we assume that $k < 1$, because as we discussed in the previous section, transparency has no impact on the bargaining process when $k = 1$ (i.e. $\Omega = 1$ for all Λ).

Supply and Demand effects Workers initially offer premia over their outside options, $w_i^* - \theta_i \geq 0$. Similarly, the firm sets a markdown below its value for labor, $v - \bar{w} \geq 0$. We show that both $\bar{w}$ and w_i^* are decreasing in Λ; with increased transparency the firm reduces the highest wage offer it accepts to avoid information spillovers across workers (which we call the **demand effect**), and workers make more conservative initial offers as they anticipate quickly, and risklessly renegotiating and receiving $\bar{w}$ (which we call the **supply effect**).

Proposition **5.** *Consider the unique linear equilibrium given the family of distributions in Equation 6. $\bar{w}(v)$ and $w_i^*(\theta_i)$ are weakly decreasing functions of Ω for all v and θ_i. As $\Omega \to 0$, $\bar{w} \to v$ for all v. As $\Omega \to 1$, $w_i^* \to \theta_i$ for all θ_i.*

Pay Inequality Initial wages are more dispersed with higher transparency. Over time, wages are equalized as workers renegotiate to a common, higher wage. Ultimately expected earnings are more equal under transparency.

[19] The approach of making parametric assumptions to ensure linear equilibrium is common. One recent example on CEO pay is Edmans et al. (2012). Power law distributions are commonly observed in economic situations such as ours, including worker income and firm productivities. See Gabaix (2009, 2016) for details.

Fix $k < 1$. We compare the earnings of workers i and j who are hired in equilibrium under both of two transparency levels, $\Lambda' < \Lambda''$, so we do not confound employment effects of increasing transparency.[20] For any two workers i and j with $\theta_i > \theta_j$ who are hired under both Λ' and Λ'', there are two effects. First, the **supply effect** incentivizes workers to reduce initial wage offers. We find that in equilibrium, since j has a lower outside option than i, j reduces her initial offer more than i. Second, higher transparency decreases the expected time it takes before both workers renegotiate to $\bar{w}$, reducing dispersion of their earnings as $\bar{w} - w_j^* > \bar{w} - w_i^*$. The first effect increases the initial wage gap between i and j, however, the latter effect dominates in the long run, leading to more equalized expected earnings, regardless of δ.

THEOREM **1.** *Consider the unique linear equilibrium given the family of distributions in Equation 6. Let $\theta_i > \theta_j$, and $1 > \Omega'' > \Omega'$. For almost all workers i and j hired in equilibrium under both Ω' and Ω'' :*

1. *The difference in initial offers $w_i^* - w_j^*$ is strictly higher under Ω'' than Ω', and*
2. *Let $T(\Omega, v, \theta_k)$ be the equilibrium expected discounted total earnings of a worker k with outside option θ_k under transparency level Ω and firm value v conditional on k being employed at the firm. Then $T(\Omega'', v, \Omega_i) - T(\Omega'', v, \theta_j) < T(\Omega', v, \theta_i) - T(\Omega', v, \theta_j)$ and $T(\Omega'', v, \theta_i) - T(\Omega'', v, \theta_j) \to 0$ as $\Omega'' \to 1$.*

Note that the first point in the above theorem does not apply to full transparency (and similarly when $k = 1$); there is a discontinuity because all workers make their initial wage offers *after* seeing the wages of their coworkers. Therefore $T(1, v, \theta_i) - T(1, v, \theta_j) = 0$, so there is never wage dispersion among employed workers.

Pay equalization does not imply equalized worker surplus. We show workers offer premia that are decreasing in outside option, implying that more equal earnings caused by transparency result in greater inequality of worker surplus.

COROLLARY **1.** *Consider the unique linear equilibrium given the family of distributions in Equation 6. Let $\theta_i > \theta_j$, $1 > \Omega'' > \Omega'$. For almost all workers i and j hired in equilibrium under both Ω' and Ω'' :*

[20] The restriction that workers be hired under both transparency levels is necessary, as we show in Theorem 3, because increasing transparency can increase the hiring rate. A previously unemployed, high outside option worker may find employment only when transparency is increased. To make this point concrete, take some small $\epsilon > 0$ and consider increasing transparency from Λ' to $\Lambda'' = \Lambda' + \epsilon$, such that more workers are employed in equilibrium under Λ''. In Appendix A.2 we show that w_i^* and $\bar{w}$ are continuous in Λ and so the expected earnings of any worker j hired under both transparency regimes is barely affected by an ϵ increase in transparency. However, a worker i who over-negotiates at level Λ' receives her outside option θ_i for her entire duration in the market, while if she manages to find employment at the firm under Λ'' her average total earnings will be greater than, and bounded away from, θ_i (as she always asks for a premium $w_i^* - \theta_i > 0$). But note that $\theta_i > \theta_j$, so the total earnings of i and j are not equalized by increased transparency.

1. *The difference in initial surplus* $(w_j^* - \theta_j) - (w_i^* - \theta_i)$ *is smaller under* Ω'' *than* Ω' *and*
2. *Let* $S(\Omega, v, \theta_k)$ *be the equilibrium expected discounted total surplus of a worker* k *with outside option* θ_k *under transparency level* Ω *and firm value* v *conditional on* k *being employed at the firm. Then* $S(\Omega'', v, \theta_j) - S(\Omega'', v, \theta_i) > S(\Omega', v, \theta_j) - S(\Omega', v, \theta_i)$*, and* $S(\Omega'', v, \theta_j) - S(\Omega'', v, \theta_i) \to \frac{\theta_i - \theta_j}{\rho + \delta}$ *as* $\Omega'' \to 1$*.*

In sum, transparency decreases differences in expected pay across workers, but it also increases differences in expected worker surplus. We discuss these two notions of equality in Section IV.B.

Division of Surplus Increasing pay transparency increases the expected profits of the firm, decreases average worker surplus, and lowers average discounted wages. The **demand effect** causes the firm to limit its demand, similar to the pricing strategy of a monopsonist. Due to the information spillover caused by transparency, the firm can commit to reducing $\bar{w}$ as Ω increases. This restricts the extensive margin of labor (the proportion of workers it hires) and increases the intensive margin (profit per worker hired). Simultaneously, the **supply effect** reduces worker initial offers, which similarly benefits the firm.

Although raising transparency increases the rate at which workers receive wage $\bar{w}$, it lowers both w_i^* and $\bar{w}$ in equilibrium. The overall effect is to shift de facto bargaining power to the firm, benefiting the firm at the expense of workers. For clear intuition, consider the extreme cases of full privacy ($\Lambda = 0$, $k = 0$) and full transparency ($\Lambda = 1$). In the former, each worker makes a once-and-for-all offer to the firm as no worker ever renegotiates. Under full transparency, there are perfect information spillovers, and each worker learns the wages of others within the firm at the instant she is hired, before initial negotiations. Therefore, every employed worker will demand and receive exactly $\bar{w}$ for each period of her employment. This is equivalent to the firm making a once-and-for-all offer to workers. The main result of Williams (1987) implies that each party prefers to be the one making the once-and-for-all offer to the other.

THEOREM **2.** *Consider the unique linear equilibrium given the family of distributions in Equation 6. The expected equilibrium profit of the firm is strictly increasing in* Ω*. The expected average equilibrium surplus of workers and expected average discounted wages conditional on employment are strictly decreasing in* Ω*.*

This theorem takes expectations over firm and worker types. In particular, it does not imply that a move from full privacy to partial transparency improves profits for *all* firm types or decrease expected surplus for *all* worker types. Example 1 in the Appendix shows that a move from full transparency to partial transparency benefits high-value firm types, and a move from full privacy to partial transparency benefits low outside option workers.

Hiring Rate Increasing transparency has competing effects on the hiring rate. Let $\bar{w}_\Omega$ denote the maximum wage the firm pays and $w^*_{i,\Omega}$ the initial offer of worker i for given de facto firm bargaining power Ω. When transparency increases, increasing Ω' to Ω'', the **demand effect** lowers the hiring rate. $\bar{w}_{\Omega''} \leq \bar{w}_{\Omega'}$ meaning that there are fewer workers with $\theta_i \leq \bar{w}_{\Omega''}$ who are eligible for employment. The **supply effect** increases the hiring rate. $w^*_{i,\Omega''} \leq w^*_{i,\Omega'}$ for all i so fewer workers over-negotiate by initially offering $w^*_{i,\Omega''} > \bar{w}_{\Omega'}$. The primary cause of unemployment when Ω is low is that workers act too much like monopolists in initial negotiations, and when Ω is high is that the firm acts too much like a monopsonist. We show that a more even split of the de facto bargaining power is employment maximizing.

THEOREM **3.** *Consider the unique linear equilibrium given the family of distributions in Equation 6. The expected proportion of workers hired in equilibrium is concave in Ω and maximized at $\Omega^* = \frac{1-\mathbb{E}(\theta)}{1+\mathbb{E}(v)-\mathbb{E}(\theta)}$. Moreover, the ex-post hiring rate is submodular in v and Ω for the set of firm types that hire a positive mass of workers.*

An interior level of Ω maximizes the expected hiring rate. Due to the concavity of the expected hiring rate in Ω, any increase in transparency will decrease expected hiring if the firm has sufficient nominal bargaining power, i.e. when $k \geq \Omega^*$. On the other hand, when $k < \Omega^*$, then (sufficiently small) increases in transparency will increase expected hiring. In general, either full privacy or full transparency is employment *minimizing*.

Ω^* is decreasing in both $\mathbb{E}(v)$ and $\mathbb{E}(\theta)$.[21] As $\mathbb{E}(v)$ converges to 0 full transparency becomes close to employment maximizing, and as $\mathbb{E}(\theta)$ converges to 1 full privacy becomes close to employment maximizing. For intuition, we return to Proposition 5. As $\mathbb{E}(v)$ decreases, the firm's markdown $v-\bar{w}$ is likely to be small regardless of Ω. Therefore, increasing transparency does not greatly reduce the number of workers with $\theta_i < \bar{w}$. But by increasing transparency, workers will shade down their initial offers w^*_i, reducing the number of workers who over-negotiate. Similarly, as $\mathbb{E}(\theta)$ increases, most workers offer small premia $w^*_i - \theta_i$ regardless of Ω. Increasing transparency has little effect on these premia, but instead discourages the firm from setting a large markdown.

An increase in transparency increases hiring only for sufficiently low value firm types. Consider an increase in transparency leading to an increase in Ω' to $\Omega'' > \Omega'$. The submodularity of the ex-post hiring rate in v and Ω means that the firm hires more workers under Ω'' than Ω if and only if v is below a particular threshold. It also implies that the value of Ω that maximizes the ex-post hiring rate is weakly decreasing in v (Topkis, 1998). These comparative statics on ex-post hiring also hold for ex-post social surplus. In fact, the ex-post maximizer of the hiring rate also maximizes ex-post social surplus. Because each employed worker earns a wage weakly greater than her outside option, in equilibrium each

[21] The expected match surplus is $\mathbb{E}(v) - \mathbb{E}(\theta)$, so $\Omega^* = \frac{1-\text{expected outside option}}{1+\text{expected match surplus}}$.

employed worker increases social surplus by $v-\theta_i > 0$, implying that social surplus is strictly increasing in the hiring rate. Therefore, Ω'' increases ex-post social surplus if and only if v is below some threshold.

Endogenous Selection of Transparency Until now, we have been studying the effects of transparency from an ex-ante perspective, before the firm draws v. This perspective is aligned with that of a government agency instituting pay transparency measures without detailed information about the value of labor to each firm. We next ask what transparency level firms choose with access to private information about v.

We allow the firm to select Λ to maximize its profits immediately after seeing the draw of v, and at the same time that it selects $\bar{w}$. Workers observe only the choice of Λ. We find that in equilibrium the firm selects full transparency regardless of its draw of v. This result does not depend on additional linearity or distributional assumptions.

THEOREM **4.** *Suppose $k < 1$ is exogenous and the firm selects Λ as a function of v. There is an essentially unique equilibrium satisfying **A1-3**. In it, the firm selects $\Lambda = 1$ for all v.*

The firm sets $\Lambda = 1$ in equilibrium due to unraveling (Milgrom, 1981). For intuition, suppose the firm can only select $\Lambda \in \{0, 1\}$. Toward a contradiction, let v_L be the infimum value for which the firm selects $\Lambda = 0$. Then upon arriving at the firm, workers will infer that the firm's value is at least v_L, and so every worker will offer at least v_L. As a result, when the value of the firm is v_L it will make 0 profits unless it deviates to selecting $\Lambda = 1$. But if this firm type deviates, there is a new "v_L." Inductively there cannot be an equilibrium in which there is a positive measure of firm types playing $\Lambda = 0$. The equilibrium in which the firm selects $\Lambda = 1$ for all v can be supported with the off-path beliefs that a deviating firm has value $v = 1$ with probability 1.

We have not formally modeled the choice of workers to "bury their heads in the sand" and ignore wage information. Nevertheless, a richer model that allows each worker to ignore information would lead to the same equilibrium outcome: for fixed $\bar{w}$ higher transparency is better for almost every employed worker, and no single worker will affect the equilibrium payoff, and therefore actions, of the firm.

III. Empirical Evidence from TaskRabbit for Business

III.A. Institutional details

To test our theoretical predictions, we use administrative data from TaskRabbit for Business, available between 2009 and 2014. After June of 2014, TaskRabbit for Business, in the form described here, was no longer offered. TaskRabbit for Business specialized in sourcing

workers for local short-term jobs, functioning like a temp agency. The platform was active in 19 metropolitan areas across the U.S. during this period. The employers using the platform ranged from older and more established firms, including approximately 10% of Fortune 500 companies, to household employers (comprising roughly one quarter of our sample).

Our research concentrates on jobs that are posted as one-off tasks. Among the most common categories of jobs are deliveries, packing and shipping, marketing, event staffing, and user testing. Employers can observe workers' profiles, which include the number of prior jobs completed on the platform, a rating out of 5 stars, and a short bio. Previous worker wages are not shown.

Bargaining Environment Employers post a description of the job, details about the exact location, and a deadline for completion. Workers search through these postings and submit bids for jobs they are interested in completing. Alternatively, the employer can choose to post a fixed price, and the first worker to accept is matched. The platform receives a fixed percentage of the final transaction, which includes any bonuses made after the job has started. This fee changes very little over the time period we study and remains in the neighborhood of 20%.

Jobs frequently involve face-to-face contact between employers and employees, allowing for the possibility of on-the-job wage bargaining. Employers can adjust pay through the platform after the job begins, to any amount higher than the initially agreed-to wage.[22]

The initial contract agreed to can be cancelled by providing a reason to the platform. TaskRabbit for Business has a three-strike rule. After three cancellations, a user is not permitted to use the site again. During the window between when the match is made and the job is completed, money is held in escrow and released to workers by default after a predetermined deadline.

Once the job begins, several additional frictions make canceling costly for both parties. TaskRabbit for Business is a spot market designed for urgent tasks: 97% of completed tasks are finished within three days of posting. Additionally, the rate of bids received slows considerably after posting. The median job receives 1 offer in the first hour, and 1 every 4 hours over the first day. Taken altogether, finding a replacement worker after the job begins would likely result in costly delays. Similarly, workers can not costlessly transition to another job. Because these jobs involve short-term in-person tasks, travel costs are high relative to the final transaction price.

Employers have the opportunity to leave a public rating out of 5 stars for a hired worker.

[22]TaskRabbit for Business reserves the right to revoke user privileges should any activity suggest circumventing the online contract, or making payments off the platform. However, we do not rule out the possibility that working relationships continue off the platform. For robustness, we replicate results to exclude and include employers that never return to the site after their initial jobs are completed.

There is no reputation system for employers. With very few exceptions (less than 1% of the time) employers leave 4 or 5 star reviews, but they frequently decline to rate workers. Employers only observe worker average ratings, not the number of declined reviews. We add to our measure of underlying performance "missing" reviews, based on prior research that missing reviews are skewed toward negative feedback.[23]

Measuring Transparency We measure pay transparency on TaskRabbit for Business several ways. Our first measure is whether the job post itself includes a posted price publicly visible to all workers. The posted price can either be text embedded in a job description or a fixed posted price associated with the job posting format selected. We classify these as fully transparent jobs.

Our second measure is based on the physical proximity of workers in multi-worker jobs and the length of time they overlap. For example, a retail branch might outsource the boxing of holiday gifts at the store (co-located workers), or outsource the distribution of catalogues in different neighborhoods (separated workers). We classify co-located jobs as partially transparent, as workers may discuss wages while on the job.

Longer co-location on the job allows for a higher level of partial transparency. We use the street address of a job to classify proximity and we supplement it with survey evidence. We hire approximately five thousand online workers to read through the detailed job descriptions and report key attributes, including how conducive the setting is to coworker communication and length of time together. Those surveyed expected workers to learn about each other's bids 40% of the time when co-located (according to our administrative record of physical proximity), and only 7% of the time if the workers are physically separated according to the street address.

In Table A1 we compare workers and tasks on TaskRabbit for Business across different transparency classifications. Panel (a) shows the similarity of observable characteristics between multi-worker co-located jobs and multi-worker separated jobs across a number of important dimensions such as the number of workers required, the dispersion and level of bids received, and winning bids. We also show that the dispersion in performance ratings are similar across job types. Lastly, we use unsupervised natural language processing to show the text of the job descriptions is also highly similar.[24] Panel (b) compares job postings that

[23]The literature on user-generated content has identified performance metrics that the platform collects but is not visible to the users. Nosko and Tadelis (2015) show the "sound of silence," or missing reviews, on eBay is skewed toward negative feedback. We find the share of missing reviews on TaskRabbit for Business predicts whether an employer returns to the platform, TaskRabbit's central measure of employer satisfaction, and the worker rating conditional on receiving a rating.

[24]When analyzing partially transparent jobs, we restrict our sample to categories of jobs which lend themselves to both multi-worker separated and co-located jobs (a minimum of 20% of each type), but we also show our results are not sensitive to this restriction.

do and do not mention price in the text that workers read before placing a bid. We show balance in the length of the job description and text used. We find a large and significant difference in the submitted bids to job postings that do and do not mention price, a result we discuss in Section II.C. Panel (c) describes local market characteristics at the city-month level. The mean price for a job is $56 conditional on closing, and 46% of jobs close on average. Over the panel, the average age of a market is 17 months. On average, 43% of jobs are listed with a posted take-it-or-leave-it price.

III.B. Empirical Tests and Results

Empirical Test of Model Bargaining Assumptions Our model has two clear empirical implications for the outcome of a re-bargaining process in partially transparent jobs. First, our model predicts that workers who learn about the pay of coworkers use this to ask for the same pay as the highest earner observed, having interpreted this as evidence the employer values the work at least that much. Empirically, this implies that workers who renegotiate on the job will frequently receive final pay equal to the highest paid coworker on-site at the same job (SF1).[25] Second, we posit that the act of discovering the pay of coworkers, and therefore renegotiating, does not depend on the level of initial pay inequality there is among coworkers. This implies that differences in initial wages do not predict whether a worker receives a raise (SF2). This prediction contrasts with those of models where the employer directly cares about the extent of inequality and proactively levels the playing field when differences are too extreme, which we study in Section V.A.

SF1: In co-located jobs, employers equalize wages if they compress pay at all.

We estimate the relationship between the amount that the worker bids below the highest bidder and the amount of the raise paid at the end of the job, conditional on receiving a raise. In this and subsequent regressions, we index workers by i, employers by j and jobs by k. We include job-specific fixed effects α_k and individual performance controls X_i at the time of job k in the equation below.

$$\text{Raise}(\%\text{ above bid}|\text{Raise} > 0)_{ik} = \beta_0 + \beta_1 \cdot (\text{Amt. under top bid } (\%))_{ik} + \phi \cdot X_i + \alpha_k + \epsilon_{ik}$$

$\beta_1 = 1$ implies that the raise amount is equivalent to the initial gap between bids, so the employer equalizes pay. In Table I, Col. 6, we estimate a precise relationship, $\hat{\beta}_1 = 0.96$ (s.e. $= 0.074$), statistically indistinguishable from 1 (p-value=0.51).

[25]The average number of workers overlapping on the job is 2.5. Unlike in our theoretical model, there are not infinitely many workers per job, implying that workers do not learn the maximum wage the employer is willing to pay for labor. Nevertheless, our model's prediction on the direct effect of transparency holds with finitely many workers: in equilibrium upon seeing the pay of her peers, a worker will demand the highest wage she observes in renegotiations.

SF2: In co-located jobs, initial inequality is uncorrelated with employers' decisions to compress pay.

We test whether the difference between the highest accepted bid and a given worker's bid predicts whether this worker receives a raise.

$$\text{Any Raise(Yes=1)}_{ik} = \beta_0 + \beta_1 \cdot (\text{Amt. under top bid (\%)})_{ik} + \phi \cdot X_i + \alpha_k + \epsilon_{ik}$$

$\beta_1 = 0$ implies no conditional correlation between a worker's bid and the probability of receiving a raise.[26] We estimate a small effect that is statistically indistinguishable from 0. In Table I, Col. 3, $\hat{\beta}_1 = 0.025$ (s.e.=0.019) implying that a 10% increase in the pay gap would only increase the likelihood of a raise by 0.2 percentage points or less than a 1% increase in the baseline rate of raises, 22%.

We argue that renegotiations exhibit a pattern that strongly suggests coworker communication about pay is a key driver. Face-to-face communication channels are critical: in higher transparency, co-located settings 22% of workers receive pay that is higher than their bids, as opposed to 7% in lower transparency, physically-separated work settings. While there are other candidate mechanisms for why pay is renegotiated when workers are co-located, such as productivity spillovers, they fall short of matching key facts. Importantly, performance ratings given by the employer are no less dispersed when workers are co-located (Table A1, Table A4), suggesting that perceived performance differences cannot explain the wage equalization we observe. More generally, there is an economically small and statistically insignificant correlation between bids and persistent performance differences between individuals. As a result, any systematic pattern of spillovers from high to low productivity types does not appear to raise the performance of the low bidder or the pay of the low bidder per se (see Appendix A.3 for more details).

Wage Equalization Theorem 1 states that increased transparency leads to equalization in the wages of employed workers. We present a visual depiction of wage equalization in partially transparent, co-located multi-worker jobs, compared to wage equalization in pay-private, physically-separated multi-worker jobs.

Figure I shows the variance in wages for workers assigned to the same job on TaskRabbit for Business. Each dot represents a multi-worker job. The x-axis represents the variance of accepted bids, and the y-axis represents the variance of ex-post payments of all workers in a particular job. Observations that fall beneath the 45 degree line depict employers who reduce the variance in ex-post payments relative to the variance of ex-ante bids. There is more wage equalization in co-located jobs than in physically separated jobs, as indicated by the fact that

[26] We exclude the highest bidder in this analysis.

Figure I: Variance in final pay vs. bids

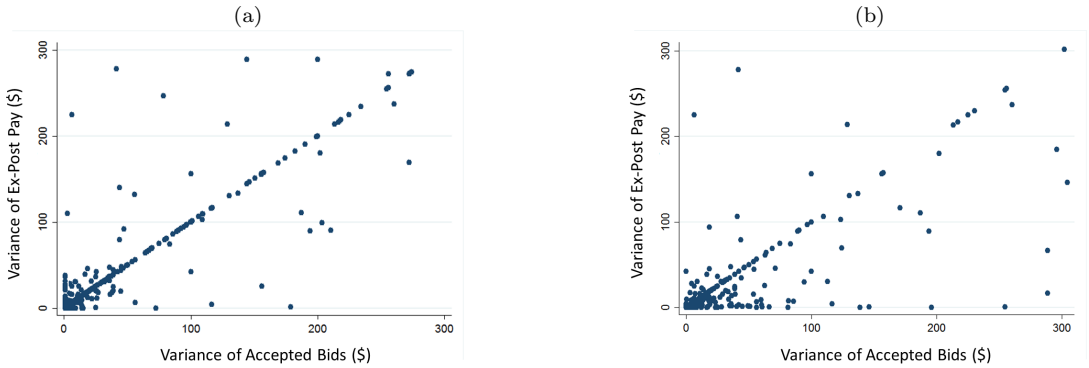

Notes: Each observation summarizes the variance in pay among the workers that have been selected for multi-worker separated jobs (Panel (a), N=184) and multi-worker co-located jobs (Panel (b), N=386). The x-axis is the variance in the bids accepted for a job, in dollars. The y-axis is the variance in the final payout. An observation below the 45 degree line indicates that wages are compressed on the job, while and observation above the 45 degree line indicates that wages become dispersed on the job. Observations along the 45 degree line generally capture the decision of the employer not to raise the wage of any worker. Since initial bids are binding, the employer cannot pay less than the bid. Marketing, user-testing and delivery are among the most common task categories in Panel (a). Packing and shipping, delivery, and event staffing are among the most common categories in Panel (b). We reject the null hypothesis that the slopes of these two plots, 0.96 (0.17) and 0.56 (0.08) are equal, P-value = 0.046.

the slope of the best-fit line for physically separated jobs is close to 1 (slope=0.96 (s.e=0.17)) while we observe a much shallower slope for the best-fit line among co-located jobs, slope =0.56 (s.e.=0.08). We reject the null hypothesis that the slopes of these two plots are equal (p-value = 0.046). We do not expect full wage equalization because, as we illustrate in Table I, the workers who renegotiate are a subset of all workers, who learn about the bids of their coworkers on the job. In Appendix Table A5 we carry out a similar test at the level of individual bids with controls for worker performance and again find wage equalization only when workers are physically co-located. We show that a 10% gap between a worker's initial bid and that of the highest bidder results in a 5% raise in final pay for the low bidder when workers are co-located (p-value<0.001) and this estimate is stable with the addition of controls. The impact of coworker bids on final pay when workers are physically separated cannot be statistically distinguished from 0. A 10% raise in the pay gap results in a 0.3% decrease in the pay for the low bidder (p-value=0.24). This result includes job-specific fixed effects, isolating variation in the co-location of workers within employer and job posting.

We interpret coworkers' bids as having a causal effect on each other's final pay when co-located. While co-location is not randomly assigned per se, the overlap between any two workers and the variation in the privately submitted bids of coworkers prior to the start of the job is as good as random.[27] We find that even within a single job posting that may be

[27]That is, a worker is unlikely to initially know whether a coworker has bid more than she has, and if so, by how much. From any worker's point of view, a coworker's pay is a non-degenerate random variable

repeatedly posted by the same employer, differences in the initial bids of coworkers predicts final pay, if and only if workers are co-located.

We provide empirical support for the assumptions necessary for a causal interpretation. We show in Table A1 that observables between separated and co-located jobs are similar. Moreover, employers do not appear to select different types of workers as a function of co-location and workers do not appear to adjust their initial bidding strategy as a function of co-location with other workers. Workers very likely do not anticipate multi-worker settings. These jobs comprise fewer than 5% of all jobs posted on the platform and appear as a single post just like any other single-worker job posting. Employers might not anticipate the extent of communication about pay on the job either. We describe these points in detail in Appendix A.5.

Wage Level Theorem 2 predicts higher levels of transparency are associated with lower average wages. In TaskRabbit for Business, we compare the difference in bids, sent from the same workers for comparable standardized jobs, in the case where the employer does mention the price of the job in the posting directly (high transparency) and the case where the employer does not mention the price in the posting (low transparency).[28]

Table II reports our estimates of the following specification, which includes worker fixed effects, α_i, and past performance evaluations X_i, as well as job category fixed effects ω_k.

$$\{\log(\text{Bid Amount})_{ik}, \log(\text{Winner's Pay})_{ik}, \log(\text{Winner's Hourly Pay})_{ik}\} = \beta_0 + \beta_1 \cdot \text{transparent}_{ik} + \phi \cdot X_i + \omega_k + \alpha_i + \epsilon_{ik}$$

When the employer mentions a price in the job description ($\text{transparent}_{ik} = 1$) the same worker carrying out the same type of job earns hourly wages that are 7.8% lower, total pay is 13.6% lower, and bids are 9.6% lower as captured by β_1. All point estimates are statistically significant (p-values < 0.001).

Endogenous Choice of Transparency Theorem 4 states the market unravels toward full transparency in equilibrium. The intuition is that employers with low values for labor benefit the most from selecting transparency. Therefore, the choice of privacy signals to workers that the employer has a high value. This leads to higher bids, which induces more employers to prefer transparency.

that can be either larger or smaller than the worker's own pay, when the employer has a per-worker budget drawn according to a distribution with full support over an open set.

[28] In Table A1 we show jobs with and without price mention are comparable along observable characteristics including the length of the description and text topics identified by unsupervised machine learning. Our preferred specification still includes worker fixed effects and job category fixed effects, further restricting the comparison to jobs that attract the very same workers.

We observe the predicted unraveling toward greater transparency in TaskRabbit for Business. Local markets where TaskRabbit for Business has been active for a longer period of time are associated with a higher proportion of transparent posted price jobs. Analyzing a balanced panel of city-months and adjusting for changes in job-category composition, Table V shows that the proportion of transparent posted price jobs rises by 1% per month.[29] The market-level movement toward transparency cannot be explained by the changing composition of jobs (Col. 2, 4) or size of the market (Col. 3, 4).[30] It also is not consistent with employer experimentation as an explanation.[31]

Also as predicted by our theory, we find empirically that employers with below-median value are most likely to use transparent posted prices. In Table A9 we find that an indicator for below-median value predicts between 4.5 to 6.5 percentage point increases in the likelihood of choosing a transparent posted price.

IV. Empirical Evidence from a Field Experiment

IV.A. Field Experiment Research Design

Design Overview We conduct a field experiment to further test our findings in a controlled environment where we can directly measure worker productivity, outside options, and employer profits. This allows us to explore additional notions of equality, such as parity in worker surplus. In our main treatments, we hire 262 managers and 655 workers from an online labor market, Amazon Mechanical Turk, who are tasked with negotiating wages for transcription services that take on average three hours to complete. We randomize transparency by restricting wage negotiations between a manager and workers to either a common chat room, or separate chat rooms. We document all interactions in these chat rooms, and we place no restrictions on the ways in which workers and managers can bargain. Common chatrooms introduce high transparency, which bear a greater resemblance to the non-anonymous posted salaries of public sector employees or posted wage formulas[32] than to the informal on-the-job discussions of TaskRabbit for Business workers.

[29] Our theory predicts immediate unraveling to full transparency, but observed unraveling in labor markets typically takes time, as discussed in Roth and Xing (1994).

[30] Einav et al. (2018) argue that a shift toward posted prices on eBay is largely due to changes in user preferences. While we cannot rule out a change in user preferences, we do not believe it to be the chief driver of this unraveling. The shift we observe is far faster than that in Einav et al.. Furthermore, TaskRabbit staggered entry into different markets, and therefore, we observe wide variance in market age. Despite this, we observe a strong linear trend toward posted price in markets of different ages.

[31] Over half of employers only post a single task on the platform, and only 10% post more than 6 tasks in their lifetime on the platform. Nearly 80% of employers use either posted price for all jobs or accept bids for all jobs.

[32] See, for example, https://open.buffer.com/salary-formula/ accessed 2/13/2019.

In Section V.A we discuss an additional cross-randomized treatment arm varying whether a manager and worker are able to renegotiate once the job begins and further wage information becomes visible. This treatment arm is used to test the relative importance of social concerns surrounding pay transparency as we are able to capture the morale consequences if the manager were unable to equalize pay.

Procedure Participants are recruited from Amazon's Mechanical Turk between October 2016 and May 2017. Participants are assigned to either the role of worker or manager and informed that their participation is voluntary and part of an academic experiment. All participants are given the following instructions: managers and workers are tasked with negotiating a per-page rate for completing text-to-text transcription of US Census tables from the 1940s. If a worker and manager agree on a wage and the worker completes a page above a stated accuracy threshold, the worker receives the agreed-upon wage. Each worker can complete up to 5 pages of transcription. Each manager is assigned to 3 workers and privately given a per-worker-page budget of $5. Managers are the residual claimant of this budget, after wages are paid, for completed work. We believe our experiment to be informative for understanding the effect of transparency at standard market wages. We selected the budget to approximate wages on the platform in general.[33]

Before interacting, workers are shown a sample transcription page, and are asked to provide an estimate for how long it would take them to transcribe the page. We collect data about the minimum price a worker would accept to transcribe a similar page, and make it clear to workers that this information will not be shared with managers. Similarly to Becker et al. (1964), we make truth-telling a dominant strategy for workers. They are asked to make several selections between receiving $X for completing one transcription page, up to a maximum of 5 pages, or $9 for doing nothing. We vary X and randomly select one choice for 1 in 10 workers and give the worker their reward (either $9 or the opportunity to complete 5 additional pages at $X per page) after their initial assignment has been completed. We calculate a worker's outside option using the lowest wage at which a worker is willing to transcribe a page and the time it takes that worker to transcribe a page.[34]

We next ask workers to place a bid for completing each page of transcription, to be used in the pay negotiation. In some treatments arms this bid is binding and will be either accepted or rejected by managers, in other treatment arms the bid is a non-binding opening

[33]The distribution of worker hourly pay in our experiment has mean $4.14 and median $6.30. Mean and median hourly wages on the platform are $3.18 and $6.19, respectively (Hara et al., 2018). Moreover, the $5 budget matches "expected" wages in our experiment: $5 is the modal, and 46$^{\text{th}}$ percentile bid placed by workers.

[34]Letting x_i represent the lowest wage at which i is willing to complete the transcription, and t_i the time it takes i to transcribe a page, we calculate i's outside option as $\theta_i = \frac{2x_i - 1}{2} - \frac{9}{5} \cdot t_i$.

offer submitted to the manager. Managers then meet with workers to bargain over per-page pay in anonymous online chatrooms. No other communication occurs between participants. We place no restrictions on the way in which participants bargain, only that they indicate in the chatroom a final agreed upon wage. Upon finishing wage negotiations, we survey demographic information including education, gender, and age. Workers who agree to a wage with a manager have 48 hours to complete up to 5 pages of transcription.

Transcription accuracy is calculated using the Levenshtein distance measure (Levenshtein, 1966), defined as the minimum number of single-character edits (substitutions, deletions, or insertions) necessary to change one string into another. Each submitted page with a Levenshtein distance from the original document of fewer than 5% of the total number of characters on the page meets our accuracy threshold. Participants were made aware of this threshold at the onset.

Treatments Our experiment has two main treatment arms. One treatment is the public visibility of wage negotiations. Managers either negotiate wages with each worker in a private chat room, or the manager negotiates over a common, transparent chat room with all workers. The only difference between the private chat room and the transparency chat room is a division between chat forums allowing (or preventing) workers to observe all communication between coworkers and the manager. In both cases, the manager must confirm wages individually with each worker in the chat room. The second treatment arm introduces a constraint on the negotiation options for the manager. In one scenario the manager must accept all bids less than or equal to the budget and reject all others, and in another, the manager can actively bargain with workers. The number of participants assigned to different treatments is displayed in Table A2. Subject characteristics across transparent and non-transparent assignments are similar in terms of age, share female, and education, within both negotiable and non-negotiable treatment arms.

User interface differences across treatments are minimal as the multiple manager chat rooms in the privacy treatment all appear on one page side by side. Pictures of the experimental interface can be found on the authors' websites. We describe an additional treatment arm involving 103 managers and 309 workers in Appendix A.8.

Replicability The experiment is automated and can be replicated using code publicly available on the authors' websites. All interactions occur through a single web interface programmed in oTree (Chen et al., 2016).

Example Transcript Workers often use the wages of others when bargaining in the transparency treatment. Below, we provide a portion of one wage negotiation as an example.

Employee 1: Minimum is 3 for me
Manager: Yes, I will agree on 3, thank you
Employee 3: $4, I got puppy to feed
Manager: Employee 3, I will agree on 4.00
Employee 3: thank you very much
Employee 1: Well that's bull I want 4 now
Manager: Employee 1, I will go up to 4.00 as well for you

IV.B. Empirical Tests and Results

Here, we discuss results from our experimental treatment in which managers can freely negotiate with workers, as this allows us to study the equilibrium responses to transparency. In Section V.A, we discuss results from the treatment in which managers cannot freely negotiate with workers.

Wage and Surplus Equalization: In what follows, we use the subscript j to denote a manager, who is either assigned to a separate chat room or a common chat room (transparent $\in \{0, 1\}$). X_i refers to age, college attainment, and gender of the worker. S_j refers to a set of demographic controls, including the age, college attainment and gender of the manager as well as average characteristics of workers eligible for hire by the manager. We run the following specifications to test the effect of transparency on equity:

$$\{Gini(\text{Wages})_j, Gini(\text{Hourly Pay})_j, Gini(\text{Worker Surplus})_j\} = \beta_0 + \beta_1 \cdot \text{transparent}_j + \psi \cdot S_j + \epsilon_j$$

Wages are significantly equalized in the transparency treatment compared to the privacy treatment. $\hat{\beta}_1 = -0.021$ (s.e.=0.005) while the mean Gini coefficient among non-transparent jobs is 0.023. Of the managers allowed to negotiate wages with workers, 100% pay common wages to all workers in the pay transparent treatment, compared to only 60% of managers in the private pay treatment.

We measure worker surplus as the agreed-to wage less the elicited reservation value. While we observe significant equalization of wages under transparency, we observe amplification in the dispersion of worker surplus, as predicted by Corollary 1.

On average, the Gini coefficient for worker surplus is twice as large in the transparency treatment relative to the non-transparent arm (Table III, Col. 5-6). $\hat{\beta}_1 = 0.084$ (s.e.=0.040) while the mean Gini coefficient for worker surplus among non-transparent jobs is 0.085. Dispersion in worker surplus arises from the fact that workers submit bids for work that

are approximately a fixed mark-up above their outside option. When the employer equalizes wages, she grants more surplus to low outside option workers, exacerbating surplus inequality.

These two findings suggest that transparency closes pay gaps caused by differences in outside options, while simultaneously widening the gap between wage and outside option. If workers' outside options reflect external wages, then equalizing pay may offset disparities in external opportunities. If the outside option reflects cost of effort, rising dispersion in worker surplus reflects a shift of surplus towards those who (are fortunate to) have low effort costs and away from those who have high effort costs. We find evidence that outside options partially reflect cost of effort by eliciting the time it takes to complete a page of transcription. When we convert the piece-rate contracts into an hourly wage, we do not find evidence of hourly-wage equalization (Table III, Col. 3-4). $\hat{\beta}_1 = 0.021$ (s.e.=0.039) while the mean Gini coefficient for hourly wages among non-transparent jobs is 0.14.

Profit and Wages Levels: Theorem 2 predicts that profits will be higher and that wages for employed workers will be lower under full transparency than under full privacy, because employers bargain more aggressively under full transparency. We run the specifications below with outcome variables: wages agreed to, wages agreed to for the subset of workers who completed at least one page of transcription at 95% accuracy, number of transcription pages completed within manager group (productivity), and overall manager profit.[35]

$$\{\log(\text{Wages}|\text{Hired})_{ij}, \log(\text{Wages}|\text{Positive Output})_{ij}\} = \beta_0 + \beta_1 \cdot \text{transparent}_{ij} + \phi \cdot X_i + \psi \cdot S_j + \epsilon_{ij}$$

$$\{sinh^{-1}(\text{Productivity})_j, sinh^{-1}(\text{Profit})_j\} = \beta_0 + \beta_1 \cdot \text{transparent}_j + \psi \cdot S_j + \epsilon_j$$

In Table IV estimates show that manager profits are 27% higher in the transparent group than in the private group (p-value=0.012), and wages are 7.4% lower in the transparency group conditional on reaching an agreement (p-value=0.065).[36] There is no statistical difference between productivity in the two groups albeit estimates are imprecise; the point estimate for productivity is a 13% increase under transparency, with an equally large standard error (p-value=0.309). These findings are consistent with the theoretical intuition that the manager has a credible way to commit to a wage in a common chat room where employees

[35] We use the log transformation of wages as the dependent variable, allowing β_1 to be interpreted as the elasticity of wages with respect to transparency. To accommodate 0 profit outcomes, we use the inverse hyperbolic sine of productivity and manager profit. The transformation arcsinh(x) down-weights treatment effects at small values, is linear for x close to 0 and approximates log(2x) for x greater than 3. For an in-depth discussion about this transformation, see Kline et al. (2019) Appendix D, page 65.

[36] There are at least two reasons why we do not expect these two numbers to be identical. First, transparency increases the hiring rate, mechanically implying higher profits without any effect on wages. Second, if wages are greater than \$2.50+ϵ for some positive ϵ then decreasing the wage by no more than ϵ will increase manager profit proportionally more than it decreases wages, because the wage is greater than half the total budget. Table IV, Col. 4 shows that the average wage in the private treatment is roughly \$4.

can see her turn down bidders demanding more.

Hiring Rate We define the hiring rate as the proportion of workers who agree on a wage with their manager. Theorem 3 states that transparency increases hiring when the value of labor is sufficiently low. The reason is that a worker will demand to high a wage for a job with a low-value employer if wages are not transparent. The wage demand includes a premium above the outside option which is optimal in expectation without knowledge of the employer's value of labor.

This experiment provides a direct test of this prediction under an open bargaining protocol.

$$\text{Hired Worker (Yes=1)}_{ij} = \beta_0 + \beta_1 \cdot \text{transparent}_{ij} + \phi \cdot X_i + \psi \cdot S_j + \epsilon_{ij}$$

We estimate a linear probability model of hiring under transparency versus privacy. We include as controls the age, college attainment, and gender of the worker and manager, X_i and S_j respectively, and allow for correlated errors at the level of the manager.[37] β_1 measures the increased likelihood that a manager hires a worker under pay transparency compared to pay privacy. We show in Col. 1 of Table IV that the hiring rate is 10% higher when negotiations are held in the common chat forum (p-value=0.030). We present additional empirical support for Theorem 3 in Appendices A.6 and A.8.

V. Key Extensions

V.A. Effects of Social Concerns

We incorporate an additional channel, social concerns about relative pay, that is a prominent explanation in the literature for equalized wages in transparent environments. Several papers (Akerlof and Yellen, 1990; Breza et al., 2018; Card et al., 2012) argue that a worker exerts less effort upon learning she is underpaid relative to her peers. A proactive employer who observes the spread of wage information (or observes decreases in effort) may augment the wages of workers who learn they are underpaid in order to avoid low effort provision (Eliaz and Spiegler, 2013). We test the hypothesis that employers optimally equalize wages due to social concerns in this section.

Proactive Employer Model Each worker i makes a wage offer w_i^* at $t = 0$. If she is hired ($w_i^* \leq \bar{w}$) she thereafter only chooses effort $e_{i,t} \in [0, 1]$ at each time t which is the probability of successfully completing her time t duties and receiving her flow wages. All workers have an outside option normalized to 0 and have to pay a linear effort flow cost

[37] Each participant is assigned to only one job in our experiment.

$\theta_i \cdot e_{i,t}$. Each worker i's expected flow payoff is $(w_{i,t} - \theta_i) \cdot e_{i,t}$.[38] At each time t the firm observes whether i receives wage information and can pay any employed worker i a wage $w_{i,t} \geq w_i^*$. We include a morale cost to a worker for learning she is underpaid, modeled as in DellaVigna et al. (2016) and Breza et al. (2018). Workers face a higher cost to effort upon learning they are paid less than any peer.[39] We refer to this as the *proactive employer model.*

Formally, let the morale cost be $m(e_{i,t}, d) \in [0, 1]$ where $d = \bar{w} - w_{i,t}$. We assume $m(\cdot, \cdot)$ is non-decreasing in both arguments and that $m(0, \cdot) = m(\cdot, 0) = 0$. As before, worker i's flow payoff is $(w_{i,t} - \theta_i) \cdot e_{i,t}$ prior to learning about the wages of her peers, so she will put in full effort in equilibrium. Upon seeing the wages of her coworkers and learning $\bar{w}$, the worker's flow payoff becomes $(w_{i,t} - \theta_i) \cdot e_{i,t} - m(e_{i,t}, d)$. Depending on $m(e_{i,t}, d)$, the worker may optimally shirk. The firm will increase the wage of a worker i at time t only if i learns the wages of her coworkers at time t.

We now formally state conditions on the morale function for the proactive employer model to generate the same equilibrium outcome as that of our original model, and in particular, fit our key empirical findings of full pay equalization under transparency, both in TaskRabbit for Business and our field experiment.

PROPOSITION **6.** *Consider any equilibrium satisfying conditions **A1-A3**. The firm always sets $w_{i,t} = \bar{w}$ for every worker i who learns $\bar{w}$ at time t for all Ω if and only if $(w_{i,t} - \theta_i) \cdot e_{i,t} - m(e_{i,t}, d) \leq 0$ for every $e_{i,t} \in (0, 1]$ and every d occurring on equilibrium path.*

Only large and discontinuous morale cost functions result in the same predictions as the bargaining model; unless social concerns reduce worker effort to 0 upon receiving even slightly less than $\bar{w}$, the firm will not equalize the wages of all workers who observe peer wages. Consider the case where $k = 0$, corresponding to our empirical setting. When transparency is low, firms make close to zero profit from their highest paid worker ($v - \bar{w} \approx 0$) so even if a worker drastically reduces her effort, a proactive firm would still prefer to pay her less than $\bar{w}$ unless she quits entirely.[40]

[38] We could replace $e_{i,t}$ with a productivity function $z(e_{i,t})$, where $z(e) > 0$ for all $e > 0$, $z(0) = 0$, and $z(\cdot)$ has a unique global maximizer at 1 without altering the results of this section.

[39] If we layer in the morale specification in this section and give the workers the ability to initiate renegotiations, then workers optimally request $\bar{w}$ upon seeing peer wages. This means that on equilibrium path, workers will never pay the morale cost or put in low effort. Therefore, the presence of a morale cost does not affect the predictions of our bargaining model.

[40] Fanning and Kloosterman (2019) produce experimental evidence that social concerns in dynamic bargaining may lead to a Coasian result of full wage equalization; that is, if with positive probability each worker refuses to work for anything less than full wage equalization due to morale concerns, all other worker types will mimic this behavior, and the firm will (almost) immediately offer wages to workers that are (almost) equal. This type of model would have nearly identical predictions as our bargaining model. However, we note that we do not observe workers rejecting all wages which are bounded away from full equality in our data, as Fanning and Kloosterman (2019) predict.

Empirical evidence of the morale channel versus bargaining channel We do find evidence that relative pay concerns lead to lower effort when a worker learns she is underpaid. To show this, we analyze data from our experimental treatment in which managers are not allowed to renegotiate wages with their workers. Instead, they must accept all initial wage offers less than their given budget, and reject all others. We randomize 36 managers into transparent, public chatrooms, and 34 managers into non-transparent, separate chatrooms. Average output is 26% lower in the transparent no negotiation treatment than the non-transparent no negotiation treatment, meaning that when inequality becomes known to workers and there is no avenue for wage equalization, effort declines. The magnitude of productivity loss is in line with a field experiment in Breza et al. (2018) who find a 22% reduction in output among workers who learn each other's salaries. Importantly, the extent of this drop in productivity depends on the level of inequality between workers who discover each others' pay. In Figure II we show that output is decreasing as a function of the difference between a worker's pay and that of her highest paid peer under transparency, compared to under pay privacy. Breza et al. (2018) also present a similar trend of a smooth decrease in output as inequality rises.

The implication of these empirical findings, in light of Proposition 6, is that a profit maximizing employer would be unlikely to proactively fully equalize wages of workers who learn the pay of peers unless there is a renegotiation channel. Without renegotiations, an employer anticipating the productivity consequences of low morale, would be more likely to equalize wages when inequality among worker pay is high at the time the job begins. Additionally the employer would also choose to only partially close the gap (rather than equalizing wages) in some cases.

Empirical support for **SF1** and **SF2** further supports bargaining, and not morale, as the main driver of wage equalization. We do not find a correlation between the extent of inequality and the likelihood of a raise in TaskRabbit for business (Col 1-3, Table I). We also find pay is equalized whether the gap is large or small in our field experiment and essentially equalized in co-located tasks on TaskRabbit for Business data when there is renegotiation (Col. 4-6, Table I).

In light of these empirical findings and Proposition 6, we conclude that re-bargaining is likely an important mechanism equalizing wages in the transparent pay environments, even in the presence of morale concerns.

V.B. Heterogeneous Worker Qualities

Until now we have assumed that all workers are equally productive. This is based on our empirical setting in which jobs consist of simple and standardized tasks. Here we discuss our findings in contexts where there may be significant heterogeneity in worker productivities.

First, consider a generalization of our base model in which there are different classes of workers, and where these differences are known. Each worker $i \in I$ has a publicly observable type $\tau \in \mathcal{T}$ where $\mathcal{T}$ is some countable set, each containing a positive measure of workers. Let $v_\tau \sim F_\tau[0,1]$ be the productivity of type τ workers, which is known only to the firm. Each worker i of type τ also has a private outside option $\theta_i \overset{iid}{\sim} G_\tau[0,1]$, which is the flow payment i receives when not matched to the firm. Our base model is a special case in which $|\mathcal{T}| = 1$, that is, all workers are equally productive with productivity $v \sim F[0,1]$ and outside option distribution G. As before, each worker i of type τ makes an initial wage offer $w^*_{i,\tau}$, and then an additional wage offer after observing the wages of peers. The firm picks a maximum wage $\bar{w}_\tau(v_\tau)$ for each type τ.

The results of our paper go through within type τ, that is, each τ forms a different market. On equilibrium path, the firm picks the maximum wage for type τ workers $\bar{w}_\tau(v_\tau)$ as in the base model given distributions F_τ and G_τ, and each worker i of type τ picks an initial offer $w^*_{i,\tau}$ as in the base model given distributions F_τ and G_τ.

We therefore focus on the more interesting case in which workers are differentially productive, but do not know their own productivity type. Suppose there are two types of workers, with productivities v and V, respectively. v and V are drawn independently from the same distribution F. Each worker is equally likely to have productivity type v or V. The firm knows each worker's productivity type, but workers do not. To highlight mechanisms at play, we study the extreme case in which outside options are distributed independently of productivity so that workers do not receive a signal of their relative productivities.

Under full privacy, the equilibrium outcome mirrors that of the base model. Therefore, firm profits, the expected hiring rate, and wage dispersion are the same as before.

For tractability, we consider only the effects of full transparency. At $t = 0$, each worker will observe $\bar{W}_v$ and $\bar{W}_V$. Without loss of generality, we assume that $v < V$ so that $\bar{W}_v \leq \bar{W}_V$. We note that if $k > 0$, this model is again not meaningfully different than before. All workers i will initially offer $w^*_i = \bar{W}_v$. Based on this, each i will perfectly learn her productivity type: if her initial wage is equal to $\bar{W}_v$, she knows she knows she has productivity type v and will not renegotiate, however, if her initial wage is greater than $\bar{W}_v$ then she knows she has productivity type V and will (almost, see Footnote 16) immediately renegotiate to wage $\bar{w}_V$.

Therefore, suppose $k = 0$. If $\bar{W}_V < \theta_i$, then worker i will remain unemployed. Otherwise, she will offer $\bar{W}_v$ (and be employed with probability 1) if $\bar{W}_v > \frac{1}{2}\bar{W}_V + \frac{1}{2}\theta_i$ and she will offer $\bar{W}_V$ if $\bar{W}_v \leq \frac{1}{2}\bar{W}_V + \frac{1}{2}\theta_i$.

When a worker demands $\bar{W}_V$ the firm will reject her offer with probability $\frac{1}{2}$, which clearly reduces the hiring rate and firm profits compared to the baseline model where there is no uncertainty about worker productivity. On the other hand, low outside option, productivity V workers will offer $\bar{W}_v$, meaning that the firm is able to hire some high productivity workers

at low wages, increasing profits. We show that, because of this latter effect, the firm sets $\bar{W}_v$ higher than it would have for the same v with known worker productivities.

PROPOSITION **7.** *In equilibrium under full transparency, $\bar{W}_V = \bar{w}(V)$ and $\bar{W}_v > \bar{w}(v)$, where $\bar{w}(\cdot)$ is the maximum wage the firm sets in the baseline model where productivity differences are observable.*

In this setting, transparency leads to wage compression as opposed to complete wage equalization. All employed, low-productivity workers earn $\bar{W}_v$ as the firm rejects all such workers who demand more. Employed, high-productivity workers earn either $\bar{W}_v$ or $\bar{W}_V$. Since $\bar{W}_v > \bar{w}(v)$, and $\bar{W}_V = \bar{w}(V)$, the gap in pay between low- and high-productivity workers is smaller than in the base model. Interestingly, we show that the firm may set $\bar{W}_v > v$ when v is sufficiently small, incurring a loss on low-productivity workers!

Because $\bar{W}_v > \bar{w}(v)$, more low-productivity workers are hired than if productivity differences were observable. This completely offsets the reduction in the hiring rate caused by high outside option, low type workers requesting $\bar{W}_V$. The fact that the firm is able to secure low outside option, high productivity workers at wage $\bar{W}_v$ also offsets the profit loss caused by missing out on certain low quality workers.

PROPOSITION **8.** *For any values v and V, firm profit and the hiring rate are the same as in the baseline model with observable productivity differences.*

VI. CONCLUSION

Although pay transparency has been in the political and popular spotlights, its effect on wages, hiring, and profits are not well understood. Our theoretical and empirical analyses of equilibrium wages, hiring rate, and profits under greater transparency reveals consequences that are counterintuitive and economically significant in a market for low-skill tasks.

Pay transparency improves the bargaining power of the employer. We observe this in both observational panel data and a field experiment, leading to wages that are 7-8% lower than under pay privacy. Employer profits are even larger as transparency raises the hiring rate. Increasing transparency increases hiring at equilibrium wages by approximately 10% in our experiment. We show that transparency equalizes wages while increasing dispersion in worker surplus by shifting surplus toward low outside option workers. We present relative worker surplus as an additional measure of fairness, which should be interpreted in light of what outside options represent. If outside options represent cost of effort, then more equal wages resulting from transparency shift worker surplus toward workers most capable of completing the job. If outside options represent previous wages, then transparency lessens disparities in opportunity in other labor markets.

Empirical studies from other settings offer corroborating evidence of the central role that bargaining plays in mediating wages. Without an equilibrium response through bargaining, we would expect wages to rise after transparency is introduced, as transparency's direct effect allows low-wage workers to negotiate higher pay. Baker et al. (2019) and Bennedsen et al. (2019) show that wage growth falls after transparency mandates. These results point to the first-order effect of transparency on bargaining power.

Our finding that pay transparency is profitable for employers appears at odds with conventional wisdom that transparency is rare and that most firms rally against it (Hegewisch et al., 2011). One reason for this is the narrow framing of transparency in the popular debate. In this paper, we show that higher transparency is a substitute for increased firm bargaining power; the equilibrium outcome under full transparency is equivalent to the firm posting an optimal wage. Redefining transparency to include not only pay discussions on the job, but also the commitment of employers to pay a posted wage, implies that there is actually a high degree of transparency in the economy, and employers are actively selecting it. Hall and Krueger (2012) show that nearly three-fourths of workers face a high degree of transparency.[41] This share is higher among jobs with a high degree of standardization, such as entry-level jobs and occupation-certified roles, and these are precisely the jobs for which learning the pay of a coworker is likely most informative in wage renegotiations.[42] Other studies place transparency within specific occupations at even higher levels.[43]

Employers who favor pay transparency during initial wage negotiations need not favor it at *all times*. Transparently posting a wage during the initial hiring process strengthens the employer's bargaining position, but renegotiations following on-the-job wage gossip directly lower profits. Employers may also be concerned about other consequences of pay transparency when there are wage disparities across *different* jobs within the firm, such as unfavorable press, or worker morale costs. As a result, the public perception that employers discourage transparency could be in large part driven by their desire to prevent gossip after wages have been initially negotiated.

While the average worker's wage falls with greater transparency, partial transparency increases the pay of the lowest outside option employees (Section II.C). These workers (or their advocates) may explain the public perception that employees demand greater transparency. Employees may also have additional reasons to seek transparency that have been outside the scope of our study. For example, knowledge of a firm's pay structure can provide

[41] In their sample, two-thirds of employees faced posted wages. Of the remaining third who did not face a posted wage, 15% knew what their salary would be when they entered negotiations.

[42] Similarly, Caldwell and Harmon (2019) find that 70% of employees face posted wages and that higher-skilled and more specialized workers are much more likely to increase their wages through renegotiations.

[43] Niederle et al. (2006) consider entry-level jobs for gastroenterologists and find that 94% of employers pay common wages and "offers are not adjusted in response to outside offers and terms are not negotiable."

information about a worker's potential career trajectory (Cullen and Perez-Truglia, 2019) and guide them in making informed career investment decisions (Faleye et al., 2013).

Our model can be adapted to evaluate recent salary history bans, which prohibit employers from asking job applicants their past wages. Consider a worker who bargains with different employers at various times during her life. Verifiable past salary information strengthens the bargaining power of worker by allowing workers to credibly commit to rejecting low wage offers, leading to higher wages. Employers, recognizing this, will respond by benchmarking their wage offers on the (lowest) previous salary the worker has accepted. Therefore, public salary history has a similar effect for the worker as full transparency does for the firm. In equilibrium, workers should voluntarily reveal previous wages due to unraveling: workers with the highest past wages will reveal this information during a negotiation to secure a higher wage, which incentivizes lower outside option workers to do the same. We believe that testing the efficacy of these laws empirically is an important avenue for further research.

References

George A. Akerlof and Janet L. Yellen. The Fair Wage-Effort Hypothesis and Unemployment. *The Quarterly Journal of Economics*, 105(2):255–283, 1990.

Linda Babcock and Sara Laschever. *Women Don't Ask: Negotiation and the Gender Divide*. Princeton University Press, 2003.

Kyle Bagwell and Asher Wolinsky. Game theory and industrial organization. *Handbook of Game Theory with Economic Applications*, 3:1851–1895, 2002.

Michael Baker, Yosh Halberstam, Kory Kroft, Alexandre Mas, and Derek Messacar. Pay Transparency and the Gender Gap. *NBER Working Paper No. 25834*, 2019.

Gordon M. Becker, Morris H. DeGroot, and Jacob Marschak. Measuring Utility by a Single-Response Sequential Method. *Behavioral Sciences*, 9(3):226–232, 1964.

Morten Bennedsen, Elena Simintzi, Margarita Tsoutsoura, and Daniel Wolfenzon. Do Firms Respond to Gender Pay Gap Transparency? *NBER Working Paper No. 25435*, 2019.

Dirk Bergemann and Johannes Hörner. Should First-Price Auctions Be Transparent? *American Economic Journal: Microeconomics*, 10(3):177–218, 2018.

Francis Bloch and Harl Ryder. Two-Sided Search, Marriages, and Matchmakers. *International Economic Review*, 41(1):93–115, 2000.

Anat Bracha, Uri Gneezy, and George Loewenstein. Relative pay and labor supply. *Journal of Labor Economics*, 33(2):297–315, 2015.

Giulia Brancaccio, Dan Li, and Norman Schürhoff. Learning by Trading: The Case of the U.S. Market for Municipal Bonds. *Mimeo*, 2020.

Vera Brenčič. Wage posting: evidence from job ads. *Canadian Journal of Economics/Revue canadienne d'èconomique*, 45(4):1529–1559, 2012.

Hanna Brenzel, Hermann Gartner, and Claus Schnabel. Wage bargaining or wage posting? evidence from the employers' side. *Labour Economics*, 29:41–48, 2014.

Emily Breza, Supreet Kaur, and Yogita Shamdasani. The Morale Effects of Pay Inequality. *Quarterly Journal of Economics*, 133(2):611–663, 2018.

Kenneth Burdett and Melvyn G. Coles. Long-Term Partnership Formation: Marriage and Employment. *The Economic Journal*, 109(456):F307–F334, 1999.

David A. Butz. Durable-Good Monopoly and Best-Price Provisions. *American Economic Review*, 80(5): 1062–1076, 1990.

Sydnee Caldwell and Nikolaj Harmon. Outside Options, Bargaining, and Wages: Evidence from Coworker Networks. *Mimeo*, 2019.

David Card, Alexandre Mas, Enrico Moretti, and Emmanuel Saez. Inequality at Work: The Effect of Peer Salaries on Job Satisfaction. *American Economic Review*, 102(6):2981–3003, 2012.

Hector Chade. Matching with Noise and the Acceptance Curse. *Journal of Economic Theory*, 129(1):81–113, 2006.

Kalyan Chatterjee and William Samuelson. Bargaining under Incomplete Information. *Operations Research*, 31(5):835–851, 1983.

Daniel L. Chen, Martin Schonger, and Chris Wickens. oTree-An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97, 2016.

Alain Cohn, Ernst Fehr, Benedikt Herrmann, and Frédéric Schneider. Social Comparison and Effort Provision: Evidence from a Field Experiment. *Journal of the European Economic Association*, 12:877–898, 2014.

Thomas E. Cooper and Timothy L. Fries. The most-favored-nation pricing policy and negotiated prices. *International Journal of Industrial Organization*, 9:209–223, 1991.

Zoë Cullen and Ricardo Perez-Truglia. How Much Does Your Boss Make? The Effects of Salary Comparisons. *NBER Working Paper No. 24841*, 2019.

Stefano DellaVigna, John A. List, Ulrike Malmendier, and Gautam Rao. Estimating Social Preferences and Gift Exchange at Work. *NBER Working Paper 22043*, 2016.

Arindrajit Dube, Laura Giuliano, and Jonathan Leonard. Fairness and Frictions: The Impact of Unequal Raises on Quit Behavior. *American Economic Review*, 109(2):620–663, 2019.

Alex Edmans, Xavier Gabaix, Tomasz Sadzik, and Yuliy Sannikov. Dynamic CEO Compensation. *Journal of Finance*, 67(5):1603–1647, 2012.

Liran Einav, Chiara Farronato, Jonathan D. Levin, and Neel Sundaresan. Auctions versus Posted Prices in Online Markets. *Journal of Political Economy*, 126(1):178–215, 2018.

Kfir Eliaz and Ran Spiegler. Reference Dependence and Labor-Market Fluctuations. *NBER Macroeconomics Annual*, 28:159–200, 2013.

Tore Ellingsen and Åsa Rosén. Fixed or Flexible? Wage-Setting in Search Equilibrium. *Economica*, 70(278): 233–250, 2003.

European Commission. Pay transparency in the EU: A legal analysis of the situation in the EU Member States, Iceland, Liechtenstein and Norway. Technical report.

Olubunmi Faleye, Ebru Reis, and Anand Venkateswaran. The determinants and effects of CEO-employee pay ratios. *Journal of Banking & Finance*, 37(8):3258–3272, 2013.

Jack Fanning and Andrew Kloosterman. A Simple Experimental Test of the Coase Conjecture: Fairness in Dynamic Bargaining. *Mimeo*, 2019.

Chiara Farronato, Andrey Fradkin, Bradley J Larsen, and Erik Brynjolfsson. Consumer Protection in an Online World: An Analysis of Occupational Licensing. *NBER Working Paper 26601*, 2019.

Michael Faulkender and Jun Yang. Is Disclosure an Effective Cleansing Mechanism? The Dynamics of Compensation Peer Benchmarking. *Review of Financial Studies*, 26(3):806–839, 2013.

Drew Fudenberg and David K. Levine. Maintaining a Reputation when Strategies are Imperfectly Observed. *The Review of Economic Studies*, 59(3):561–579, 1992.

Xavier Gabaix. Power Laws in Economics and Finance. *Annual Review of Economics*, 1(1):255–294, 2009.

Xavier Gabaix. Power Laws in Economics: An Introduction. *Journal of Economic Perspectives*, 30(1): 185–206, 2016.

Simon Gächter and Christian Thöni. Social comparison and performance: Experimental evidence on the fair wage-effort hypothesis. *Journal of Economic Behavior & Organization*, 76(3):531–543, 2010.

Avi Goldfarb and Catherine Tucker. Technology, Age, and Shifting Privacy Concerns. *Mimeo*, 2011.

Wulong Gu and Peter Kuhn. Centralization and Strikes. *Labour Economics*, 5(3):243–265, 1998.

Robert Hall and Alan Krueger. Evidence on the Incidence of Wage Posting, Wage Bargaining, and On-the-Job Search. *American Economic Journal: Macroeconomics*, 4(4):56–67, 2012.

Kotaro Hara, Abigail Adams, Kristy Milland, Saiph Savage, Chris Callison-Burch, and Jeffrey P. Bigham. A Data-Driven Analysis of Workers' Earnings on Amazon Mechanical Turk. *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 2018.

Ariane Hegewisch, Claudia Williams, and Robert Drago. Pay Secrecy and Wage Discrimination. *Institute for Women's Policy Research*, 2011.

Johannes Hörner and Nicolas Vieille. Public vs. Private Offers in the Market for Lemons. *Econometrica*, 77 (1):29–69, 2009.

International Labour Organization. Global wage report 2018/19: What lies behind gender pay gaps. Technical report.

Kenneth L. Judd. The law of large numbers with a continuum of IID random variables. *Journal of Economic Theory*, 35(1):19–25, 1985.

Ayça Kaya and Qingmin Liu. Transparency and price formation. *Theoretical Economics*, 10(2):341–383, 2015.

Marlene Kim. Pay Secrecy and the Gender Wage Gap in the United States. *Industrial Relations*, 54(4): 648–667, 2015.

Patrick Kline, Neviana Petkova, Heidi Williams, and Owen Zidar. Who Profits from Patents? Rent-sharing at Innovative Firms. *Quarterly Journal of Economics*, 2019.

David M. Kreps and Robert Wilson. Reputation and imperfect information. *Journal of Economic Theory*, 27(2):253–279, 1982.

Peter Kuhn and Wulong Gu. A Theory of Holdouts in Wage Bargaining. *American Economic Review*, 88 (3):428–449, 1998.

Peter Kuhn and Wulong Gu. Learning in Sequential Wage Negotiations: Theory and Evidence. *Journal of Labor Economics*, 17(1):109–140, 1999.

Bradley Larsen. The Efficiency of Real-World Bargaining: Evidence from Wholesale Used-Auto Auctions. *Review of Economic Studies, forthcoming*, 2020.

Bradley J. Larsen and Anthony Lee Zhang. A Mechanism Design Approach to Identification and Estimation. *NBER Working Paper 24837*, 2018.

Edward P. Lazear. Retail Pricing and Clearance Sales. *American Economic Review*, 76(1):14–32, 1986.

Wolfgang Leininger, Peter B. Linhart, and Roy Radner. Equilibria of the Sealed-Bid Mechanism for Bargaining with Incomplete Information. *Journal of Economic Theory*, 48:63–106, 1989.

Vladimir I. Levenshtein. Binary Codes Capable of Correcting Deletions, Insertions, and Reversals. *Soviet Physics Doklady*, 10(8):707–710, 1966.

Alan Manning. *Monopsony in Motion: Imperfect Competition in Labor Markets*. Princeton University Press, 2005.

Alexandre Mas. Does Disclosure affect CEO Pay Setting? Evidence from the Passage of the 1934 Securities and Exchange Act. *Mimeo*, 2016.

Alexandre Mas. Does Transparency Lead to Pay Compression? *Journal of Political Economy*, 125(5): 1683–1721, 2017.

Guido Menzio. A Theory of Partially Directed Search. *Journal of Political Economy*, 115(5):748–769, 2007.

Claudio Michelacci and Javier Suarez. Incomplete Wage Posting. *Journal of Political Economy*, 114(6): 1098–1123, 2006.

Paul Milgrom. Good News and Bad News: Representation Theorems and Applications. *The Bell Journal of Economics*, 12(2):380–391, 1981.

Paul Milgrom and John Roberts. Predation, reputation and entry deterrence. *Journal of Economic Theory*, 27(2):280–312, 1982.

Muriel Niederle, Deborah D. Proctor, and Alvin E. Roth. What Will Be Needed for the New Gastroenterology Fellowship Match to Succeed? *Gastroenterology*, 130:218–224, 2006.

Chris Nosko and Steven Tadelis. The limits of reputation in platform markets: An empirical analysis and field experiment. *NBER Working Paper No. 20830*, 2015.

Tomasz Obloj and Todd Zenger. The influence of pay transparency on inequity, inequality, and the performance-basis of pay in organizations. *Mimeo*, 2019.

Kathleen Pender. New law bans California employers from asking applicants their prior salary. *https://www.sfgate.com/business/networth/article/New-law-bans-California-employers-from-asking-12274431.php*, 2017.

Ricardo Perez-Truglia. The effects of income transparency on well-being: Evidence from a natural experiment. 2019.

Macon Phillips. Now Comes Lilly Ledbetter. *https://obamawhitehouse.archives.gov/blog/2009/01/25/now-comes-lilly-ledbetter*, 2009.

Roy Radner and Andrew Schotter. The Sealed-Bid Mechanism: An Experimental Study. *Journal of Economic Theory*, 48:179–220, 1989.

Barbara Robles and Marysol McGee. Exploring Online and Offline Informal Work: Findings from the Enterprising and Informal Work Activities (EIWA) Survey. *Finance and Economics Discussion Series Divisions of Research & Statistics and Monetary Affairs Federal Reserve Board, Washington, D.C.*, 2016.

Alvin E. Roth and Xiaolin Xing. Jumping the Gun: Imperfections and Institutions Related to the Timing of Market Transactions. *American Economic Review*, 84(4):992–1044, 1994.

Mark A. Satterthwaite and Steven R. Williams. Bilateral Trade with the Sealed Bid k-Double Auction: Existence and Efficiency. *Journal of Economic Theory*, 48:107–133, 1989.

Cornelius Schmidt. Does Transparency Increase Executive Compensation? *Mimeo*, 2012.

Fiona Scott Morton. The Strategic Response by Pharmaceutical Firms to the Medicaid Most- Favored-Customer Rules. *The RAND Journal of Economics*, 28(2):269–290, 1997a.

Fiona Scott Morton. The Interaction between a Most-Favored-Customer Clause and Price Dispersion: An Empirical Examination of the Medicaid Rebate Rules of 1990. *Journal of Economics & Management Strategy*, 6(1):151–174, 1997b.

Leo K. Simon and Maxwell B. Stinchcombe. Extensive Form Games in Continuous Time: Pure Strategies. *Econometrica*, 57(5):1171–1214, 1989.

Gary R. Siniscalco, Erin M. Connell, and Chad Smith. State Pay Equity Laws: Where a few go, many may follow. *Mimeo*, 2017.

Yeneng Sun. The exact law of large numbers via Fubini extension and characterization of insurable risks. *Journal of Economic Theory*, 126(1):31–69, 2006.

Jean Tirole. From Bottom of the Barrel to Cream of the Crop: Sequential Screening with Positive Selection. *Econometrica*, 84(4):1291–1343, 2016.

Donald M. Topkis. *Supermodularity and Complementarity*. Princeton University Press, 1998.

Steven R. Williams. Efficient Performance in Two Agent Bargaining. *Journal of Economic Theory*, 41: 154–172, 1987.

Ahu Yildirmaz, Mita Goldar, and Sara Klein. Illuminating the Shadow Workforce: Insights Into the Gig Workforce in Businesses. *ADP Research Institute*, 2020.

TABLE I: Bonuses Among Co-located Workers, TaskRabbit for Business

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.	Any Raise (Yes = 1)			Raise (% Above Bid) | Any Raise		
$\hat{\beta}_1$ Amt. under top bid (%)	0.0105	0.0257	0.0294	0.883***	0.958***	0.947***
	[0.0181]	[0.0192]	[0.0193]	[0.105]	[0.0740]	[0.0791]
Years experience		0.0191	0.0436**		-0.126	0.0334
		[0.0369]	[0.0215]		[0.112]	[0.0527]
Effective percent positive overall		0.000788	-0.00100		0.0184	0.0326
		[0.0108]	[0.00570]		[0.0278]	[0.0243]
Effective percent positive in cat.		0.0349	-0.0120		0.0482	-0.0632
		[0.0312]	[0.0240]		[0.0703]	[0.100]
No. reviews		-0.0135	-0.00474		0.0453	0.0215
		[0.00899]	[0.00609]		[0.0302]	[0.0176]
No. reviews cat.		0.0792	-0.00761		-0.375**	-0.0614
		[0.0546]	[0.0390]		[0.170]	[0.144]
Mean rating		-0.000217	0.0130		-0.0306	0.0145
		[0.0169]	[0.0103]		[0.0479]	[0.0399]
Mean rating in category		0.000598	-0.00112		0.0271	-0.0216
		[0.0131]	[0.00863]		[0.0360]	[0.0262]
No. workers (log)		-0.128***			-0.00728	
		[0.0293]			[0.197]	
Mean bid (log)		0.0138			-0.171	
		[0.0304]			[0.176]	
Constant	0.215***	0.234*	0.206***	0.267***	0.934	0.333
	[0.0256]	[0.129]	[0.0467]	[0.0625]	[0.761]	[0.201]
>1 hour overlap		✓	✓		✓	✓
Job-specific FE			✓			✓
P-value Test: $H_0 : \beta_1 = 1$				0.267	0.570	0.509
Mean Dep. Var.	0.22	0.22	0.22	0.43	0.44	0.46
SE Dep. Var.	0.41	0.41	0.41	0.38	0.41	0.42
Observations	930	708	708	203	154	116
Clusters Jobs	390	299	299	123	92	54
Clusters Workers	627	481	481	165	130	96
Clusters Employers	293	237	237	117	88	52
R^2	0.001	0.047	0.839	0.665	0.730	0.963

Notes: Each model is estimated by OLS. Col. 1-3 are linear probability models. An observation is an accepted worker-bid for jobs with co-located workers. Our main sample is restricted to job categories with at least 20% of each of separated and co-located multi-worker jobs, however results are robust to looking across all categories with co-located jobs. The dependent variable equals one if the particular worker earns more than their agreed to bid, and 0 otherwise. Following theoretical predictions SF1 and SF2, Col. 4-6 are restricted to those workers that receive a raise (final pay higher than their agreed-to bid). The dependent variable is the size of the raise, as percent above bid. The primary explanatory variable, amount below maximum bid, is equal to $(bid^{max} - bid^i)/bid^i$ for person i. Reviews are in units of 1000. ">1 hour overlap" is equal to 1 if the survey response to the question "How many hours is it necessary for workers to overlap in the same place at the same time in order to complete this job?" is greater than 1 hour on average. Singleton observations are dropped when job fixed effects are added. Standard errors are three-way clustered at the level of the job, worker, and employer.

TABLE II: Lower Wages Under Transparency, TaskRabbit for Business

	(1)	(2)	(3)
Dep. Var.	All Bids (log $)	Winners' Pay (log $)	Winners' Hourly Pay (log hourly wages)
Transparent	-0.0959***	-0.136***	-0.0778***
(Job mentions price)	[0.0136]	[0.0142]	[0.0281]
Exp. on platform (Days)	0.378**	0.493***	1.048***
	[0.175]	[0.123]	[0.393]
No. ratings in category	-0.991***	-0.512***	1.921***
	[0.0859]	[0.172]	[0.308]
No. ratings overall	1.656***	1.694***	0.361
	[0.0546]	[0.123]	[0.400]
Mean rating in category	0.0476***	0.0792***	0.0230
	[0.00733]	[0.0199]	[0.0363]
Mean rating overall	-0.00385	-0.0168	-0.0591
	[0.0114]	[0.0330]	[0.0796]
Job Category FE	✓	✓	✓
Worker FE	✓	✓	✓
Mean Dep. Var.	3.37	3.69	3.02
SE Dep. Var.	0.946	0.854	0.624
Observations	>100k	>100k	>20k
R^2	0.277	0.381	0.607

Notes: Each model is estimated by OLS. An observation is a worker-bid on TaskRabbit for Business. The dependent variable is the log bid in Col. 1, and log final pay in Col. 2 and 3. In Col. 3 we restrict our attention to the sample of jobs that solicit hourly wage bids rather than piece rate. The independent variable, transparent, is an indicator equal to one if there is any mention of price in the job post. Only job posts that accept private bids are included in these regressions. Platform tenure is measured in days. Performance covariates include the square of all ratings covariates. Two-way clustered standard errors at the worker and employer level are in square brackets. We do not reveal observation counts for aggregate activity on the platform at the request of TaskRabbit for Business.

TABLE III: Dispersion in Wages and Worker Surplus, Experiment

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.	Wages (Gini)		Hourly Wage (Gini)		Worker Surplus (Gini)	
Transparent	-0.0237***	-0.0211***	0.0395	0.0217	0.0784*	0.0849**
(Public chat)	[0.00721]	[0.00570]	[0.0383]	[0.0393]	[0.0400]	[0.0406]
Manager Some College		0.0335**		0.158		0.0730
		[0.0167]		[0.0965]		[0.0953]
Manager Age		-0.000522		0.00167		0.000961
		[0.000460]		[0.00230]		[0.00193]
Manager Female		-0.000686		-0.0492		0.0488
		[0.00892]		[0.0459]		[0.0488]
Age (Worker Avg.)		-0.000275		-0.00265		-0.00380
		[0.000474]		[0.00328]		[0.00279]
Some College (Worker Avg.)		-0.0963**		-0.323		-0.0814
		[0.0477]		[0.266]		[0.247]
Female (Worker Avg.)		-0.0243		0.0966		-0.150*
		[0.0190]		[0.0826]		[0.0867]
Constant	0.0237***	0.128***	0.147***	0.321	0.0882***	0.259
	[0.00721]	[0.0419]	[0.0222]	[0.200]	[0.0222]	[0.214]
Mean Dep. Var.	0.016	0.016	0.161	0.161	0.115	0.115
SE Dep. Var.	0.041	0.041	0.151	0.151	0.159	0.159
Observations	70	70	70	70	70	70
R^2	0.077	0.211	0.016	0.094	0.056	0.154

Notes: Each model is estimated by OLS. An observation is a manager. Data is from the sample assigned to treatment groups where workers were allowed to negotiate with the manager. For participants who opted not to report certain demographics, we impute the missing value using the average of the non-missing values and include an indicator variable equal to 1 if the value has been imputed. The dependent variable in Col. 1-2 is the dispersion in final wages agreed-to between the worker and manager. The dependent variable in Col. 3-4 is dispersion in expected hourly wage, defined as the per page wage agreed to divided by the estimated time to complete each page. The dependent variable in Col. 5-6 is dispersion in worker surplus, defined as the difference between the per page rate agreed to and the outside option. Since the Gini coefficient is only defined if more than one worker is hired by the manager, we exclude from this table employers who hire one or zero workers. We also exclude workers with inconsistent responses to the Becker-DeGroot-Marschak (BDM) outside option elicitation method. Covariates with "worker avg." refer to the mean demographic characteristic for all workers assigned to a particular manager. Robust standard errors are displayed in square brackets.

TABLE IV: Higher Profits and Hiring Under Transparency, Experiment

	(1)	(2)	(3)	(4)	(5)
Dep. Var.	Worker Outcomes			Manager Outcomes	
	Hiring (yes =1)	Wages|Hire (log $)	Wages|+ Output (log $)	Total Output (arc sinh)	Profit (arc sinh)
Transparent	0.106**	-0.0740*	-0.130***	0.136	0.271**
(Public chat)	[0.0485]	[0.0396]	[0.0476]	[0.133]	[0.106]
Worker Some College	-0.121	0.123	0.242**		
	[0.194]	[0.104]	[0.106]		
Worker Age	-0.00518	0.00274**	-0.00366*		
	[0.00338]	[0.00131]	[0.00212]		
Worker Female	0.0470	-0.0718	0.0513		
	[0.0736]	[0.0466]	[0.0550]		
Manager Some College	0.0896	0.0745	0.281**	-0.298	-0.242
	[0.106]	[0.1000]	[0.123]	[0.270]	[0.249]
Manager Age	0.00507*	0.00179	-0.00275	0.0186***	0.0154***
	[0.00292]	[0.00201]	[0.00240]	[0.00701]	[0.00556]
Manager Female	-0.106	-0.0483	0.0142	-0.212	-0.207*
	[0.0672]	[0.0522]	[0.0581]	[0.163]	[0.124]
Age (Worker Avg.)	-0.00537	-0.00370	0.00167	-0.0141	-0.00759
	[0.00437]	[0.00319]	[0.00414]	[0.0107]	[0.00623]
Some College (Worker Avg.)	-0.169	-0.225	-0.574	-0.206	-0.184
	[0.209]	[0.252]	[0.349]	[0.691]	[0.489]
Female (Worker Avg.)	-0.00532	0.271***	0.0486	0.160	0.0179
	[0.103]	[0.0918]	[0.108]	[0.259]	[0.179]
Constant	0.866***	1.397***	1.681***	0.901	0.402
	[0.220]	[0.180]	[0.230]	[0.691]	[0.554]
Mean Dep. Var.	0.525	4.276	4.377	0.603	0.271
Observations	303	159	57	101	101
Clusters	101	84	49	101	101
R^2	0.273	0.135	0.370	0.117	0.212

Notes: Col. 1-5 are estimated using ordinary least squares. Data is from the sample assigned to treatment groups where workers were allowed to negotiate with the manager. For participants who opted not to report certain demographics, we impute the missing value using the average of the non-missing values and include an indicator variable equal to 1 if the value has been imputed. An observation is a worker (Col. 1-3) or manager (Col. 4-5). The dependent variables (moving left to right) are hiring (equal to 1 if the worker and manager agree on a wage and 0 otherwise), log wage agreed to with the manager, log wage agreed to with the manager conditional on receiving a payout by submitting at least one page at 95% accuracy, inverse hyperbolic sine of total pages completed by workers assigned to a manger, and inverse hyperbolic sine of profits a manager earns. We use inverse hyperbolic sine transformation to accommodate 0 outcomes. This transformation down-weights treatment effects at small values, is linear for x close to 0 and approximates log(2x) for x greater than 3 (for more details see Kline et al. (2019) Appendix D, page 65). Covariates with "worker avg." refer to the mean demographic characteristic for all workers assigned to a particular manager. Clustered standard errors at the manager level are in square brackets.

TABLE V: Share of Jobs with Posted Price, TaskRabbit

	(1)	(2)	(3)	(4)
Dep. Var.	Proportion of Jobs with Posted Price			
Market age (months)	0.0110***	0.0133**	0.0120***	0.0140**
	[0.0000196]	[0.00500]	[0.00102]	[0.00511]
Number of posts per month			0.00000761	0.00000754
			[0.00000731]	[0.00000728]
City FE, Month FE	✓	✓	✓	✓
Category job-share control		✓		✓
Observations	417	417	417	417
P-Value-Wild-Bootstrap	0.132	0.213	0.113	0.206
Clusters	19	19	19	19
R^2	0.668	0.668	0.671	0.671

Notes: Each model is estimated by OLS. An observation is a city-month in TaskRabbit. The dependent variable is the proportion of tasks that use the transparent posted price scheme. Market size is measured as the number of job posts. Category share is measured by the share of job posts that fall into the top three categories at the month-city level. Standard errors, and significance stars, from OLS are clustered at the city level and displayed in square brackets below the point estimates. To address the small number of clusters, we also report p-values after wild cluster bootstrapping, drawing from the six-point Webb distribution, 1 million bootstrap samples.

Figure II: Productivity consequences of transparency when pay is non-negotiable, Experiment

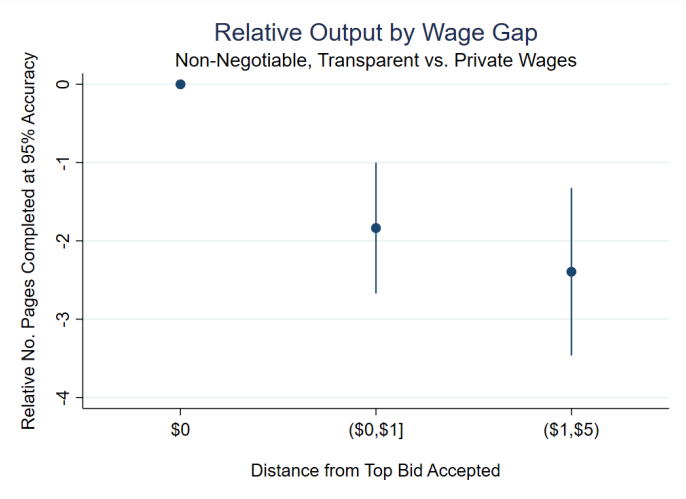

Notes: We plot OLS coefficients and robust standard errors from regressing the number of pages completed on the interaction between co-worker bid differences (amt. below highest bid) and an indicator equal to one if co-workers are in a transparent common chat, 0 otherwise. We only include the treatment group that was *not* allowed to renegotiate, so initial bids were equal to the final pay (conditional on satisfying the manager's budget). We group the 'amount from highest bidder' accepted into three bins: exactly equal, between 0 and $1 difference, and $1 upwards. The data include 205 workers who bid less than or equal to the $5 budget.