

EQUILIBRIUM EFFECTS OF PAY TRANSPARENCY IN A SIMPLE LABOR MARKET*

ZOË B. CULLEN[†] AND BOBAK PAKZAD-HURSON[‡]

First draft: November 2015

This draft: August 2018

Public discourse on pay transparency has not focused on equilibrium effects: how greater transparency impacts hiring and bargaining. To study these effects, we combine a dynamic wage-bargaining model with unique data of temporary work arrangements that differ in their level of transparency. Full transparency lowers wages by up to 25% and increases hiring by similar magnitudes. Earnings inequality falls, and employer profits rise significantly. A key intuition is high transparency commits employers to negotiating aggressively, because a highly paid worker's salary affects negotiations with other workers. 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, Eric Budish, Gabriel Carrol, Arun Chandrasekhar, Kalyan Chatterjee, Isa Chaves, Ken Chay, Bo Cowgill, Piotr Dworzak, Jack Fanning, Chiara Farronato, 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, Catherine Tucker, Emmanuel Vespa, Alistair Wilson, and seminar attendees at Stanford, UCLA Anderson, NYU Stern, HBS, Dartmouth, Brown, USC, UCSD, Penn State, Michigan SI, AEA 2017, SIEPR, NBER SI (Digitization and Personnel) for helpful comments and suggestions. We appreciate TaskRabbit for providing essential data access. Pakzad-Hurson acknowledges financial support by the B.F. Haley and E.S. Shaw Dissertation Fellowship through a grant to the Stanford Institute for Economic Policy Research. This paper subsumes two earlier drafts: *Equal Pay for Unequal Work?* and *Is Pay Transparency a Good Idea?*.

[†]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

Private information can have a large effect on the outcomes of negotiations, and in imperfectly competitive labor markets, it can lead to high levels of inequality (Manning, 2005). 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” through negotiations, by informing them of their employer’s willingness to pay for labor. (Phillips, 2009). But the equilibrium impacts of increasing transparency, namely how firms might change their hiring and wage-setting policies and how workers might adjust their bargaining strategies, are not well understood.¹ Despite this, 13 states have passed laws preventing employers from punishing workers who discuss wages with coworkers (Siniscalco et al., 2017) and California now requires employers to provide the range of salaries it pays current workers to prospective employees (Pender, 2017). The level of transparency in jobs arranged through online labor markets—including 7% of the US workforce (14 million people) (Robles and McGee, 2016)—remains unregulated. Among the design choices that online labor platforms make is whether to make past contracts visible, and who can communicate with each other about pay.

Our paper tests the equilibrium impact of pay transparency policies on wage negotiations, hiring, and profits in the context of online labor markets. To do so, we combine a dynamic wage-setting model with administrative data from TaskRabbit, an online platform over which employers and workers negotiate temporary work contracts, typically carried out locally. While the platform itself is 10 years old, the employers using the platform range from older and more established firms, including more than 30 Fortune 500 companies, to household employers (comprising roughly half of our sample). This setting allows us to study the impact of platform design choices affecting transparency, and it allows us to observe a level of detail about the wage negotiations between thousands of employers and workers that offer insights into other parts of the labor market as well.

To introduce the mechanisms we study in this paper, consider an employer and a worker when they first bargain over a wage. Neither party knows the exact value of the relationship to the other. If pay is transparent within the firm, the employer reasons that if she pays a high wage to the worker, other workers may learn this information and use it in future negotiations. Therefore, the employer will bargain aggressively to drive the wage down, and may decline to hire the worker at all if he asks for too high a wage. The worker similarly anticipates information spillovers; he expects to learn the wages of his peers quickly, and will then use the information of what the employer has paid others to renegotiate a higher wage

¹In 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.”

for himself. Therefore, it is less important for the worker to secure a high wage in this initial negotiation, so he bargains more passively to ensure that he is hired.

In a controlled setting where transparency is exogenously varied, we show that when workers can communicate with each other about pay, nearly all employers pay equal wages to all workers. Higher transparency also lowers wages by 25%, and increases the chances of hiring an additional worker by up to 40%. This leads to employer profits that are on average 60% higher under transparency.²

To lay out testable predictions of the equilibrium effects of transparency, we propose a model of dynamic wage negotiations.³ A continuum of workers individually bargain for wages with a firm(s).⁴ Time is continuous. Bargaining takes the form of directed take-it-or-leave-it (TIOLI) offers from each worker to the firm. Once employed, workers are able to renegotiate with the firm at will, again through TIOLI offers. Whenever an offer is rejected, the worker making that offer is permanently unmatched with the firm and receives her heterogeneous, exogenous outside option.⁵ We make the assumption of TIOLI offers, both in initial negotiations and renegotiations, to mirror our empirical setting, but we extend our results to various bargaining protocols in Appendix G. Over time workers stochastically learn the wages of their peers, and the level of transparency affects the rate at which this information arrives. Workers do not know their value to the firm, therefore, learning that an equally productive co-worker is earning a higher wage is an indication of being underpaid.

In equilibrium, and regardless of the level of transparency, a worker will only choose to renegotiate her wage once she learns the pay of her peers. When she learns this, she will (successfully) use this information in renegotiations to raise her own wage to the highest wage she observes. Therefore, transparency causes an information externality, as one worker’s wage can affect the renegotiations of others.

Transparency affects the frequency of renegotiations, which alters 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 a raise in pay for one worker is learned by others and used by them to renegotiate (demand effect).⁶ At the same

²As a guide to readers, Table I provides a roadmap for our theoretical and empirical results on each of these effects of transparency.

³The insights of our model can be adapted to understand the effect of hiding or revealing previous worker wages. Between 2016 and 2018, 9 states and municipalities passed laws that prohibit employers from requesting past salary information of potential employees during the hiring process (Cain et al., 2018), while internet platforms have done the opposite—eLance and UpWork now both include accounts of past contract payments on worker profiles. Please see Section III.E and Appendix H for details.

⁴For simplicity, we initially present our model as containing a single firm. We generalize our results to the case with multiple firms in Appendix D.

⁵Farrell and Greig (2016) find that online labor platform users earn on average one-third of their total income on platform, leading to heterogeneous outside options based on earnings off the platform.

⁶Another interpretation of the demand effect is that with increased transparency the firm becomes more

time, workers make lower initial offers to increase their chances of getting a foot in the door because they know they will quickly learn about the wages of others and be able to risklessly renegotiate (supply effect). These effects are present in static double auction models, first studied by [Chatterjee and Samuelson \(1983\)](#).⁷ 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). [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 effects, which cause simultaneous adjustment of bargaining strategies by workers and the firm in response to changes in transparency. This leads to novel model predictions:

Earnings are more equal with higher transparency, as workers renegotiate to a common wage.⁹ Because workers have heterogenous outside options, this implies worker surplus, the difference between earnings and outside option, is more dispersed with higher transparency.

We show that a higher level of transparency is more effective at increasing the hiring rate when the firm’s value of labor is low. Lower initial offers mean that a low-value firm can (and does) feasibly hire more workers, while a high-value firm will decline to hire workers with high wage offers to avoid information spillovers. This suggests that transparency mandates will be more effective at creating low-paying jobs within a particular industry. We also show more generally that either full privacy or full transparency minimizes the expected hiring rate, and we find the level of partial transparency that maximizes the expected hiring rate.

Pay transparency also changes the division of surplus between workers and the firm, leading to lower wages which benefit the firm at the workers’ expense.¹⁰ The demand effect

concerned with the ratchet effect ([Weitzman, 1980](#))—the firm mimics a lower productivity type to avoid workers discovering the true productivity of the firm and demanding a higher wage in future negotiations. For references to recent work on the ratchet effect, please see [Cardella and Depew \(Forthcoming\)](#).

⁷In their model, a buyer and a seller have private values for a particular good. Both agents simultaneously place bids, and if the bid of the buyer is higher than that of the seller, the two exchange the good at a price that is a predetermined convex combination of the two bids. We show a connection between the equilibria of our dynamic bargaining model and the equilibria of the static double auction weight–transparency affects the “weight” placed on workers’ bids when bargaining with a firm. [Larsen \(2015\)](#) empirically studies a related model using data from post-auction bargaining over used cars.

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.

⁹[Bewley \(1999\)](#), [Abowd et al. \(1999\)](#), and, more recently, [Song et al. \(2018\)](#) show that wages within firms are compressed relative to the wages across firms. Our paper may provide a microfoundation for these findings: wages are compressed within firm because employees rebargain to a common wage upon learning the firm’s willingness to pay for labor through transparency.

¹⁰[Faminow and Benson \(1987\)](#), [Kühn and Vives \(1995\)](#), [Alæk et al. \(1997\)](#), and [Nilsson \(1999\)](#) demonstrate instances in which price transparency among consumers can lead to increased prices (and firm profits) due to

allows the firm to commit to rejecting high wage offers to strengthen its bargaining position in renegotiations. Simultaneously, the supply effect causes workers to offer smaller markups over their outside options as transparency increases. The combination of these two forces shifts the de facto bargaining power to the firm. Under full transparency, workers learn the maximum wage the firm is willing to pay immediately upon matching with the firm. Each worker will either choose to work at this wage, or take her outside option. Therefore, the equilibrium outcome under full transparency is the same as a monopsonistic firm posting a wage.¹¹ Because of this, full transparency maximizes firm expected profits and minimizes worker expected surplus (Myerson, 1981; Williams, 1987).

If instead the firm endogenously picks the level of transparency then this choice signals the firm’s value for labor, affecting how workers negotiate. This leads to a unique equilibrium outcome in which the firm pools on full transparency regardless of its value of labor. The cause of this is unraveling (Milgrom, 1981). If low-value firm types were to maintain pay privacy, they would earn zero profits, as workers would always offer more than the firm’s value for labor. Instead, any of these low-value firm types could deviate to full transparency, effectively posting a wage below its value to ensure positive profits. But this causes workers to raise their offers when they do not encounter full transparency, resulting in new “low-value firm types” that receive zero profits by maintaining pay privacy. This logic can be iteratively applied until the firm chooses full transparency for any value.

Back-end data from TaskRabbit allows us to study and verify these equilibrium effects in an environment uncomplicated by pecuniary benefits or career concerns. Most jobs on the platform are low-skill, standardized tasks that clear quickly through a worker-bidding assignment process. Each accepted bid serves as a binding wage floor, but the employer can pay a bonus to any worker after the job is completed, allowing for the possibility of wage renegotiations on the job. 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.

We argue that the amount of pay transparency varies across jobs depending on two factors. First, there is partially transparency when co-workers are physically co-located while carrying out the task and can therefore communicate about pay through word of mouth. Second, some jobs are advertised with publicly-posted wages, making pay fully transparent.

collusion in oligopolistic settings. This mirrors our finding that transparency increases firm profits, although our result is not due to collusion between competing firms.

¹¹Because of this equilibrium outcome, our theory unifies previous results from a variety of models (and our empirical evidence corroborates these findings). For example: Michelacci and Suarez (2006) show that bargaining leads to more dispersed wages (our Theorem 1); Ellingsen and Rosén (2003) find wage posting is more effective when reservation wages are low (an implication of our Theorem 2); Brenzel et al. (2013) suggest that bargaining may lead to higher average worker wages (our Theorem 3).

The bargaining mechanism in our model has testable implications using TaskRabbit data. First, workers under partial transparency should receive the same pay as the highest earner, conditional on receiving a raise. Second, initial wage differences between coworkers should be uncorrelated with whether or not a worker receives a raise. Comparing the difference between final wages and the initial bids agreed to verifies these claims. When workers are co-located and employers adjust pay above the initial bid, final wages are nearly completely equal. Also, inequality between workers’ initial bids does not predict whether or not a worker receives a raise. These facts suggest that there are information spillovers about pay between workers (uncorrelated with bids), and after learning about relative pay, the information is effectively used to renegotiate up to the highest wage.

We rule out alternative mechanisms for the wage equalization patterns we observe, such as employer preference for equal wages and productivity spillovers between workers equalizing (observed) output. We also endogenize effort costs and embed worker social preferences into our model to study the psychological cost workers face after learning they are underpaid, which potentially leads to low morale and low effort (higher probability of leaving the job) (Akerlof and Yellen, 1990).¹² As a result, proactive employers may raise and equalize the wages of workers in pay transparent environments. We model these morale concerns as in DellaVigna et al. (2016) and Breza et al. (2017), and find that nearly all specifications predict a relationship between the extent of inequality of co-worker initial wages and the likelihood of wage equalization, which do not observe in our data. Also, without the presence of a bargaining channel, only very extreme and discontinuous morale cost functions replicate our empirical findings of full wage equalization. Morale costs must be so severe that all workers quit the job (expend 0 effort) upon finding out they are paid even small amounts less than a peer. We produce evidence (described in Section VI) that documents a morale effect on effort, but not enough of one to explain the observed equalization patterns. Previous work (egs. Breza et al. (2017); Charness and Kuhn (2007); Gächter and Thöni (2010)) also does not find the required extreme morale responses to explain the wage equalization we observe.

TaskRabbit data also allow us to study the endogenous choice of employers to select full transparency. In single-worker jobs, the employer can choose to be fully transparent in two ways: by publicly posting a price along with the job description, or by mentioning a wage in the text of a job description, which we can identify with natural language processing techniques. In line with our model, 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 price expectation up front, and

¹²Card et al. (2012), Mas (2017), Mas (2016), Perez-Truglia (2016), and Breza et al. (2017) conduct field and natural experiments on transparency, and document worker dissatisfaction upon learning of peers with higher pay. Charness and Kuhn (2007), Gächter and Thöni (2010), and Greiner et al. (2011) investigate similar claims in laboratory settings.

the gap in piece-rate work is even larger. Moreover the employers who choose transparent posted prices tend to be themselves lower income households, consistent with the theoretical prediction that high levels of transparency is optimal for employers with low values for labor.

We also observe the market unraveling towards higher levels of pay transparency across all employers, as our theory predicts. TaskRabbit staggered its entry into metropolitan areas across the U.S., creating segregated markets for local temp work of various levels of maturity. Across all markets, 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 of jobs or employers on the platform, nor do we find evidence that this is caused by employer learning.

We run a field experiment to measure worker welfare, and manipulate employer value for labor to observe the division of surplus. We also adopt an open bargaining protocol, allowing workers and managers to engage in conversation. We hire 293 managers and 920 workers from multiple online labor markets who are tasked with negotiating wages for, and carrying out, transcription services that take on average one day to complete.¹³ We elicit outside options from workers, and randomly assign a budget to managers, who are the residual claimants of that budget after paying wages. We randomize transparency by restricting wage negotiations to either a common chat room or split chat room, where the only difference is whether co-workers 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. We directly measure worker outside options, productivity, and employer profits.

Our field experiment results corroborate our theoretical findings and TaskRabbit analyses: pay is equalized when workers negotiate in a common chat room 97% of the time, compared to 64% of the time in private chat rooms. In the common chat room, manager profits are 60% higher, wages 25% lower, and hiring (the share of workers who reach an agreement with the employer) rises by more than 40% for managers with a small budget (low value of labor). Worker surplus is lower and more dispersed, both of which are predicted by our theory. Productivity levels are similar across the two groups.

We consider the effects of worker heterogeneity along gender lines. We elicit worker outside options in an incentive compatible manner in our field experiment and find that women have outside options on average approximately 10% lower than those of men. Bids for work reflect these differences. We find that the gender pay gap caused by this difference in outside options is mitigated by higher levels of transparency. We also find men are more likely to receive bonuses than women when there are communication channels between

¹³Another motivation for the experiment is replicability. The experiment is automated, and can be replicated using our code in a matter of days.

workers on TaskRabbit.¹⁴ We nest this privacy difference between genders into our model and show that intermediate levels of transparency can lead to very different arrival rates of wage information for men and women, potentially increasing the gender pay gap. This finding may give pause to those who advocate communication about pay to mitigate the gender pay gap.

The closest empirical literature to our paper focuses on publicly mandated full-transparency of individual pay. [Card et al. \(2012\)](#) experimentally nudge some University of California employees to visit a list of salaries published online and return to participants one week later to find aggregate worker satisfaction falls, and self-reported job search rises. [Mas \(2017\)](#) shows that top earners in municipal jobs experience a drop in nominal wages following the public revelation of wages, which he argues is primarily due to public aversion to visibly exorbitant salaries.¹⁵ Our analysis is based on a simpler labor market than in these prior studies. There are at least two important features of our setting to consider when determining where our results apply beyond the context of online labor markets. First, work is short term and relationships developed on the job may be correspondingly short lived. Second, workers complete more standard, low-skill tasks than in the general population, and their relative productivities may be more easily observed. We extend our model to allow for unobserved productivity differences between workers in [Section VIII](#). We show that the presence of these unobserved productivity differences does not change (quantitatively or qualitatively) the expected impact of pay transparency. However we do not have an empirical analogue to test these predictions.

The remainder of the paper is organized as follows. [Section II](#) lays out our model. [Section III](#) presents our main theoretical findings. [Section IV](#) describes the TaskRabbit market and contains empirical tests of our main findings using TaskRabbit data. [Section V](#) discusses our field experiment and related findings. [Section VI](#) examines an alternative model based on the fair wage-effort hypothesis and morale costs associated with pay transparency. [Section VII](#) investigates heterogeneity of workers across gender lines, and the effects of transparency on the gender pay gap. [Section VIII](#) extends our model and main results to settings in which workers have heterogeneous, and unobservable productivity differences. [Section IX](#) concludes. Omitted proofs, regression tables, and extensions are contained in the Appendix.

¹⁴This is consistent with survey evidence that women are more private about wages than men ([Goldfarb and Tucker, 2011](#)) and that women are less informed about the market value of labor than men ([Babcock and Laschever, 2003](#)).

¹⁵Another literature focuses on public disclosure of executive compensation in private sector jobs ([Faulkender and Yang, 2013](#); [Mas, 2016](#); [Schmidt, 2012](#)). Among the key mechanisms identified is the benchmarking of executive pay against rival executives, and the fear of public backlash against extreme inequality.

II. MODEL

II.A. Preliminaries

We propose and analyze a simple model of dynamic wage negotiations with some important elements: workers effectively bid for jobs by placing TIOLI offers and are able to renegotiate their pay, workers do not initially know their value to the employer, but learn how much they can obtain in renegotiations upon observing the pay of their peers. Our model of wage negotiation mirrors the bargaining process on TaskRabbit. The relative simplicity of TaskRabbit allows for many generalizations and extensions of our model which preserve our main findings, as we discuss in Sections VII and VIII, and Appendices D - G.

Time is continuous, and is indexed by $t \in \mathbb{R}_+$. There is a single firm in the economy.¹⁶ At each time t , a mass of workers enters the market, and each existing worker exogenously departs the market via a Poisson process with rate ρ . Each worker i has a private outside option $\theta_i \stackrel{iid}{\sim} G[0, 1]$, which is the flow payment i receives when not matched to the firm. Each worker also has a publicly observable type $\tau \in \mathcal{T}$ where \mathcal{T} is some countable set. Let $v_\tau \sim F_\tau[0, 1]$ be the productivity of type τ workers, which is known only to the firm. To simplify notation, we assume for the majority of the analysis that $|\mathcal{T}| = 1$, that is, all workers are equally productive with productivity $v \sim F[0, 1]$. This simplification does not meaningfully impact our findings. We discuss the case in which workers are heterogeneously productive and can be misinformed about their productivities in Section VIII.

The firm has a constant returns to scale production function, and receives a flow surplus v from each employed worker. 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 over the interior of their supports, 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 .

Workers bargain through a series of take-it-or-leave-it (TIOLI) offers to the firm. 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. Each worker i makes an initial offer w_i^* to the firm at the first moment she is hired, and can elect to make further TIOLI offers at any point during her employment, potentially renegotiating infinitely often. Two things can happen when a worker i makes a wage offer $w_{i,t}$ at time t . If $w_{i,t} \leq \bar{w}$ then i receives a flow wage $w_{i,t}$ for all time periods $t' > t$ until she departs the market or attempts to renegotiate. If $w_{i,t} > \bar{w}$ then i is permanently unmatched with the firm for all time periods $t' > t$ and consumes her outside option until she departs.¹⁷

¹⁶We generalize our model to include multiple firms in Appendix D.

¹⁷Clearly some bargaining friction must exist, or else workers could always ensure payments of \bar{w} arbitrarily

We model transparency as the random arrival of information about current wages. At time t each matched worker observes the set of wages the firm pays to currently 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 .¹⁸ For convenience, we assume that $W_t \equiv \bar{w}$ if the firm does not have any currently employed workers.¹⁹ Therefore, higher λ corresponds to more transparency.²⁰

The timing of the stage game is as follows at each time $t \geq 0$:

1. New workers enter the market and are matched with the firm.
2. Each matched worker i learns W_t independently with arrival rate λ .
3. Workers bargain according to the TIOLI protocol laid out above.
4. Existing workers depart independently at rate ρ .

In Appendix D we expand our model to allow workers to search for work across multiple firms, and show that many results are robust to this extension. In Appendix E we allow the firm to accept or reject each offer individually as it arrives, instead of picking a single \bar{w} at $t = 0$, and show that all results of this paper are unchanged under a selection of “Markov Perfect” equilibria. In Appendix F we allow the firm to reject worker renegotiation offers at cost but without permanently unmatching, and find an equilibrium of this game that mirrors that of the base model. In Appendix G we analyze alternative bargaining protocols, including one that allows for firm counteroffers, and show that our results are robust to these settings.

quickly by starting at wage 0, and continuously increasing offers until exceeding \bar{w} and getting rejected, then renegotiating for exactly that amount. The current assumption adds a draconian friction, but one that fits our context—if initial wages are set through a bidding process (as in TaskRabbit), those who bid too high will be rejected without the ability to readjust the bid. In Appendix F we generalize the game to allow for smaller frictions in the rebargaining stage, and show that all of our results are unchanged in this setting.

¹⁸Note that we are implicitly assuming that transparency is constant over time. This is perhaps a more reasonable assumption in short-term labor situations. In long-term labor markets with changes in transparency over time, our model predictions are unchanged if we assume that the firm picks a new anytime λ changes, and that workers who observe the wages of their peers also observe the λ at the time the previous wages were set.

¹⁹Without this assumption, the first cohort of workers in the full transparency equilibrium face an openness issue of wanting to renegotiate wages at the earliest time $t > 0$. It is also possible to complicate the timing of the game to resolve this issue, but at the expense of clarity.

²⁰For much of the paper, we abstract away from the genesis of this arrival process. We discuss endogenous information arrival in Section III.D, and differential rates of information arrival in Section VII. Alternatively, if workers are not able to initiate renegotiation at will (e.g. workers have to wait for a scheduled performance review to renegotiate), λ can be thought of as the rate at which workers both learn the wages of their peers and have the ability to renegotiate.

II.B. Equilibrium

We investigate pure strategy perfect Bayesian equilibria (PBE) of the game. We restrict our attention to equilibria satisfying the following conditions:

- A1** \bar{w} and w_i^* are strictly increasing and continuous functions of v and θ_i , respectively.
- A2** $0 \leq \bar{w} \leq v$ for all v . If $v \leq w_i^*$ for every worker i according to equilibrium strategies then $\bar{w} = v$.
- A3** $\theta_i \leq w_i^* \leq 1$ for all i . If there is no v such that $\theta_i \leq \bar{w}$ according equilibrium strategies then $w_i^* = \theta_i$.
- A4** At any time t , any worker i who makes an offer $w_{i,t}$ that should have been rejected according to equilibrium strategies but is not, believes with probability 1 that $\bar{w} = w_{i,t}$ until observing the wages of others via the transparency process.
- A5** Along equilibrium path, if any worker i renegotiations at time $t' > t$ then $w_{i,t'} > w_{i,t}$.

A1 limits our study to continuous equilibria, in which the highest wage the firm accepts (\bar{w}) and the initial wage offers of workers (w_i^*) are continuous functions of agents' types. This removes equilibria in which workers and the firm pool on a predetermined wage from consideration.²¹ **A2** and **A3** 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, $\bar{w} = 0$ and all workers choose $w_i^* = 1$.

A4 deals with histories that do not occur in equilibrium, and pins down off-equilibrium-path beliefs. If a worker makes an offer that is higher than the highest offer the firm is supposed to accept in equilibrium, yet the offer is accepted, the worker believes she is extremely lucky and is receiving the highest possible wage until she is presented with evidence to the contrary. These are the most favorable beliefs to the firm allowable in a PBE, so any equilibrium sustainable under **A4** is sustainable for any off-path beliefs.

A5 rules out a multiplicity of essentially equivalent equilibria in which workers "renegotiate" infinitely often by offering the same wage over and over again.

II.C. Renegotiation

We find that in equilibrium, workers renegotiate wages at most once.

²¹Leininger et al. (1989) suggest similarities between the set of continuous equilibria and a set of discontinuous equilibria of a game similar to our own, and so we do not believe this to be a conceptually limiting constraint. We discuss this similarity in Section II.D.

PROPOSITION 1. *The set of equilibria is non-empty. In any equilibrium workers only renegotiate wages in the first instant they receive information about wages of co-workers. Upon renegotiating, workers offer and receive \bar{w} .*

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 t and offer $w' > w$ at time $t' > t$ if I have not learned the wages of my co-workers” 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 t .²²

Due to the continuum of workers entering the market at each time, in addition to our equilibrium selection criteria, workers trace out the entire set $[0, 1]$ with their initial offers at each time t . Therefore, the highest wage paid by the firm is \bar{w} for all $t \geq 0$. As a result, the maximum wage that any worker observes upon information arrival is \bar{w} . Clearly, a worker will then demand this amount from the firm.

II.D. Solving for equilibrium

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 C that each worker solves

$$w_i^* = \operatorname{argmax}_{w_i} \int_{w_i}^1 ((1 - \Lambda) w_i + \Lambda x - \theta_i) \bar{f}(x) dx \quad (1)$$

and the firm solves

$$\bar{w} = \operatorname{argmax}_w \int_0^w (v - (1 - \Lambda) y - \Lambda w) \bar{g}(y) dy \quad (2)$$

where $\Lambda = \frac{\lambda}{\rho + \delta + \lambda}$ for all $\lambda \in [0, \infty)$ and $\Lambda \equiv 1$ for $\lambda = \infty$. λ is the arrival rate of information, and Λ is the *effective* level of transparency.

With the exception of Section VII, we use Λ to represent transparency in the rest of the paper. The reason is that this parameter encapsulates the relevant information from agents’ perspectives. A high rate of information arrival λ will be unimportant to workers if the discount rate δ and/or departure rate ρ are sufficiently higher than λ . $\lambda = \Lambda = 0$ corresponds to full privacy, while $\lambda = \infty$ and $\Lambda = 1$ correspond to full transparency. There are uncountably many combinations of parameters δ, ρ , and λ that correspond to any

²²This reasoning is shared in Lazear (1986) and Tirole (2016).

$\Lambda \in (0,1)$. However, fixing ρ and δ , there is a bijection between λ and Λ with higher λ corresponding to higher Λ .

Equations 1 and 2 are the same objective functions as those in the well-known [Chatterjee and Samuelson \(1983\)](#) double auction in which a seller (worker) with a private value for a good (θ_i) and a buyer (firm) with a private value for a good (v) submit sealed bids. If the bid of the buyer is at least as large as that of the seller, the good switches hands at a price set by a predetermined convex combination of the two bids (determined by Λ). Therefore, the equilibria of our model coincide with the equilibria of [Chatterjee and Samuelson \(1983\)](#), in which higher transparency shifts the average wage of employed workers toward the maximum wage set by the firm, \bar{w} . The first order conditions for workers and the firm are, respectively:

$$w_i^* - \theta_i = \overbrace{(1 - \Lambda)}^{\text{direct effect}} \cdot \underbrace{\frac{1 - \bar{F}(w_i^*)}{\bar{f}(w_i^*)}}_{\text{indirect effect}} \quad (3)$$

$$v - \bar{w} = \overbrace{\Lambda}^{\text{direct effect}} \cdot \underbrace{\frac{\bar{G}(\bar{w})}{\bar{g}(\bar{w})}}_{\text{indirect effect}} \quad (4)$$

The set of equilibria corresponds to solutions of the first order equations, and given the equilibrium strategy of the firm, workers have a unique best response, and vice versa ([Satterthwaite and Williams, 1989](#)).

The optimal bidding and wage-setting policies of workers and the firm, respectively, are interdependent. 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. [Satterthwaite and Williams \(1989\)](#) show that there exists a continuum of equilibria satisfying Equations 3 and 4. Our 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 produce similar evidence in a setting similar to TaskRabbit, shown in Figure B4. We use this evidence as our equilibrium selection criterion, and therefore, we focus our analysis on linear equilibria. To do so, we restrict attention to a two-parameter family of power law distributions of worker outside options and firm values, which we show admit a unique linear equilibrium.²³ We then study the properties of this equilibrium, and analyze the effects of transparency. The class of distributions we study are:

²³The approach of making parametric assumptions to ensure linear equilibrium is common. One recent example on CEO pay is [Edmans et al. \(2016\)](#). Power law distributions are commonly observed in economic situations such as ours, including worker income and firm productivities. See [Gabaix \(2009, 2016\)](#) for details.

$$\begin{aligned} F(v) &= 1 - (1 - v)^r, & r > 0 \\ G(\theta) &= \theta^s, & s > 0 \end{aligned} \tag{5}$$

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.

DEFINITION 1. *A linear equilibrium is a pure strategy perfect Bayesian equilibrium satisfying **A1-5**, 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 2. *For any pair of distributions within the family described in Equation 5 there exists a unique linear equilibrium.*

II.E. 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**). We graphically represent the **demand** and **supply effects** in Figure II as Λ increases, and the following proposition formalizes these arguments.

PROPOSITION 3. *\bar{w} and w_i^* are decreasing functions of Λ . As $\Lambda \rightarrow 0$, $\bar{w} \rightarrow v$ for all $v \in [0, 1]$. As $\Lambda \rightarrow 1$, $w_i^* \rightarrow \theta_i$ for all $\theta_i \in [0, 1]$.*

III. MAIN RESULTS - EFFECTS OF TRANSPARENCY ON EQUILIBRIUM

We analyze the equilibrium effects of increasing transparency along three dimensions: pay inequality, the hiring rate, and division of surplus.

III.A. 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.

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.²⁴ 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 . Figure II shows that the relative impact of transparency on w_i^* is smaller the larger θ_i is. 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 compressed expected earnings, regardless of δ and ρ . We document these two effects by plotting the expected difference in wages between workers i and j over time and for different levels of transparency in Figure III.

THEOREM 1. *Let $\theta_i > \theta_j$, $1 > \Lambda'' > \Lambda'$, and suppose workers i and j are both hired in equilibrium under Λ' and Λ'' .*

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

Note that the first point in the above theorem does not apply to full transparency. Under full transparency, there is a discontinuity because workers make their initial wage offers *after* seeing the wages of their coworkers, therefore $w_i^* - w_j^* = 0$ and $T(\Lambda'', v, \theta_i) - T(\Lambda'', v, \theta_j) = 0$ when $\Lambda'' = 1$, 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. *Let $\theta_i > \theta_j$, $1 > \Lambda'' > \Lambda'$, and suppose workers i and j are both hired in equilibrium under Λ' and Λ'' .*

1. *The difference in initial surplus $(w_j^* - \theta_j) - (w_i^* - \theta_i)$ is smaller under Λ'' than Λ' and*

²⁴The restriction that workers be hired under both transparency levels is necessary, as we show in Theorem 2, 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 C 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 compressed by increased transparency.

2. Let $S(\Lambda, v, \theta_k)$ be the equilibrium expected discounted total surplus of a worker k with outside option θ_k under transparency level Λ and firm value v conditional on k being employed at the firm. Then $S(\Lambda'', v, \theta_j) - S(\Lambda'', v, \theta_i) > S(\Lambda', v, \theta_j) - S(\Lambda', v, \theta_i)$, and $S(\Lambda'', v, \theta_j) - S(\Lambda'', v, \theta_i) \rightarrow \frac{\theta_i - \theta_j}{\rho + \delta}$ as $\Lambda'' \rightarrow 1$.

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 equalization in Section [V.E](#).

III.B. Hiring Rate

Increasing transparency has competing effects on the hiring rate. Let \bar{w}_Λ denote the maximum wage the firm pays and $w_{i,\Lambda}^*$ the initial offer of worker i for transparency level Λ . When transparency increases from Λ' to Λ'' , the **demand effect** lowers the hiring rate. $\bar{w}_{\Lambda''} \leq \bar{w}_{\Lambda'}$ meaning that there are fewer workers with $\theta_i \leq \bar{w}_{\Lambda''}$ who are eligible for employment. The **supply effect** increases the hiring rate. $w_{i,\Lambda''}^* \leq w_{i,\Lambda'}^*$ for all i so fewer workers over-negotiate by initially offering $w_{i,\Lambda''}^* > \bar{w}_{\Lambda'}$.

THEOREM 2. *The expected proportion of workers hired in equilibrium is concave in Λ and maximized at*

$$\Lambda^* = \frac{1 - \mathbb{E}(\theta)}{1 + \mathbb{E}(v) - \mathbb{E}(\theta)} \quad (6)$$

and the ex-post employment maximizing level of Λ is weakly decreasing in v .²⁵

An interior level of transparency maximizes the hiring rate. Due to the concavity of the hiring rate in Λ either full privacy or full transparency is employment *minimizing*.

Λ^* is decreasing in both $\mathbb{E}(v)$ and $\mathbb{E}(\theta)$. 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 [3](#). As $\mathbb{E}(v)$ decreases, the firm's markdown $v - \bar{w}$ is likely to be small regardless of Λ . 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 Λ . Increasing transparency has little effect on these premia, but instead discourages the firm from setting a large markdown.

Both the ex-ante and ex-post optimal levels of transparency are decreasing in v . We have already discussed how Λ^* is strictly decreasing in $\mathbb{E}(v)$, and increasing transparency is

²⁵The expected match surplus is $\mathbb{E}(v) - \mathbb{E}(\theta)$, so $\Lambda^* = \frac{1 - \text{expected outside option}}{1 + \text{expected match surplus}}$.

more beneficial for the hiring rate when v is small. Given that workers' initial offers are not affected by the realization of v (they do not observe it), high transparency causes firms to reduce \bar{w} more significantly when v is high, leading to relatively less hiring. Additionally, this comparative static on the ex-post employment maximizing level of transparency also holds for the ex-post *social surplus* maximizing level of transparency. 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 employed worker increases social surplus by $v - \theta_i > 0$, implying that social surplus is strictly increasing in the hiring rate. Therefore, increasing transparency is also more beneficial from a social surplus perspective when the firm has a small realization of v .

III.C. Division of Surplus

Increasing pay transparency benefits the firm and hurts workers, and lowers 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 Λ 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 increasing Λ 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 ($\Lambda = 0$) and full transparency ($\Lambda = 1$). In the full privacy equilibrium, 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 every worker learns the wages of others within the firm at the instant they are hired, before their 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 [Myerson \(1981\)](#) implies that each party prefers to be the one making the once-and-for-all offer to the other.

THEOREM 3. *The ex-ante expected equilibrium profit of the firm is strictly increasing in Λ . The expected equilibrium surplus of workers and expected discounted wages conditional on employment are strictly decreasing in Λ .*

III.D. Endogenous selection of transparency

Until now, we have been studying the effects of transparency from an ex-ante perspective, or 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 .

THEOREM 4. *Suppose the firm selects Λ as a function of v . There is an essentially unique equilibrium satisfying **A1-5**. In it, the firm selects $\Lambda = 1$ for all v .*

For intuition, suppose the firm can only select $\Lambda \in \{0, 1\}$. We show that the firm cannot set $\Lambda = 0$ in equilibrium due to unraveling (Milgrom, 1981). 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 bid 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. As $\bar{w} = v$ when $\Lambda = 0$, a deviating firm will make zero profits.

Fixing \bar{w} , higher transparency is better for every worker, so every worker has an incentive to seek out the wages of their coworker to the highest possible degree. We have not formally modeled the choice of workers to "bury their heads in the sand" and ignore information about their coworker's wages. Nevertheless, a richer model that allows each worker i to specify Λ_i such that the effective transparency to worker i is $\min\{\Lambda_i, \Lambda\}$ will result in each worker i setting $\Lambda_i \geq \Lambda$ because no single worker will affect the equilibrium payoff, and therefore actions, of the firm.

III.E. Worker wage history

Our model can be adapted to understand the effect of hiding or revealing previous worker wages. Suppose a worker receives a job offer from a different firm at each time t and can either accept the wage and complete the job at time t or consumer her outside option. In either case, the job finishes at time t and the game continues. Suppose, for clarity of argument, that there is no transparency of wages within the firm at each time t .

We formalize this model in Appendix H. We find that many of the effects we note in our base model mirrored in this setting, most prominently, that the revelation of previous worker wages increases the bargaining position of the worker and increases her surplus at the expense of firms.

IV. EMPIRICAL EVIDENCE FROM TASKRABBIT

IV.A. Platform

We use an administrative dataset of temporary workers and employers, matched on an online labor platform, TaskRabbit, between January, 2009 and June, 2014 to test our theoretical predictions. TaskRabbit differentiates itself from other online labor platforms by specializing in local jobs which account for 89% of jobs completed. The platform is active in 19 metropolitan areas across the U.S. during this period.

Our research concentrates on jobs that are posted as one-time tasks. Among the most common categories of jobs are deliveries, shopping, 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 five stars, and a short bio.

Employers post a description of the task, details about the exact location, number of workers needed, frequency of task, and a deadline for completion. Workers search through these postings and submit bids for tasks they are interested in completing. Alternatively, the employer can choose to post a TIOLI price, and the first worker to accept is matched.

IV.B. Bargaining Environment

Employers can elect to increase wages through the platform once the job is completed. As jobs frequently involve face-to-face contact between employers and employees, this allows for the possibility of on-the-job wage bargaining.²⁶

At the time that a worker is assigned a job, the worker and employer enter a contract that can be cancelled by contacting the platform and providing a reason. TaskRabbit has a three strike rule. After three cancellations, a user will not be permitted to use the site again. During the window between when the match is made and the job is complete, money is held in escrow and is released to workers by default after a predetermined deadline.

Once the job begins, several additional frictions make canceling costly for both parties. TaskRabbit is a spot market designed for urgent tasks. Conditional on completion, 97% of

²⁶TaskRabbit 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.

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 cannot costlessly transition to another job. Because these are 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 five stars for a hired worker. There is no reputation system for employers. With very few exceptions (less than 1% of the time), employers do not leave reviews lower than 4 stars, but they frequently decline to rate workers. Employers only observe worker average rating, not the number of declined reviews. We add to our performance measure “missing” reviews, based on prior research that missing reviews are skewed toward negative feedback.²⁷

IV.C. Measuring Transparency

We measure pay transparency on TaskRabbit 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 TIOLI 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 in the same location, with longer co-location leading to higher transparency. We distinguish settings inherently suited for either co-located workers and physically separated workers, 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 use the street address 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 co-worker communication and length of time together. The survey evidence suggests that workers learn about each others’ bids 47% 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.

²⁷The literature on user generated content has identified a number biases and manipulation techniques that we can address using data about performance 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 show the share of missing reviews on TaskRabbit predicts whether an employer returns to the platform, TaskRabbit’s central measure of employer satisfaction, and the worker star-rating conditional on receiving a rating (Table A1 Col. 1 and Col. 2 respectively). The share of positive reviews a worker receives is correlated with ex-post pay, but not the ex-ante bid accepted (Table A2), suggesting that we are really detecting the performance that the employer observes on the job.

In Table I we summarize these transparency measures and how they are used in our empirical analysis. In Table II we report characteristics of workers and tasks on TaskRabbit according to these transparency classifications.

IV.D. Verifying Bargaining Assumptions

The premise of our model has two clear empirical implications for the outcome of a re-bargaining process that are demonstrated empirically using TaskRabbit data in Table IV.

SF1: Initial inequality is uncorrelated with employers’ decisions to compress pay.

SF2: Employers equalize wages if they compress pay at all.

Conditional on a worker receiving higher final pay than her initial bid, she receives final pay that is statistically no different than that of the highest accepted bidder (SF2). In Col. 4-6 of Table IV we estimate the relationship between the amount that the worker bid below the highest bidder and the amount of the raise paid at the end of the job, conditional on receiving a raise. A coefficient equal to 1 (on “amt. under top bid %”) implies that the raise completely equalizes pay. We estimate a precise relationship statistically indistinguishable from 1. At the same time, the difference in initial wages does not predict whether or not a worker receives any raise at all (SF1). In Col. 1-3 of Table IV, we test if the amount the worker bids less than the highest accepted bid predicts whether the worker will receive a raise. A coefficient equal to 0 (on “amt. under top bid %”) implies no conditional correlation between a worker’s bid and the probability of receiving a raise.²⁸ We estimate a small effect that is statistically indistinguishable from 0.

We argue that renegotiation results from communication about pay. While there are other candidate mechanisms, such as productivity spillovers when co-located or employer preferences for equity, they fall short of matching key facts.

Perceived productivity differences as the explanation for the wage equalization we observe requires that assessed performance of co-workers is less dispersed when workers are together. We find ratings employers give to workers are no more dispersed or compressed when workers are co-located, detailed in Appendix A. More generally, there is a small and statistically insignificant correlation between bids and ratings. As a result, any systematic pattern of spillovers from high to low productivity types does not necessarily raise the performance of the low bidder or the pay of the low bidder per se.

Another potential channel is employer preferences for equity *only when workers are co-located*. We collect survey evidence (Appendix J.1) of how likely workers are to talk to and ask about another worker’s pay on the job and find the frequency of renegotiation is positively

²⁸We exclude the highest bidder in this analysis.

correlated with the probability of communication even within co-located jobs. Duration as well as isolation together are both predictive of exchanging information. For example, when workers are hired to transcribe in a classroom, the likelihood that they talk to each other about pay is relatively small compared to those packing boxes, and so too is the likelihood that their pay is equalized. Therefore, this explanation does not appear plausible.

In Section VI we analyze and describe the shortcomings of an alternative model in which worker preferences for equality drives wage compression. In Section V we compare results from our analysis of TaskRabbit side-by-side with the results from our experiment in which we exogenously vary transparency and observe negotiations directly. Our experimental results are consistent with our findings from observational data.


IV.E. Wage Compression

Theorem 1 states that increased transparency leads to compression in the wages of employed workers. We present a visual depiction of wage compression across two types of tasks, higher transparency jobs that require multiple workers to be situated together and lower transparency jobs that can be carried out separately.

Figure I offers a visual depiction of the variance in wages for workers assigned to the same job on TaskRabbit. Each dot represents a multi-worker job. Observations that fall beneath the 45 degree line illustrate the tendency for employers to reduce the variance in ex-post payment relative to the variance of ex-ante bids. In higher transparency, co-located work settings 22% receive pay that is higher than their bids, as opposed to 7% in lower transparency, physically separated work settings. We do not expect full wage equalization because, as we illustrate in Table IV, the workers who renegotiate are a subset of all workers, who learn about the bids of their co-workers on the job. In Appendix Table B4 we carry out a similar test at the level of individual bids with controls for worker performance and again find wage compression only when workers are physically co-located. We find that a 10% gap between a worker’s initial bid and that of the highest bidder results in an approximately 4% increase in final pay on average.

We interpret co-workers’ bids as having a causal effect on each others’ final pay when co-located. 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 may not anticipate multi-worker settings. They comprise fewer than 5% of all posted jobs and appear as a single post rather than multiple postings. For more detailed analyses of non-strategic bidding and worker selection, see Appendix Section B.1.

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=291) and multi-worker co-located jobs (Panel b, N=399). 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 an 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 fail to accept 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.

IV.F. Hiring

TaskRabbit administrative data includes both job posts that are successfully matched with a worker and those that go unmatched, offering us a measure of unmet labor demand.²⁹ We refer to hiring rate, in the context of TaskRabbit, as the proportion of posted positions which are ever filled.

Our model of transparency illuminates an important reason for vacancies: a worker may overbid for a job posted by a low-value employer if wages are not transparent. The bid includes a premium above the outside option which is optimal in expectation without knowledge of v , but sub-optimal if v were known to be low.

As a result, we expect low-value employers' positions to be vacant more often, and we also expect low-value employers to endogenously choose to announce a wage publicly more often than other employers (Theorem 2).

We use annual earnings of household employers as a proxy for willingness to pay. One reason to favor this measure is that, from a survey of employers conducted by TaskRabbit, the most common alternative to using TaskRabbit to complete a task is to do it oneself. Using money as measure of the opportunity cost of time, higher income employers are more likely to have higher time costs (i.e. leisure is a normal good).

²⁹Cullen and Farronato (2016) find TaskRabbit to be a slack labor market with highly elastic labor supply, supporting the notion that unfilled tasks reduce total work completed by workers and their wages on platform.

We find that below-median earners in each city are slightly less likely to fill their jobs. Col. 1-3 of Table VI shows employers are 2% less likely to fill their positions. They are also 6% more likely to choose a transparent posted price (Col. 1-5 of Table B9). Finally, conditional on choosing posted price, below-median earners gain the greatest boost to their vacancy fill rates. Our preferred specification, Col. 3 of Table VI, shows low-value employers gain a relative boost of 4% more than other employers when they choose to publicly post a price for their jobs.

IV.G. Profits

Theorem 3 predicts higher levels of transparency are associated with higher expected employer profits and lower wages. In TaskRabbit, we do not directly observe employer profits, but we can observe labor costs. We compare the difference in bids, sent from the same workers for similar 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). Private auction jobs with and without price mention appear similar along observable characteristics, which we test extensively by analyzing the text of the jobs through unsupervised machine learning. We show the comparison of traits in Table II. In Table VII, Col. 1-3 we include worker fixed effects and compare the bids and final pay of the same worker as a function of whether the job posting contained any information about the expected price of the job before the bid was submitted. In Col. 3, we restrict the sample to only those employers who pay hourly wages (rather than a piece rate for the entire job), and include worker fixed effects and performance evaluations. We show earnings are 7.8% lower when the employer mentions a price in the job description.

IV.H. 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. As marginal employers switch from choosing privacy to transparency, privacy signals an even higher valuation and attracts even higher bids, until all switch to full transparency.

We observe the predicted unraveling in TaskRabbit by studying segregated city-markets for local tasks. For most jobs, an employer chooses full transparency through posting a price, or full privacy by soliciting private bids.³⁰ Figure IV shows the share of posted price jobs in each metropolitan area in June, 2014. Markets where TaskRabbit has been active for

³⁰Fewer than 5% of jobs involve multiple workers whose interactions cause partial pay transparency.

longer 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, Figure V shows that the proportion of transparent posted price jobs rises by 1% per month.³¹ This trend is not explained by the changing composition of jobs or employers on the platform (for more details, see Table B11), nor do we find evidence that this is caused by employer learning, as employers rarely experiment by switching between posting prices and accepting private bids.³²

V. EMPIRICAL EVIDENCE FROM A FIELD EXPERIMENT

We conduct a field experiment to further test our findings in a controlled environment. We hire 293 managers and 920 workers from multiple online labor markets who are tasked with negotiating wages for transcription services that take on average one day to complete.³³ We randomize transparency by restricting wage negotiations to either a common chat room containing a manager and multiple workers, or separate each worker into a private chat room with the manager. We document all interactions in these chat rooms, and we place no restrictions on the ways in which workers and managers can bargain.

We 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 Section VI we use this additional data to test the relative importance of social concerns surrounding pay transparency, and run a horse race between such a model and our re-bargaining model.

V.A. Procedure

Participants are recruited from TaskRabbit, Mechanical Turk, Upwork, and Craigslist 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

³¹Our theory predicts immediate unraveling to full transparency, but observed unraveling in labor markets typically takes time, as discussed in Roth and Xing (1994). This dynamic is consistent with a richer, bounded rationality formulation of our model, a la Kandori et al. (1993), in which employers select a level of transparency (to post a price or to accept private bids) based on which scheme would have maximized their expected profits in the previous period.

³²Einav et al. (2018) study a shift in eBay sales away from an auction format and toward posted prices. We discuss their model, its relation to our setting, and the differences between TaskRabbit and eBay in Appendix B.2.

³³This experiment is completely automated and can be replicated using our code in a matter of days.

wage. Each worker can complete up to 5 pages of transcription. Each manager is assigned to an average of 3 workers and privately given a per-worker-page budget, either \$5 or \$9. Managers are the residual claimant of this budget, after wages are paid for completed work.

Before interacting, workers are shown a sample transcription 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 also survey other characteristics including the expected time it takes to complete each page, daily household income, management/transcription experience, gender, location, and age.

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 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. Workers who agree to a wage with a manager have 48 hours to complete up to 5 pages of transcription.

V.B. Treatments

The experiment follows a $2 \times 2 \times 2$ design. The number of participants assigned to different treatments is displayed in [Table III](#). One treatment is the public visibility of wage negotiations. Managers either negotiate wages with each worker in a separate, private chat room (“privacy” treatment), or the manager negotiates over a common chat room (“transparency” treatment) with all workers. The only difference between the the transparency and privacy treatments is a division between chat forums allowing (or preventing) workers to observe all communication between co-workers and the manager. In both cases, the manager must confirm wages individually with each worker in the chat room. User interface differences are minimal as the multiple manager chat rooms in the privacy treatment all appear on one page side-by-side. For pictures of the interfaces, please see [Appendix I](#).

The second condition varies the budget assigned to managers, either \$5 per worker per page or \$9 per worker per page.

The third treatment either requires managers to accept all bids less than or equal to the budget, or allows them to actively bargain with workers.

V.C. Administrative Details

We present results from two rounds of experimentation that follow the same procedure and differ only in the degree of automation. One version relies on us, the experimenters, to invite workers to chat rooms and collect transcriptions via email following the intake survey. Another is completely automated in this dimension, and all interaction occurs through a single web interface programmed in oTree (Chen et al., 2016). oTree became available after the initial round of our experiment. More than three-quarters of our participants interacted through the automated oTree system, and results are comparable across the two interfaces. Table III shows that, along ex-ante observable characteristics, the workers and managers randomly assigned to different treatments are comparable.

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.

V.D. 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.

Manager: You agreed to \$1 per page?

Worker: I really don't remember, it sounds good but I suppose you would give the same to everyone? I see you gave \$5 to [other worker].

⋮

Manager: Okay, \$5 per page!

V.E. Analysis and Main Findings

Wage and Surplus Equalization:

Pay is significantly equalized in the transparency treatment compared to the privacy treatment. In the common, transparent chat room, pay is completely equalized 97% of the time. Of a subset of managers allowed to negotiate wages with worker, 33 of 37 in the transparent pay treatment pay common wages to all workers, compared to only 32 of 53

managers in the private pay treatment.³⁴ This further corroborates theoretical predictions SF1 and SF2.

We measure worker surplus as the final payout received less the elicited reservation value. While we observe significant compression in final pay under transparency, we observe amplification in the dispersion of worker surplus, as predicted by Corollary 1. On average, the Gini coefficient (a measure of inequality) for worker surplus nearly doubles in the transparency treatment (Table V, Col. 5-6). 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.

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 compression (Table V, Col. 3-4).

Hiring Rate:

We define the hiring rate as the proportion of workers who agree on a wage with their manager. Theorem 2 states that the positive effect of transparency on hiring is higher when the value of labor is low. The reason is that a worker will overbid for a job posted by a low-value employer if wages are not transparent. The bid includes a premium above the outside option which is optimal in expectation without knowledge of v , but sub-optimal if v were known to be low.

This experiment provides a direct test of these predictions under an open bargaining protocol. Managers with a low value for labor, or a \$5 budget per page of transcription, are less likely to reach an agreement with workers than high budget managers. We show in Col. 4-6 of Table VI that the hiring rate rises by approximately 40% for low-budget managers when negotiations are held in the common chat forum (full transparency) as opposed to the private forum. On the other hand, when the manager's budget is \$9, employment remains statistically unchanged when negotiations are transparent than when they are private. This provides evidence that higher transparency is more effective at improving the hiring rate when the employer value (manager budget) is low.

³⁴In order to maintain consistency throughout the analysis of compression, we restrict the sample of managers to those with a \$5 budget, and who are assigned to multiple workers with reported time estimates for completing the task.

Profits:

Our theory predicts that profits will be higher under full transparency than under full privacy. In Col. 6 of Table VII, manager profits are over 60% higher in the transparent group than the private group. This is consistent with the strong theoretical prediction that workers will bargain more aggressively under conditions of pay privacy, asking for a higher premium over their outside option. The employer does not have a credible way to commit to a maximum willingness to pay, as she does in the common chat room where employees can see her turn down previous bidders demanding more than a chosen limit.

VI. EFFECTS OF SOCIAL CONCERNS

We incorporate an additional channel, social concerns about relative pay, that is a prominent explanation in the literature for compressed wages in transparent environments. Several papers (Akerlof and Yellen, 1990; Breza et al., 2017; Card et al., 2012; Mas, 2017; Perez-Truglia, 2016) argue that a worker exerts less effort upon learning she is underpaid relative to her peers. A proactive employer may augment the wages of workers who learn they are underpaid in order to avoid low effort provision (Eliaz and Spiegler, 2013).³⁵

VI.A. Endogenous Worker Effort

To test this competing mechanism, we build a model of bargaining under transparency that endogenizes worker effort. As before, workers make TIOLI offers $w_{i,t}$ and observe the wages of their peers at rate λ . Now each worker i also selects $e_{i,t} \in [0, 1]$ 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 $\theta_i \cdot e_{i,t}$. Each worker i 's expected flow payoff is $(w_{i,t} - \theta_i) \cdot e_{i,t}$. We show that workers bargain in the same way as in the original game in equilibrium. Therefore, all of the results of the paper carry through when effort is endogenous.

PROPOSITION 4. *There is a unique linear equilibrium of the endogenous worker effort game. In it all workers set $e_{i,t} = 1$ for all periods of employment. All other actions are the same along the equilibrium path as those in the equilibrium of the original game.*

VI.B. Proactive Employer Model

We now remove workers' ability to renegotiate in order to study employer decisions to raise wages to avoid effort reduction. Workers make a wage offer w_i^* when they are initially

³⁵This, of course, hinges on the ability of the firm to monitor a worker's output or knowledge of the wages of others within the firm.

hired, and thereafter only choose effort. We include a morale cost to a worker for learning she is underpaid, modeled as in DellaVigna et al. (2016) and Breza et al. (2017). Workers face a higher cost to effort upon learning they are paid less than any peer, and this cost is non-decreasing in the effort the worker provides and the difference between her wages and that of her highest paid coworker.³⁶ At each time t the firm can observe whether a worker has received wage information and can elect to unilaterally increase her wage for all times $t' \geq t$ to reduce the morale cost. 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 co-workers 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. It is easy to see that the firm will increase the wage of a worker i at time t only if i learns the wages of her co-workers 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 and our field experiment.

PROPOSITION 5. *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 $d \in (0, 1]$ does there exist an equilibrium satisfying regularity conditions **A1-4** in which the firm sets $w_{i,t} = \bar{w}$ for each worker i who learns \bar{w} at time t for all Λ .*

Only large and discontinuous morale cost functions result in the same predictions as the bargaining model. Unless workers optimally reduce 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. Intuitively, 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 upon learning she is underpaid.³⁷

³⁶If we use the morale specification in this section and give the workers the ability to make TIOLI offers to the firm, 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.

³⁷Fanning and Kloosterman (2017) 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 all workers rejecting all wage offers by managers which are bounded away from full equality in our experiment, as Fanning and Kloosterman (2017) predict.

We do find evidence that relative pay concerns lead to lower effort when a worker learns she is underpaid, but output declines proportionally with the extent of inequality.³⁸ We show in Figure VI that when we do not allow for any renegotiation, output is decreasing in the difference between a worker’s pay and that of her highest paid peer in a common chat room compared to separate chat rooms.³⁹ In our data, effort reduction appears convex in the degree of inequality, implying morale concerns would make it profitable to raise wages more often when inequality is small, and to potentially only partially close the gap (rather than equalizing wages) in some cases. Recalling our findings of full wage equalization, regardless of initial wage differences in both TaskRabbit data and our experiment, and in light of Proposition 5, reduced effort due to relative pay concerns do not unilaterally generate pay equalization for a profit maximizing employer. Our findings suggest that re-bargaining is an important mechanism equalizing wages in the transparent pay environments, even in the presence of morale concerns.

VII. GENDER DIFFERENCES AND THE GENDER PAY GAP

We study the effects of differences between men and women along two dimensions, outside options and the spread of peer wage information. We extend our empirical and theoretical analyses to shed light on how transparency affects wage inequality when genders differ along these dimensions. We believe this to be an important avenue of inquiry as pay transparency is commonly cited as a way to close the gender wage gap.⁴⁰

We elicit the outside option of workers in our experimental setting as discussed in Section V. We find that women have on average 9.7% lower outside options than men, and we show that increasing pay transparency can close a pay gap caused by differences in outside options.

To see this theoretically, let there be two types of workers, m (male) and f (female), such that $qG_m(x) + (1 - q)G_f(x) = G(x)$ for all $x \in [0, 1]$, where $q \in [0, 1]$ is the proportion of men

³⁸Robert Duvall refused to take part in *The Godfather Part III* stating, “if they paid Pacino twice what they paid me, that’s fine, but not three or four times, which is what they did” (<http://www.imdb.com/name/nm0000380/bio> accessed 11/7/2016). Similarly, in our experimental treatment in which managers must accept worker bids as final wages, one low-bidding worker remarked, “yeah, I won’t be working for less than a third of what others are getting for the same amount of work” before ending his participation.

³⁹Average output is 18% lower in the transparent no negotiation treatment than the non-transparent no negotiation treatment. This magnitude is in line with a field experiment in Breza et al. (2017) who find a 22% reduction in output among workers who learn they are paid less than their peers. Moreover, they also present a similar trend of a smooth decrease in output as inequality rises.

⁴⁰For example, Hillary Clinton championed transparency during her presidential campaign. Her website reads, “As president, Hillary will: Work to close the pay gap...We should promote pay transparency across the economy and work to pass the Paycheck Fairness Act—a bill Hillary introduced as senator—to give women the tools they need to fight discrimination in the workforce.” <https://www.hillaryclinton.com/issues/womens-rights-and-opportunity/>, accessed 3/20/2018.

in the market. The first, and simplest, result is an application of Theorem 1. Similarly to above, we denote the average equilibrium expected earnings of an employed worker of type $\ell \in \{m, f\}$ as $T(\Lambda, v, \theta_i, \ell)$.

COROLLARY 2. *If $G_m(\cdot)$ first-order stochastically dominates $G_f(\cdot)$ then $\frac{\mathbb{E}_{G_f}[T(\Lambda, v, \theta_i, f)]}{\mathbb{E}_{G_m}[T(\Lambda, v, \theta_i, m)]}$ converges monotonically to 1 as Λ converges to 1 for all v .*

In words, this result says that the average earnings of employed women is rising relative to the average earnings of employed men as transparency increases, and reaching full transparency completely equalizes earnings. The proof of this result follows from Theorem 1. When $G_m(\cdot)$ first-order stochastically dominates $G_f(\cdot)$, it is possible to pair up every f type worker with an m type worker with a higher outside option. Formally, let $\mu : [0, 1] \rightarrow [0, 1]$ define for each f type worker i an m type worker j such that $\theta_j \geq \theta_i$ and $\mu(\theta_i) \neq \mu(\theta_{i'})$ for any $i \neq i'$. We know by Theorem 1 that $\frac{T(\Lambda, v, \theta_i)}{T(\Lambda, v, \theta_j)}$ converges monotonically to 1 in Λ , which implies that the average earnings of each gender also converges monotonically to 1 in Λ .

However, if transparency spreads through word of mouth, men’s and women’s differing propensities to gossip about wages may affect who receives what information, and hence who renegotiates. Empirically, we find evidence that a permissive communication environment affects men’s and women’s wages differently. In TaskRabbit, men receive raises above their initial bids more often than women on average when workers are co-located, and not when they are separated. In Col. 1-3 of Table VIII we regress receipt of any raise on the interaction between female and co-location. While co-location increases the likelihood of a raise for men by more than 13% across all specifications, for women the likelihood is statistically smaller than half this magnitude. We also show in Col. 4-7 that conditional on a raise women also renegotiate to equality. This is consistent with survey evidence that men are more informed about the wages of their peers than women.⁴¹ In a survey we conduct presenting vignettes of TaskRabbit workers in a co-located job, participants (of both genders) believe men are more likely than women to ask and discover a co-worker’s wage (Figure B1).

We make simple adjustments to the model to capture differences in rates of wage communication across genders. Let $\alpha_m > \alpha_f > 0$ be the rates at which men and women “speak about wages,” respectively. Let the arrival rate of information for a worker of gender $\ell \in \{m, f\}$ be $\alpha_\ell \lambda$. Then

$$\Lambda_\ell = \frac{\alpha_\ell \lambda}{\rho + \delta + \alpha_\ell \lambda} \quad \text{for } \ell \in \{m, f\} \quad \text{and } \lambda \in [0, \infty) \quad (7)$$

⁴¹Babcock and Laschever (2003) find that women are less informed of the market value of their work than men and are less likely to negotiate. Hall and Krueger (2012) (Table 4) find that women are roughly half as likely as men to bargain for wages. Goldfarb and Tucker (2011) show that women are more private about their pay than men.

This communication heterogeneity causes the de facto arrival rate of information of men to be greater than that of women, that is, $\Lambda_m - \Lambda_f \geq 0$ for all λ . We plot $\Lambda_m - \Lambda_f$ as a function of λ for arbitrary parameters in Figure VII. $\Lambda_m - \Lambda_f$ is initially increasing but converges to 0.

PROPOSITION 6. *Let λ_c solve $\frac{\alpha_m}{\alpha_f} = \frac{(\rho+\delta+\alpha_m\lambda)^2}{(\rho+\delta+\alpha_f\lambda)^2}$. $\Lambda_m - \Lambda_f$ is strictly increasing in λ for all $\lambda < \lambda_c$ and strictly decreasing for all $\lambda > \lambda_c$. As $\lambda \rightarrow \infty$, $\Lambda_m - \Lambda_f \rightarrow 0$.*

Compare the effects of moving from full privacy to some $\lambda > 0$. When λ is low, information transmission between workers rarely happens through word of mouth. If men are more likely to speak about wages, they disproportionately benefit from low levels of transparency. However, when λ is high, men gain less compared to women because all workers learn information quickly. In extreme cases of transparency in which the firm posts a price ($\lambda = \infty$) any communication advantage men have completely disappears as information arrival is not based on word of mouth transmission.

An important implication of these findings is that pay gossip between co-workers may exacerbate pay discrepancies. Advocates of gender equality may reconsider whether promoting worker discussions about pay will further their goal.

VIII. HETEROGENOUS AND UNKNOWN WORKER QUALITIES

Until now we have assumed that relative productivities are known by workers and the firm. This is based on our empirical setting in which many of the tasks are simple and standardized tasks, with binary performance outcomes, complete or incomplete. Here we discuss our findings in contexts where there may be significant heterogeneity in worker productivities, and where these differences are only observed by the firm. We find that equilibrium expected wages, firm profit, and the hiring rate level under full transparency are unchanged from the base model when productivity differences are observed. Transparency does not fully equalize wages, due to different worker productivities. But, earnings between high productivity workers and low productivity workers are compressed relative to the case where productivity differences are observable.

Suppose there are two types of workers, with productivities v and V , respectively. v and V are drawn independently from the same distribution. Each worker is equally likely to have productivity type v or V . The firm knows each worker's productivity type, but workers do not. Worker productivity is independent of outside option. The firm initially sets two maximum wages, \bar{W}_V and \bar{W}_v , and the rest of the game is as before.

Under full privacy, the equilibrium outcome mirrors that of the base model: $\bar{W}_v = v$ and $\bar{W}_V = V$, and workers make the same initial offers. 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. Upon meeting the firm, each worker sees the wages of other employed workers, and, in particular, will observe \bar{W}_v and \bar{W}_V . Without loss of generality, we assume that $v < V$ so that $\bar{W}_v \leq \bar{W}_V$. Worker i 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 $\bar{W}_V \geq \theta_i$ and

$$\bar{W}_v \leq \frac{1}{2}\bar{W}_V + \frac{1}{2}\theta_i$$

When a worker requests \bar{W}_V the firm will reject her offer with probability $\frac{1}{2}$, which clearly causes a reduction in the hiring rate and firm profits compared to the 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 lock in 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 model where productivity differences are observable.*

This proposition implies that there is wage compression, as opposed to complete wage equalization. All employed, low-productivity workers will earn \bar{W}_v as the firm rejects all low productivity workers who offer 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 productivity workers 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!

The fact that $\bar{W}_v > \bar{w}(v)$ results in more low productivity workers being hired than if productivity differences were observable, and 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 losing 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 model with observable productivity differences.*

This analysis shows that transparency has similar effects even when relative worker productivities are not known. These results hinge on the assumption that workers correctly estimate their chances of being a high productivity worker. Recent research has argued that

some workers may be overconfident in their abilities compared to those of their peers (Hoffman and Burks, 2017). In our context, we can model worker confidence with parameter γ , which is the probability that each worker places on being type V upon seeing the wages of her peers. $\gamma = \frac{1}{2}$ is as above, where workers are neither over nor underconfident. Higher levels of γ represent more confidence.

PROPOSITION 9. *Firm profit under full transparency is decreasing in γ .*

Unobservable worker productivities coupled with high levels of overconfidence may lead transparency to reduce profits and the hiring rate below full privacy levels. As workers become extremely overconfident, i.e. γ converges to 1, the firm must either employ all workers at essentially the same wage, or only hire high productivity workers. Therefore, transparency may cause high turnover with high overconfidence.

IX. CONCLUSION

Pay transparency and norms about pay privacy have been in the political and popular spotlights, but their effect on wages, hiring rate, and profits are not well understood. Moreover, federal and local pay transparency regulations do not systematically address alternative work arrangements or online markets where platform design choices can dramatically affect transparency. We study the equilibrium effects of pay transparency in a market for short-term work by combining a dynamic bargaining model with an empirical analysis of panel data from an online labor platform, and a field experiment. Our analysis of equilibrium wages, hiring rate, and profits under greater pay transparency reveals consequences that are counterintuitive and economically large in a market for low-skill tasks.

We empirically assess pay transparency’s effects on negotiations in our administrative and experimental settings. In both, pay transparency improves the negotiating position of the employer, and this leads to wages that are approximately 25% lower. Employer profits are even larger as transparency raises the hiring rate. Increasing transparency has a stronger positive impact on the hiring rate for employers with low values for labor, and increases hiring for these employers by as much as 40%. We also show that transparency equalizes wages while increasing dispersion in worker surplus both through lowering demand for labor and 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 shifts worker surplus toward workers most capable of completing the job. If outside options represent previous wages, then transparency lessens disparities of opportunity in other labor markets.

Our finding that pay transparency is profitable for employers appears at odds with the commonly cited fact 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 posted wages are the equilibrium outcome under full transparency. Redefining transparency to include not only pay discussions on the job but also the choice of the employer to be forthcoming about a posted price at the time of hire 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 very high degree of transparency.⁴² This share is higher among jobs with a high degree of standardization, such as entry-level jobs and occupation-certified roles, as our framework predicts because these are precisely the jobs that are most conducive to informative spillovers about pay information between co-workers. Other studies place transparency at even higher levels.⁴³ While high transparency is beneficial during the initial hiring process by strengthening the employer’s bargaining position, employers may want to suppress wage gossip to limit morale responses to interpersonal comparisons across the organization, to protect wage information from competing firms, to avoid bad press from (perhaps justifiable) wage disparities, and overaggressive negotiations by overconfident workers.

The insights of our model can be adapted to evaluate many different ways of introducing (or limiting) pay transparency. For example, we can assess recent laws prohibiting employers from asking workers their past wages during the interview process. Our findings suggest that workers benefit from the disclosure of prior wages, as it strengthens their bargaining position in the same way that full transparency leads to posted prices, strengthening the bargaining position of employers. In equilibrium, all workers should voluntarily reveal previous wages due to unraveling.

We also discover through our analysis that the method of implementing transparency can be consequential. Some policies operate by protecting the rights of workers to discuss pay. We show that when pay information travels through gossip networks, encouraging pay discussions can exacerbate the gender pay gap because men and women may not access the same information at the same rate. In our empirical setting we see that transparency disproportionately leads to renegotiations among men, widening the gender pay gap. This may be an important consideration in how transparency policies are enacted.

⁴²In 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.

⁴³Niederle 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.”

REFERENCES

- John M. Abowd, Francis Kramarz, and David N. Margolis. High Wage Workers and High Wage Firms. *Econometrica*, 67(2):251–333, 1999.
- George A. Akerlof and Janet L. Yellen. The Fair Wage-Effort Hypothesis and Unemployment. *The Quarterly Journal of Economics*, 105(2):255–283, 1990.
- Avend Alæk, Peter Møllgaard, and Per B. Overgaard. Government-Assisted Oligopoly Coordination? A Concrete Case. *The Journal of Industrial Economics*, 45(4):429–443, 1997.
- Linda Babcock and Sara Laschever. *Women Don't Ask: Negotiation and the Gender Divide*. Princeton University Press, 2003.
- Gordon M. Becker, Morris H. DeGroot, and Jacob Marschak. Measuring Utility by a Single-Response Sequential Method. *Behavioral Sciences*, 9(3):226–232, 1964.
- Truman Bewley. *Why Wages Don't Fall During a Recession*. Harvard University Press, 1999.
- Francis Bloch and Harl Ryder. Two-Sided Search, Marriages, and Matchmakers. *International Economic Review*, 41(1):93–115, 2000.
- Hanna Brenzel, Hermann Gartner, and Claus Schnabel. Wage bargaining or wage posting? Evidence from the employers' side. *IZA DP No. 7624*, 2013.
- Emily Breza, Supreet Kaur, and Yogita Shamdasani. The Morale Effects of Pay Inequality. *Quarterly Journal of Economics*, *Forthcoming*, 2017.
- Kenneth Burdett and Melvyn G. Coles. Long-Term Partnership Formation: Marriage and Employment. *The Economic Journal*, 109(456):F307–F334, 1999.
- Aine Cain, Anaele Pelisson, and Shayanne Gal. 9 places in the US where job candidates may never have to answer the dreaded salary question again. <http://www.businessinsider.com/places-where-salary-question-banned-us-2017-10>, 2018.
- 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.
- Eric Cardella and Briggs Depew. Testing for the Ratchet Effect: Evidence from a Real-Effort Work Task. *Games and Economic Behavior*, *Forthcoming*.
- Hector Chade. Matching with Noise and the Acceptance Curse. *Journal of Economic Theory*, 129(1):81–113, 2006.
- Gary Charney and Peter Kuhn. Does Pay Inequality Affect Worker Effort? Experimental Evidence. *Journal of Labor Economics*, 25(4):693–723, 2007.
- 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.
- Zoe Cullen and Chiara Farronato. Outsourcing Tasks Online: Matching Supply and Demand on Peer-to-Peer Internet Platforms. *Mimeo*, August 2016.
- Stefano DellaVigna, John A. List, Ulrike Malmendier, and Gautam Rao. Estimating Social Preferences and Gift Exchange at Work. *NBER Working Paper 22043*, 2016.
- Alex Edmans, Xavier Gabaix, Tomasz Sadzik, and Yuliy Sannikov. Dynamic CEO Compensation. *Journal of Finance*, *Forthcoming*, 2016.
- 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.

- Merle Faminow and Bruce Benson. On the Implications of Reporting Retail Food Prices over Extended Periods. *The Journal of Consumer Affairs*, 21(1):40–69, 1987.
- Jack Fanning and Andrew Kloosterman. A simple experimental test of the Coase conjecture: fairness in dynamic bargaining. *Mimeo*, 2017.
- Diana Farrell and Fiona Greig. Paychecks, Paydays and the Online Platform Economy. *JPMorgan Chase & Co. Institute*, 2016.
- 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.
- 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.
- Ben Greiner, Axel Ockenfels, and Peter Werner. Wage transparency and performance: A real-effort experiment. *Economic Letters*, 111:236–238, 2011.
- Wulong Gu and Peter Kuhn. Centralization and Strikes. *Labour Economics*, 5(3):243–265, 1998.
- Marina Halac. Relational Contracts and the Value of Relationships. *American Economic Review*, 102(2):750–779, 2012.
- 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.
- Ariane Hegewisch, Claudia Williams, and Robert Drago. Pay Secrecy and Wage Discrimination. *Institute for Women’s Policy Research*, 2011.
- Mitchell Hoffman and Stephen V. Burks. Worker Overconfidence: Field Evidence and Implications for Employee Turnover and Returns from Training. *NBER Working Paper 23240*, 2017.
- Michihiro Kandori, George J. Mailath, and Rafael Robb. Learning, Mutation, and Long Run Equilibria in Games. *Econometrica*, 61(1):29–56, January 1993.
- Patrick Kline, Neviana Petkova, Heidi Williams, and Owen Zidar. Who Profits from Patents? Rent-sharing at Innovative Firms. *Mimeo*, 2017.
- David M. Kreps and Robert Wilson. Reputation and imperfect information. *Journal of Economic Theory*, 27(2):253–279, 1982.
- Kai-Uwe Kühn and Xavier Vives. Information Exchanges Among Firms and their Impact on Competition. *Office for Official Publications of the European Communities*, 1995.
- 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. *NBER Working Paper 20431*, 2015.
- 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.
- 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.
- Roger Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6:58–73, 1981.
- 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.
- Arvid Nilsson. Transparency and Competition. *Mimeo*, 1999.
- Chris Nosko and Steven Tadelis. The limits of reputation in platform markets: An empirical analysis and field experiment. *Mimeo*, 2015.
- 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. *Mimeo*, 2016.
- 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.
- Gary R. Siniscalco, Erin M. Connell, and Chad Smith. State Pay Equity Laws: Where a few go, many may follow. *Mimeo*, 2017.
- Jae Song, David J. Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter. Firming Up Inequality. *Mimeo*, 2018.
- Jean Tirole. From Bottom of the Barrel to Cream of the Crop: Sequential Screening with Positive Selection. *Econometrica*, Forthcoming, 2016.
- Donald M. Topkis. *Supermodularity and Complementarity*. Princeton University Press, 1998.
- Joel Watson. A General, Practicable Definition of Perfect Bayesian Equilibrium. *Mimeo*, 2017.
- Martin L. Weitzman. The ‘Ratchet Principle’ and Performance Incentives. *Bell Journal of Economics*, 11 (1):302–308, 1980.
- Steven R. Williams. Efficient Performance in Two Agent Bargaining. *Journal of Economic Theory*, 41: 154–172, 1987.

X. TABLES AND FIGURES

TABLE I: ROADMAP

TABLE I.A: MAPPING BETWEEN THEORY AND EMPIRICS					
Effect of transparency on:	Theoretical reference	Empirical reference	TaskRabbit or Experiment?	Choice of empirical setting	
1	Wage equality	Theorem 1	Tables IV, V; Figure I	Both	
2	Worker surplus equality	Corollary 1	Table V	Experiment	Worker surplus unobserved in TaskRabbit data.
3	Hiring	Theorem 2	Table VI	Both	
4	Surplus share/ average wages	Theorem 3	Table VII	Both	
5	Endogenous choice of transparency/ unraveling	Theorem 4	Figures IV, V	TaskRabbit	Employers do not select transparency level in experiment.
6	Gender wage gap	Corollary 2; Proposition 6	Table VIII	TaskRabbit	Not applicable in full transparency environment of experiment.

TABLE I.B: OVERVIEW OF EMPIRICAL COMPARISON GROUPS

Comparison groups	Comparable along observables?	Exogenously determined?	Sample size (# jobs)	Additional information	
1	Multi-worker jobs that are co-located (same time and location) vs. separated (same time, different locations) (TaskRabbit)	Yes	Partially due to nature of work	690	Restricted to private auctions jobs, and job categories with minimum share 20% of both types of jobs (Tables IV, V, VIII). ⁱ
2	Posted price vs. private auction jobs (TaskRabbit)	No ⁱⁱ	No	>500k	Table VI, Fig. IV, V.
3	Price mentioned in job description vs. price not mentioned (TaskRabbit)	Yes ⁱⁱⁱ	No	>500k	Restricted to private auction jobs. Sub-sample of 19,827 with hourly wage prices only (Table VII).
4	Common public chat forum vs. private chat forums, wages negotiated with manager (Field experiment)	Yes	Yes	90 managers	Negotiable treatment arms used for Tables V, VI, VII. Non-negotiable treatment arms used in Fig. VI.

Notes: (i) Replication of Table IV results with all categories included can be found in Table B12. (ii) See Table B10 for further details about differences between private auctions and posted price jobs. (iii) See Table II for more details. Small differences are statistically significant due to the large sample size.

TABLE II: SUMMARY STATISTICS, TASKRABBIT

Multi-worker jobs	Separated (mean)	Co-located (mean)	T-Stat (Sep. – Co.)	Separated (N)	Co-located (N)
Initial bids received (\$)	52.2	52.3	-0.03	578	950
Amount below max bid (%)	0.35	0.31	1.02	578	950
Number hired	2.60	2.50	0.98	291	399
Received bids (Gini)	0.19	0.20	0.25	291	399
Chosen bids (Gini)	0.08	0.07	1.33	291	399
Positive ratings (Gini)	0.31	0.29	1.10	291	399
Length of description	732	756	-0.72	291	399
LDA 1 (eg. “would like”, “hoping”, “would be great”)	1.06	3.61	-2.37	291	399
LDA 2 (eg. “give feedback”, “product review”)	3.25	2.02	0.87	291	399
LDA 3 (eg. “assistance”, “shopping”, searching”)	5.29	6.61	-0.94	291	399
LDA 4 (eg. “planning”, “evening”)	6.60	7.52	-0.82	291	399
LDA 5 (eg. “moving”, “pickup”, “delivery”)	8.55	8.70	-0.14	291	399
LDA 6 (eg. “documents”, “photos”)	3.20	3.46	-0.18	291	399

(a) Co-located jobs are defined as jobs that require more than one person to be on-site concurrently to complete the job, according to the employer who enters number of slots to fill, the location and the time. Separated jobs also require multiple workers at a given time but not on-site together, for example marketing in different neighborhoods. We include in the main sample all multi-worker jobs in categories with at least 20% separated and co-located job types, resulting in 15 job categories. Results are robust to using all job categories, which we use in Appendix tables for additional statistical power. The explanatory variable “amount below maximum bid” is equal to $(bid^{max} - bid^i)/bid^i$ for person i . The Gini coefficient is a non-parametric measure of dispersion, $G = \frac{1}{n} \left(n + 1 - 2 \frac{\sum_{i=1}^n (n+1-i)y_i}{\sum_{i=1}^n y_i} \right)$ for person i indexed in non-decreasing order $y_i < y_{i+1}$ with outcome y_i hired alongside $n - 1$ co-workers. Gini equal to 0 indicates that all workers have equal allocations and 1 if a single worker is allocated everything. LDA (Latent Dirichlet Allocation) text components are included to compare the content of job descriptions. LDA is an unsupervised bag-of-words machine learning method. The algorithm assumes every document contains a mixture of topics and every topic is a mixture of words. Topics are assumed to have a Dirichlet prior, which intuitively means that each job description is likely to only contain a small set of topics and each topic frequently uses a small set of words. Words can belong to multiple topics and LDA estimates the probability that a particular word belongs to a particular topic. Using the probabilities that given words belong to given topics in combination with the Dirichlet priors, it is then possible to construct the probabilities that each document contains each topic.

Single-worker jobs	Price Mention (mean)	No Mention (mean)	T-Stat (Sep. – Co.)	Price Mention (N)	No Mention (N)
Initial bids received (\$)	55.5	64.5	14.08	>100,000	>100,000
Time between post and deadline (days)	3.35	3.37	-0.54	>100,000	>100,000
Length of description	770	769	0.18	>100,000	>100,000
LDA 1 (eg. “would like”, “hoping”, “would be great”)	2.73	2.48	1.45	>100,000	>100,000
LDA 2 (eg. “give feedback”, “product review”)	3.13	2.99	0.68	>100,000	>100,000
LDA 3 (eg. “assistance”, “shopping”, searching”)	10.56	10.56	0.01	>100,000	>100,000
LDA 4 (eg. “planning”, “evening”)	4.98	4.98	0.01	>100,000	>100,000
LDA 5 (eg. “moving”, “pickup”, “delivery”)	6.80	7.04	-1.34	>100,000	>100,000
LDA 6 (eg. “documents”, “photos”)	8.36	8.31	0.37	>100,000	>100,000

(b) “Price Mention” refers to the presence of text in the job description that indicates a range for acceptable bids. Initial bids are submitted after reading the job post with or without price mentions and we consider initial bids as an outcome of interest in Table VII. Observation counts obscured at the request of TaskRabbit.

City markets	N	Mean	Stand. Dev.	25th Perc.	75th Perc.
Share posted price	417	0.43	0.10	0.37	0.46
Market age (months)	417	16.6	12.2	6.1	26.1
Vacancy fill rate	417	0.46	0.11	0.41	0.53
Price (\$)	417	56.1	8.69	52.0	61.0

(c) City-month level summary statistics about the 19 active cities of TaskRabbit. Market age is the number of months since Task Rabbit became active in the city. Vacancy-fill rate is the share of posted job positions that are matched to a worker during the time period, since the platform started in 2008 until June 2014. posted price refers to jobs assigned where the employer chooses a price that is publicly posted on the platform, and the first worker to accept the price will be assigned to the job. Price is the total amount transacted for the job, including the platform fee.

TABLE III: SUMMARY STATISTICS, EXPERIMENT

		<u>Main Treatment Arms</u>			<u>Additional Treatment Arms</u>			
		Not Transparent	Transparent	T-Statistic	Budget= \$9		Non-Negotiable	
		(mean)	(mean)	(diff)	(mean)	(diff)	(mean)	(diff)
Workers	Age	37	37	-0.36	35	0.73	35	2.13
	Share female	0.56	0.61	-0.85	0.40	1.29	0.49	2.35
	Share some college	0.96	0.92	1.41	0.90	0.04	0.92	0.66
	<i>N</i>	<i>163</i>	<i>120</i>		<i>131</i>		<i>506</i>	
Managers	Age	37	39	-0.80	36	0.41	36	1.45
	Share female	0.57	0.68	-1.04	0.47	1.26	0.52	0.75
	Share some college	0.92	0.86	0.92	0.84	2.68	0.92	-0.20
	<i>N</i>	<i>53</i>	<i>37</i>		<i>42</i>		<i>161</i>	

Notes: The leftmost three columns describe the sample assigned to negotiate in either a split or common chat room with a \$5 manager budget. The rightmost four columns describe secondary treatment arms: (1) a higher manager budget (\$9) and (2) a sample who were not allowed to renegotiate after placing the initial bid. The “diff” columns report the t-statistic of a test of the null hypothesis that the difference in means between the treatment arms is 0, or between the treatment arm and the pooled main sample is 0. For a comparison of subgroups assigned to transparency or non-transparency within the secondary treatment arms, please refer to Appendix Table B8. Experimental findings are robust to excluding demographics controls.

TABLE IV: BONUSES AMONG CO-LOCATED WORKERS, TASKRABBIT

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.	Any Raise (Yes = 1)			Raise (% Above Bid) Any Raise		
Amt. under top bid (%) ⁽ⁱ⁾	0.0129 [0.0180]	0.0291 [0.0183]	0.0305 [0.0259]	0.894*** [0.0992]	0.967*** [0.0694]	0.959*** [0.131]
Years experience		0.0133 [0.0362]	0.0453* [0.0273]		-0.145 [0.111]	0.0423 [0.0866]
Effective percent positive overall		0.00250 [0.00968]	-0.000512 [0.00766]		0.0211 [0.0263]	0.0219 [0.0398]
Effective percent positive in cat.		0.0445 [0.0313]	-0.00647 [0.0316]		0.0353 [0.0704]	-0.0730 [0.167]
No. reviews		-0.0110 [0.0106]	-0.00707 [0.00778]		0.0481 [0.0302]	0.0274 [0.0308]
No. reviews in category		0.0567 [0.0653]	0.00209 [0.0519]		-0.387** [0.192]	-0.0811 [0.223]
Mean rating		-0.00649 [0.0739]	0.0616 [0.0971]		-0.190 [0.348]	0.118 [0.364]
Mean rating in category		0.113** [0.0465]	0.0777 [0.0671]		-0.0249 [0.267]	-0.689 [0.615]
No rating		-0.0320 [0.365]	0.246 [0.487]		-0.837 [1.651]	0.477 [1.799]
No rating w/in category		0.570** [0.225]	0.396 [0.342]		-0.248 [1.354]	-3.319 [3.035]
No. workers (log)		-0.108*** [0.0272]			-0.124 [0.157]	
Mean bid (log)		0.0118 [0.0265]			-0.215 [0.170]	
Constant	0.215*** [0.0201]	-0.314 [0.289]	-0.447 [0.510]	0.275*** [0.0600]	2.259 [1.727]	3.089 [3.718]
> 1 hour overlap		✓	✓		✓	✓
Job FE			✓			✓
P-value Test: $H_0: \beta^{(i)} = 1$				0.289	0.631	0.755
Mean Outcome	0.22	0.22	0.22	0.43	0.45	0.45
Std. Dev. Outcome	0.41	0.41	0.41	0.39	0.43	0.43
Observations	950	727	727	208	158	158
Clusters	410	318	318	128	96	96
R^2	0.001	0.051	0.848	0.669	0.736	0.983

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. The dependent variable equals one if the particular worker earns more than their agreed to bid, and 0 otherwise. Col. 4-6 are restricted to those workers that do receive more than their 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. Standard errors are clustered at the level of the job. “> 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. Note, in later tables we use > 1 hour to refer to duration, not necessarily overlap.

TABLE V: DISPERSION IN FINAL WAGES AND WORKER SURPLUS

Experimental Evidence						
	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.	Final Pay (Gini)		Hourly Wage (Gini)		Worker Surplus (Gini)	
Transparent (Public chat)	-0.0232*** [0.00707]	-0.0215*** [0.00802]	0.0282 [0.0363]	0.0233 [0.0376]	0.0838* [0.0495]	0.101* [0.0557]
Constant	0.0232*** [0.00707]	0.114 [0.0935]	0.144*** [0.0219]	0.0933 [0.215]	0.100*** [0.0252]	0.0646 [0.228]
Mgr. Char.		✓		✓		✓
Mean D.V.	0.015	0.015	0.155	0.155	0.131	0.131
Std. Dev. D.V.	0.040	0.040	0.150	0.150	0.195	0.195
Obs.	74	74	74	74	74	74
R^2	0.0794	0.171	0.00831	0.122	0.0434	0.0771

Notes: Each model is estimated by OLS. An observation is a multi-worker local task. Standard errors are clustered at the task level. We present experimental evidence based on the sample of employees who are assigned to treatments where negotiation is permissible and the manager budget is fixed at \$5. The dependent variable in Col. 4-5 is the dispersion in final pay, measured as the total pay the worker receives. The dependent variable in Col. 6-7 is dispersion in 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, all multiplied by number of pages eventually completed. Since the Gini coefficient on final pay and hourly wage is only defined if more than one worker transacts with the manager, we exclude from this table employers who transact with one or zero workers. In this restricted sample 27 of 27 managers assigned to the transparency treatment pay completely equal wages to all workers, and 29 of 47 pay completely equal wages to all workers in the privacy treatment. Manager characteristics in the experimental setting include gender, age, age squared, managerial experience, and years of formal education.

TABLE VI: EFFECT OF TRANSPARENCY ON HIRING, BY VALUE OF LABOR

		TaskRabbit			Experimental Evidence			
		(1)	(2)	(3)	(4)	(5)	(6)	
Dep. Var.		Vacancy filled (yes = 1)			Hired worker (yes = 1)			
Low Value	× Transparent (Employer inc.) (Posted price)	0.0203 [0.0139]	0.0229* [0.0128]	0.0409** [0.0200]	Low Value × Transparent (Mgr budget) (Public chat)	0.227** [0.0926]	0.240** [0.0972]	0.258*** [0.0976]
Low Value		-0.0210* [0.0114]	-0.0172* [0.0101]	-0.0187 [0.0140]	Low Value	-0.439*** [0.0504]	-0.437*** [0.0497]	-0.441*** [0.0503]
Transparent		0.170*** [0.0112]	0.132*** [0.0103]	0.143*** [0.0164]	Transparent	-0.0695 [0.0539]	-0.0620 [0.0589]	-0.0818 [0.0594]
Employer Char.			✓	✓	Manager Char.		✓	✓
Category FE			✓	✓	Worker Char.			✓
City FE, Month FE, Mkt. Age				✓				
Mean Dep. Var.		0.650	0.650	0.650	Mean Dep. Var.	0.622	0.622	0.622
Observations		>20k	>20k	>20k	Observations	414	414	414
Clusters		>5k	>5k	>5k	Clusters	132	132	132
R^2		0.0352	0.0664	0.0745	R^2	0.136	0.153	0.178

Notes: Each model is a linear probability model estimated by OLS. In Col. 1-3, an observation is a job posting on TaskRabbit. The sample is restricted to jobs posted by household employers with observable earnings. The dependent variable is equal to 1 if the job posting is matched to a worker before it expires on TaskRabbit. The primary explanatory variable, lower value, is an indicator equal to 1 if the employer earns less than the median earning household on the platform in each city. We mask the number of observations at the request of TaskRabbit. Col. 4-6 are from our experimental data, the sample assigned to treatment groups where workers were allowed to negotiate with the employer. An observation is a worker. The dependent variable is equal to 1 if this worker is hired, and 0 otherwise. Employer characteristics from TaskRabbit include age, age squared, gender and time since joining the platform. Manager and worker characteristics from the experiment include age, age squared, gender and education (4 levels). Standard errors in all columns are clustered at the level of the employer or manager.

TABLE VII: HIGHER PROFITS UNDER TRANSPARENCY

Dep. Var.	TaskRabbit			Experimental Evidence			
	(1)	(2)	(3)	(4)	(5)	(6)	
	All Bids (log \$)	Winners' Pay (log \$)	Winners' Hourly Pay (log hourly wages)	Wages Employed (log \$)	Productivity (arc sinh)	Profits (arc sinh)	
Transparent (Job mentions price)	-0.0959*** [0.0136]	-0.136*** [0.0142]	-0.0778*** [0.0281]	Transparent (Public chat)	-0.257** [0.0958]	0.175 [0.241]	0.617*** [0.224]
Exp. on platform (Days)	0.378** [0.175]	0.493*** [0.123]	1.048*** [0.393]	High School	0.327** [0.142]	-0.578 [0.397]	-0.746 [0.530]
No. ratings in category	-0.991*** [0.0859]	-0.512*** [0.172]	1.921*** [0.308]	Some College	0.220* [0.123]	-1.022*** [0.305]	-1.349*** [0.259]
No. ratings overall	1.656*** [0.0546]	1.694*** [0.123]	0.361 [0.400]	College	0.550*** [0.107]	-0.733*** [0.266]	-0.890*** [0.271]
Mean rating in category	0.0476*** [0.00733]	0.0792*** [0.0199]	0.0230 [0.0363]	Post Graduate	0.342** [0.129]	-0.375** [0.183]	-1.355*** [0.309]
Mean rating overall	-0.00385 [0.0114]	-0.0168 [0.0330]	-0.0591 [0.0796]	Age	0.0679 [0.0460]	0.0249 [0.0668]	0.0481 [0.0685]
Category FE	✓	✓	✓	Age Sqr.	-0.000767 [0.000539]	-0.0000685 [0.000802]	-0.0000637 [0.000832]
Worker FE	✓	✓	✓	Female	0.0382 [0.104]	-0.415* [0.249]	-0.454** [0.209]
Performance Covariates	✓	✓	✓				
Mean Dep. Var.	3.37	3.69	3.02	Mean Dep. Var.	2.883	1.677	0.666
Std. Dev. D.V.	0.946	0.854	0.624	Std. Dev. D.V.	0.337	0.907	1.062
Observations	>100k	>100k	19827	Observations	47	90	90
R^2	0.277	0.381	0.607	R^2	0.380	0.083	0.236

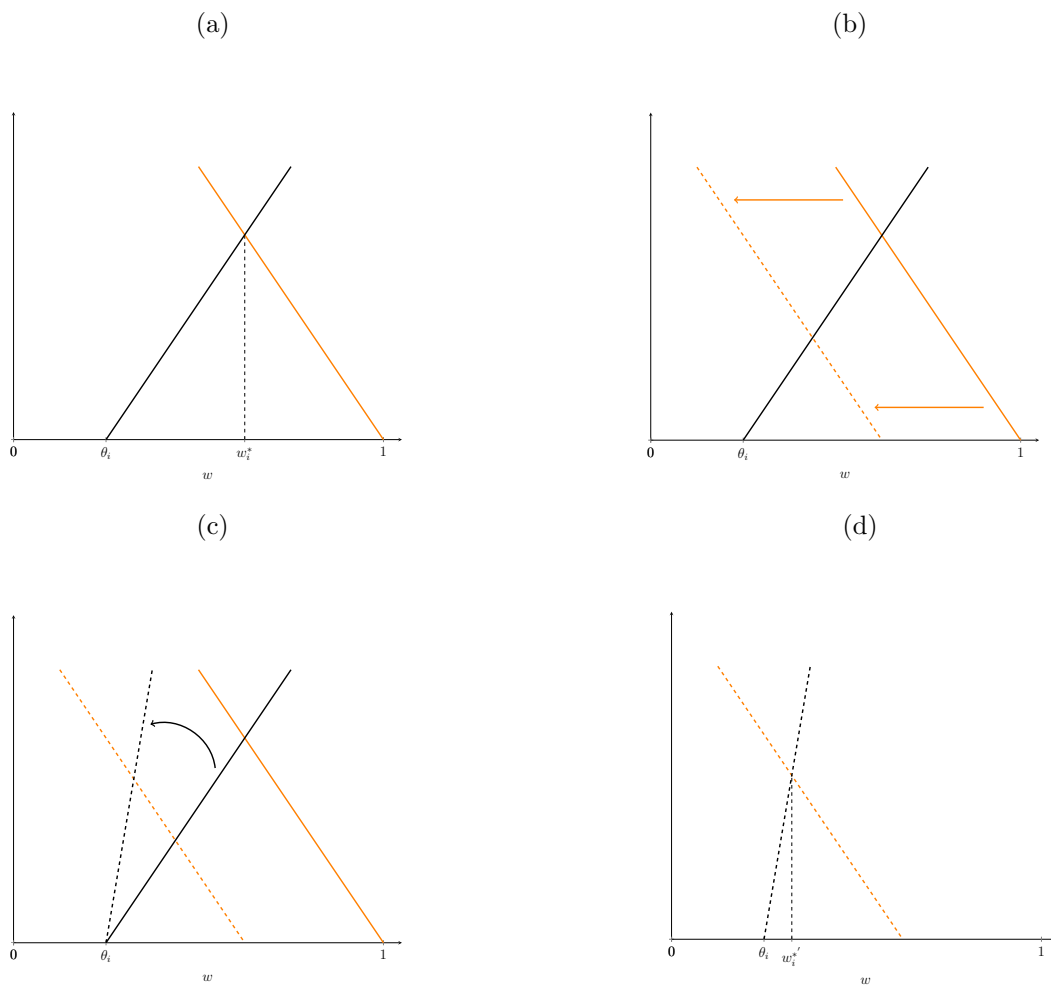
Notes: Each model is estimated by OLS. In Col. 1-3 an observation is a worker-bid on TaskRabbit. The dependent variable is the log bid in Col. 1, and final pay in Col. 2 and 3. In Col. 3 we restrict our attention to the small sample of jobs that solicit hourly wage bids rather than piece rate. Transparency in TaskRabbit refers to 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. In Col. 4-6, our sample is from our experimental setting, those who were assigned to treatment arms where negotiation with the manager was permissible and the manager budget was equal to \$5 per page transcribed. An observation is a worker (Col. 4) or manager (Col. 5-6) in our experiment. The dependent variables (moving left to right) are log wages, productivity or number of transcription pages completed in total, and inverse hyperbolic sine of profits a manager earns. We use inverse hyperbolic sine transformation to accommodate 0 outcomes. The transformation $\text{arcsinh}(x)$ 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. \(2017\)](#) Appendix D, page 65. Education is a categorical variable with reference category, some high school. Covariates refer to the worker in Col. 4, and the manager in Col. 5-6. Robust standard errors are in square brackets. We do not reveal observation counts for aggregate activity on the platform at the request of TaskRabbit.

TABLE VIII: GENDER GAP IN LIKELIHOOD OF RAISE RISES WITH PAY TRANSPARENCY, TASKRABBIT

	(1)	(2)	(3)	(4)	(5)	(6)
Dep.Var.	Any Raise (Yes = 1)			Final Pay (% Above Bid) Any Raise		
Transparent (Co-located)	0.187*** [0.0291]	0.191*** [0.0319]	0.137*** [0.0331]			
Transparent × Female	-0.141*** [0.0390]	-0.118*** [0.0436]	-0.0906** [0.0445]			
Amt. under top bid (%)				0.759** [0.305]	0.906*** [0.224]	0.902*** [0.226]
Female × Amt. under top bid (%)				-0.00437 [0.344]	0.0214 [0.241]	0.0552 [0.241]
Female	0.0154 [0.0259]	0.0123 [0.0267]	0.0640** [0.0271]	-0.149 [0.0998]	-0.137 [0.102]	-0.188* [0.111]
Years experience		0.0116 [0.0214]	0.0153 [0.0210]		-0.146* [0.0782]	-0.171** [0.0817]
Effective percent positive overall		-0.00177 [0.00601]	-0.000771 [0.00583]		0.0473 [0.0318]	0.0502 [0.0341]
Effective percent positive in cat.		0.0235 [0.0151]	0.00909 [0.0147]		-0.0762 [0.0814]	-0.0923 [0.0780]
No. reviews		-0.00905 [0.00653]	-0.00842 [0.00627]		0.0366 [0.0308]	0.0474 [0.0316]
No. reviews cat.		0.0454 [0.0424]	0.00822 [0.0392]		-0.317* [0.173]	-0.329* [0.179]
Mean rating		-0.0269 [0.0524]	-0.0255 [0.0526]		0.109 [0.342]	0.127 [0.341]
Mean rating in category		0.104*** [0.0325]	0.0847** [0.0334]		-0.413 [0.351]	-0.484 [0.364]
No rating		-0.121 [0.257]	-0.108 [0.259]		1.731 [1.484]	1.867 [1.536]
No rating w/in cat.		0.490*** [0.158]	0.411** [0.160]		-2.192 [1.730]	-2.525 [1.798]
No. workers (log)			-0.0673*** [0.0220]		-0.0918 [0.102]	-0.0890 [0.135]
Mean bid (log)			-0.0181 [0.0202]		-0.324* [0.177]	-0.378* [0.212]
Constant	0.0785*** [0.0179]	-0.330* [0.199]	-0.252 [0.226]	0.330*** [0.0785]	3.207* [1.825]	3.137* [1.895]
> 1 hour		✓	✓		✓	✓
Category FE			✓			✓
Mean Outcome	0.17	0.16	0.16	0.41	0.41	0.41
Std. Dev. Outcome	0.37	0.36	0.36	0.38	0.39	0.39
Observations	1,590	1,280	1,280	392	317	317
Clusters	658	537	537	233	178	178
R ²	0.046	0.064	0.121	0.446	0.587	0.605


Notes: Each model is estimated by OLS. An observation is a worker-bid in a multi-worker job on TaskRabbit using a private auction. The sample is restricted to those with self-reported gender information (96% of workers). Col. 1-3 are linear probability models, and are restricted to categories with (at least 20%) separated and co-located jobs. The dependent variable equals 1 if the particular worker earns more than their initial bid, and 0 otherwise. Col. 4-6 are restricted to those workers who receive more than their bid in all co-located jobs. The dependent variable is the size of the raise, as percent above bid. The primary explanatory variable is an indicator for whether the job entails co-location of workers (a measure of partial transparency) and its interaction with the gender of the worker. Standard errors are clustered at the level of the job.

FIGURE II: EFFECTS OF INCREASING Λ ON WORKER AND FIRM STRATEGIES




Notes: (a) By Equation 3, worker i picks an initial wage offer w_i^* that equalizes $\frac{w-\theta_i}{1-\Lambda}$ (the black, upward sloping line) and $\frac{1-\bar{F}(w)}{f(w)}$ (the orange, downward sloping line). (b) The **demand effect** of increasing Λ from 0 to $\frac{3}{4}$ reduces \bar{w} for each v , shifting $\frac{1-\bar{F}(w)}{f(w)}$ to the left. (c) The **supply effect** of increasing Λ from 0 to $\frac{3}{4}$ increases the slope of the function $\frac{w-\theta_i}{1-\Lambda}$. (d) The **supply** and **income effects** combine to reduce the initial wage offer of worker i to $w_i^{*'}$ when Λ increases.

FIGURE III: EXPECTED DIFFERENCE IN EQUILIBRIUM WAGES T PERIODS AFTER ENTERING THE MARKET



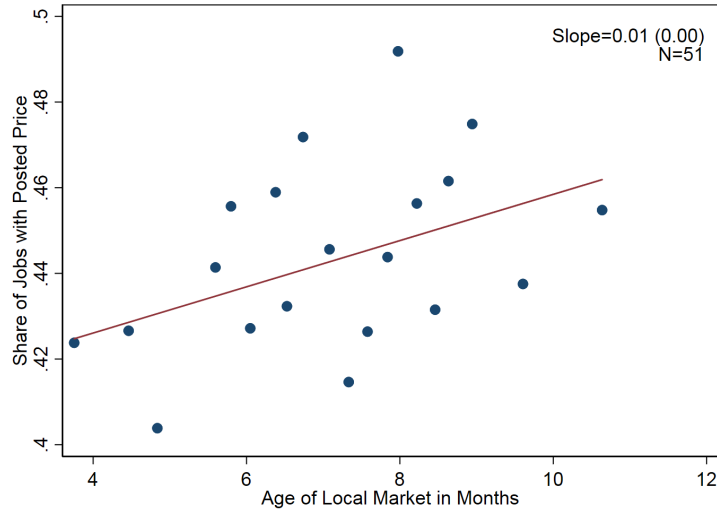
Notes: Figure III shows the expected difference in the wage of two workers, i and j , T periods after each has entered the market when $\theta_i > \theta_j$. The dashed (black) curve represents this difference when $\Lambda = \frac{1}{2}$, $\rho + \delta = 1$, $r = s = 1$, and the solid (orange) curve represents this difference when $\Lambda = \frac{1}{3}$, $\rho + \delta = 1$, $r = s = 1$. Although the dashed curve is initially above the solid one, the two curves satisfy a single-crossing condition in t .

FIGURE IV: POSTED PRICE BY CITY



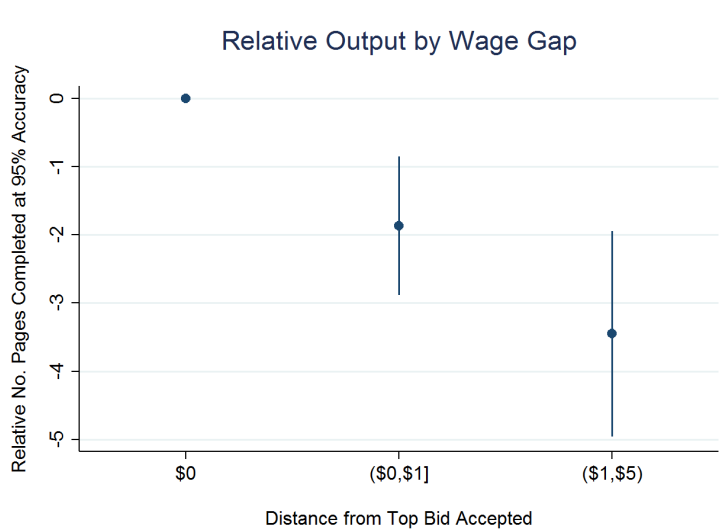
Notes: Figure IV plots the age of each TaskRabbit market (horizontal axis) and the proportion of posted price jobs in each market (vertical axis) at the end of our data sample in June, 2014. Older markets appear to be associated with a higher proportion of posted price jobs. TaskRabbit entered Boston in 2008 nearly one year before the start of our data sample. In our analysis we treat the Boston market as if it started at the same date as our data sample, but in reality, there are many months that we do not observe. Boston was where the company was founded and the technology evolved considerably during the first year when it was only available in one city. “Virtual” refers to tasks that are completed by workers online.

FIGURE V: POSTED PRICE AND MARKET AGE




Notes: Figure V is a binscatter plot of a balanced panel of 9 local markets active for longer than a year and each city-month is an observation. The bin-means are adjusted for city-level fixed effects, total size of market (measured by the number of job listings), and the share of job postings in each of the 8 largest categories. We also include in Appendix Table B11 a full table of regression results with a more complete set of controls and observations.

FIGURE VI: PRODUCTIVITY CONSEQUENCES OF TRANSPARENCY WHEN PAY IS NON-NEGOTIABLE



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. 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 in distance, and \$1 upwards. The data include 150 managers, and 267 workers who bid less than or equal to the \$5 budget.

FIGURE VII: DIFFERENCE IN DE FACTO ARRIVAL RATE OF WAGE INFO BETWEEN GENDERS AS A FUNCTION OF λ



Notes: Figure VII plots λ against the difference in the de facto rate of information arrival between men and women. This difference is initially increasing in λ , but after a single peak, it decreases toward 0. Parameters used: $\alpha_m = 2$, $\alpha_f = 1$, $\rho + \delta = 1$.