

NBER WORKING PAPER SERIES

EQUILIBRIUM EFFECTS OF PAY TRANSPARENCY

Zoe B. Cullen
Bobak Pakzad-Hurson

Working Paper 28903
<http://www.nber.org/papers/w28903>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2021

We are extremely grateful to Aviv Nevo and three anonymous referees. One referee in particular shifted the direction of this paper, and we are extremely appreciative of their contribution. 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 also thank Mohammad Akbarpour, Lint Barrage, Jose Maria Barrero, Doug Bernheim, Ben Brooks, Eric Budish, Gabriel Carrol, Arun Chandrasekhar, Kalyan Chatterjee, Isa Chaves, Ken Chay, Bo Cowgill, Piotr Dworzak, Jack Fanning, Chiara Farronato, Laura Gee, Matt Gentzkow, Avi Goldfarb, Bob Hall, Gregor Jarosch, Scott Kominers, Maciej Kotowski, Vijay Krishna, Jon Levin, Shengwu Li, Erik Madsen, Davide Malacrino, Alejandro Martinez, Paul Milgrom, Muriel Niederle, Ricardo Perez-Truglia, Kareen Rozen, Ilya Segal, Isaac Sorkin, Jesse Shapiro, Chris Stanton, Bryce Millet Steinberg, Takuo Sugaya, Neil Thakral, Catherine Tucker, Christopher Snyder, Emmanuel Vespa, and Alistair Wilson and seminar attendees at Stanford, UCLA Anderson, NYU Stern, HBS, Dartmouth, Brown, USC, UCSD, UCSB, Penn State, Michigan SI, AEA, SIEPR, EC'19, SOLE, and NBER SI (Digitization and Personnel) for helpful comments and suggestions. We are indebted to Chloe Lee, and especially Dylan Balla- Elliott and Diego Gentile for excellent research assistance. This research was supported by the Center for Comparative Studies in Race and Ethnicity and the B.F. Haley and E.S. Shaw Fellowship through SIEPR. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Zoe B. Cullen and Bobak Pakzad-Hurson. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Equilibrium Effects of Pay Transparency
Zoe B. Cullen and Bobak Pakzad-Hurson
NBER Working Paper No. 28903
June 2021
JEL No. C78,D82,D83,J31,M52

ABSTRACT

The public discourse around pay transparency has focused on the direct effect: how workers seek to rectify newly-disclosed pay inequities through renegotiations. The question of how wage-setting and hiring practices of the firm respond in equilibrium has received less attention. To study these outcomes, we build a model of bargaining under incomplete information and test our predictions in the context of the U.S. private sector. Our model predicts that transparency reduces the individual bargaining power of workers, leading to lower average wages. A key insight is that employers credibly refuse to pay high wages to any one worker to avoid costly renegotiations with others under transparency. In situations where workers do not have individual bargaining power, such as under a collective bargaining agreement or in markets with posted wages, greater transparency has a muted impact on average wages. We test these predictions by evaluating the roll-out of U.S. state legislation protecting the right of workers to inquire about the salaries of their coworkers. Consistent with our prediction, the laws lead wages to decline by approximately 2% overall, but declines are progressively smaller in occupations with higher unionization rates. Our model provides a unified framework to analyze a wide range of transparency policies, and reconciles effects of transparency mandates documented in a variety of countries and contexts.

Zoe B. Cullen
Rock Center 210
Harvard Business School
60 N. Harvard
Boston, MA 02163
and NBER
zcullen@hbs.edu

Bobak Pakzad-Hurson
64 Waterman Street
Providence, RI 02912
bph@brown.edu

A data appendix is available at <http://www.nber.org/data-appendix/w28903>

I. INTRODUCTION

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

Our paper studies both the direct and indirect impacts of pay transparency policies on wage negotiations and renegotiations. We combine a dynamic wage-bargaining model with an event-study analysis of the enactment of U.S. state laws increasing pay transparency. We find that increasing transparency decreases workers’ de facto bargaining power. Because of this, average wages are lower under higher transparency. Lower average wages distinguishes our bargaining framework from the partial equilibrium effect on which pay transparency policies are commonly premised: without an equilibrium response through bargaining, we would expect wages to rise after transparency is introduced, as transparency’s direct effect allows low-wage workers to negotiate higher pay.

To introduce the mechanisms we study in this paper, we present two scenarios. First, suppose a worker learns that a colleague with the same job is earning significantly more than she is. She reasons that her employer must be willing to pay a higher wage for the work she is doing, and will therefore seek to renegotiate her wage.³ The employer has an interest in preventing such negotiations because they raise wages. While the solution seems that the firm should restrict the flow of wage information to workers, this is only half of the story.

Consider instead the initial wage negotiations of the same worker at a firm with full pay transparency. Even before agreeing to an initial wage, she observes the salary of the highest-paid worker; the worker can demand that wage, but cannot negotiate any raise. The

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.”

²The policies in the U.S. range from punishments for employers that prevent employees from discussing salaries (Siniscalco et al., 2017) to requirements that employers inform prospective employees of the range of salaries currently paid to workers (Pender, 2017). EU policies range from mandating full disclosure by companies over a threshold number of employees, like in Germany, to publication of salary statistics such as the mean, median and gap between men and women, as in the case of the U.K. and Denmark (European Commission; International Labour Organization).

³Caldwell and Harmon (2019) present empirical evidence that workers renegotiate their wages on the job after learning new wage information.

fact that wages are public allows the employer to credibly say, “If I give you a higher salary, I’ll have to give everyone else a raise too, and I just can’t afford that.” Under pay secrecy, the worker might have been skeptical of such a claim and bargained more aggressively, but due to transparency, the worker grasps the (true and costly) ramifications of asking for more than the current maximum wage earned by others. Therefore, full transparency leads to an unintended side effect: if workers all get the same wage and cannot negotiate this wage upward, then the firm gets the power to set the wage. To maximize its profit, the firm acts like a monopsonist and sets a relatively low wage. Thus transparency increases the de facto bargaining power of the employer, becoming the enforcement mechanism for a low wage.

Transparency, in this way, resembles best-price guarantees (also known as most-favored-customer agreements). Best-price models argue that agreements to rebate existing customers if prices fall in the future allows the firm to maintain higher prices (Butz, 1990; Cooper and Fries, 1991). The firm can commit to refusing marginal buyers because of the cost of rebates to previous consumers,⁴ and empirical evidence supports these findings (Scott Morton, 1997a,b). Our study contains this effect, but the forces at play differ significantly because we study a setting of (two-sided) incomplete information.

We present a baseline model of continuous-time wage negotiation in the presence of transparency, in which a continuum of workers individually bargain with a firm. We extend this model in many ways, including to a setting in which there are many firms and workers search among them for jobs. As our key findings are preserved in these extensions, we investigate our simple model to shed light on the forces at play. A key point of departure from much of the previous wage-bargaining literature which assumes complete information (see, e.g. Hall and Milgrom (2008)) is that the potential surplus from a worker-firm match is initially unknown for both parties.⁵ Both sides of the market alter their bargaining strategies in response to transparency, and the interaction of these responses drives our predictions.

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

⁴See Bagwell and Wolinsky (2002) for a series of pricing models that demonstrate how the future observability of durable decisions can lead to increased commitment power.

⁵Manning (2005) finds that employers’ private information can affect bargaining and lead to high levels of inequality in imperfectly competitive markets, which we believe justifies this departure.

⁶As argued by Larsen and Zhang (2018), double auctions serve as useful empirical approximations to real-world bargaining in settings with private information. For example, Larsen (2021) empirically estimates a double auction model using data from post-auction bargaining over used cars. Our empirical setting provides a rare opportunity to study details of bargaining for wages, rather than goods, in settings with

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

Workers stochastically learn the wages of their peers over time. An independent arrival process reveals the entire profile of wages at the firm to a particular worker, and the arrival rate $\lambda \geq 0$ characterizes the level of transparency: with higher levels of transparency, wage information arrives more quickly. Under full transparency ($\lambda = \infty$), workers learn the wages of their peers immediately; under full privacy ($\lambda = 0$), they never do. Workers renegotiate their wage when information arrives by specifying a new minimum acceptable wage.

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

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

We show that increasing transparency has the same equilibrium effect as decreasing worker bargaining power. Formally, we show that the set of equilibria under high transparency and high worker bargaining power is identical to the set of equilibria with lower transparency and lower worker bargaining power. Regardless of workers’ nominal bargaining power, full transparency grants the firm full de facto bargaining power, as the firm

private information.

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

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

commits to, and workers immediately observe, a maximum wage; and when workers have no bargaining power, the firm makes a single take-it-or-leave-it offer, implying that all workers immediately know the wages of their peers, whether or not they directly observe these wages through another channel. As workers' individual bargaining power shrinks, the effects of higher transparency (continuously) become more muted as workers are less able to exploit differences in their outside options to secure heterogeneous wages. In markets where workers have no individual bargaining power, such as markets with posted wages or markets where wages are set by a collective or union, transparency will not affect the equilibrium outcome.

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

The amount of firm bargaining power has a non-monotonic effect on employment because employment is maximized when bargaining power is shared between workers and the firm. When bargaining power is highly skewed, either workers act like monopolists, making high wage demands that are often rejected; or the firm acts like a monopsonist, committing to low wages that deter high-outside-option workers from considering work at the firm. Granting either the firm or workers all of the bargaining power minimizes expected employment.¹⁰

We study these equilibrium effects empirically in an event study framework. Since 2004, 18 U.S. states and the District of Columbia have enacted policies specifically aimed at facilitating communication about pay between coworkers.¹¹ The policies impose punishments for employers that retaliate against workers who disclose their wages or inquire about the wages of coworkers. We refer to these as “Right of Workers To Talk” (ROWTT) policies. While we cannot directly observe how these laws affected communication between coworkers, we have suggestive evidence of their effect on the existence of firm-level policies forbidding

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

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

¹¹As early as 1935, a clause in the National Labor Relations Act (NLRA) established worker rights to discuss pay in the spirit of facilitating collective bargaining; however, these protections were described in very general terms, “protecting concerted activity,” and violators did not face punitive damages, which led to the critique of the NLRA as a “toothless tiger” (Green, 2014).

wage discussion or disclosure. [Hegewisch et al. \(2011\)](#) and [Sun et al. \(2021\)](#) survey workers about such non-disclosure policies in 2010 and 2017, respectively. Among states that enact ROWTT legislation between 2010 and 2017, the share of workers subject to non-disclosure mandates falls 42%, twice the national rate.¹² The decline of formal non-disclosure policies is only one indicator of changes in the workplace day-to-day. Anecdotally, these legal protections led to the prolific use of Google spreadsheets where employees enter their own salaries for coworkers to see. The Philadelphia Inquirer identifies local instances of shared salary spreadsheets among museum workers, lawyers, ad agencies, and journalists [Reyes \(2019\)](#).

We study the effect of ROWTT policies enacted between 2004 and 2016, at which time two similar federal transparency policies are enacted.¹³ Our identifying assumption is that the precise timing of ROWTT enactment during a sixteen-year period is uncorrelated with wages and employment dynamics. We empirically verify this assumption of parallel trends by examining how wages, employment and additional labor outcomes evolve in each state leading up to the enactment of ROWTT.

We use data from the American Community Survey (ACS) to track wages and employment. Between 2000 and 2016, the ACS surveyed over 4 million individuals living in states that enact ROWTT laws. These surveys result in a repeated cross-section that contains information on employment, sector, occupation, industry, U.S. state of work, demographic characteristics and wages earned up to \$250,000. We focus on private-sector workers and merge in unionization rates at the occupation level from the Current Population Survey.

Corroborating our main theoretical finding, wages fall as pay transparency rises. In the year following the enactment of the ROWTT, wages fall by 2.2% and are statistically different from 0, reaching 2.6% by the third year after the policy. Our analysis is robust to many alternative specifications, including allowing for heterogeneous treatment effects across cohorts; the interaction-weighted estimator proposed by [Sun and Abraham \(2020\)](#) yields a precisely estimated wage decline of 2.8% three years following the enactment of ROWTT.

Our theory posits that wage declines may come from two sources related to reduced worker bargaining power: a reduction of the wages of workers and, under some circumstances, reduced employment of high outside option workers. Private sector employment levels in the ACS are stable following the enactment of ROWTT: confidence intervals each year include

¹²The share of employees forbidden from discussing salaries falls from 16.2% of surveyed employees to 9.4% of surveyed employees. We would not expect rates to drop to zero even after the passage of ROWTT for several reasons; first, employees responding to the survey may not be aware of recent changes their employer made; secondly, employers may not have responded to the law either due to ignorance or as part of a long tradition of using gag laws to intimidate employees. Rates of non-disclosure policies decline in states that have not enacted an ROWTT by 2017 fall 22.5% from 20% to 15.5% over this time span, which also includes the 2016 federal mandate extending ROWTT to all federal contractors.

¹³Because of this concurrent policy timing, we follow [Donohue III and Heckman \(1991\)](#) and use neither event data nor outcome data after 2016.

zero change in employment and exclude changes greater than 0.4%. The counterfactual effect on wages holding fixed the composition of employees in each education-by-gender cell yields nearly equivalent results. We further conduct a decomposition exercise to liberally bound the effect of job composition changes on wage declines; we estimate less than half of the reduction in wages could stem from composition changes.

To investigate the interaction of individual bargaining power and ROWTT policies, we analyze the effect of unionization on wage declines. Our model predicts that wage declines are largest when workers have high *individual* bargaining power, and are muted when workers bargain collectively. To measure individual worker bargaining power, we order occupations by the share of the workforce that is covered by a union. We re-estimate our event-study analysis allowing for independent dynamic effects in occupations with above- and below-median unionization rates.¹⁴ We find that occupations with below-median unionization share incur a 3.2% drop in wages following the enactment of an ROWTT by the third year after the policy while occupations with above-median unionization share experience a 1.5% decline in wages over the same period. The 1.6% difference in effect is statistically significant at conventional levels. Indeed, the effect progressively grows larger moving from the upper quartile of occupations by unionization rate to the lowest quartile.

We caution that our estimation of heterogeneous treatment effects by unionization rates may not imply a causal relationship; there likely exist differences between these groups that we cannot control for using ACS worker and occupation characteristics.¹⁵ However, as we detail below, these estimates fall in line with similar figures across labor markets in five countries.

To test our theoretical predictions in other settings, we extend our framework to additional forms of pay transparency policies. Our baseline analysis primarily focuses on word-of-mouth transmission of wage information, however, our bargaining framework applies to other policy measures that reveal signals about the salaries of coworkers, including the publication of salary ranges or average wages in a position. Theoretically, we show that both of these policies shift equilibrium *de facto* bargaining power away from workers in the same way as increased frequency of word-of-mouth transparency. This allows us to extend our main results to a broader class of policies.

¹⁴The median rate of unionization across occupation in the private sector is 7%; with a five-fold difference between unionization rates below and above the median, or a 12 percentage point gap in means between groups.

¹⁵Some of these differences are indeed a part of our theoretical framework and are the very reason we predict effects of transparency to be more muted among unionized workers. Specifically, our theory predicts that, as a result of lower individual bargaining power under collective agreements, employees are less able to exploit their different outside options through individual negotiations - hence, we expect wage variation to be greatly reduced in unionized workplaces prior to ROWTT and is partly why transparency will have a muted effect.

Our framework reconciles heterogeneous impacts of pay transparency policies documented by researchers independently studying different labor markets and geographies. A comprehensive review of studies that evaluate the effect of pay transparency on wages of all workers yields nine studies in eight papers, spanning five countries (Baker et al., 2019; Bennedsen et al., 2020; Blundell, 2021; Böheim and Gust, 2021; Duchini et al., 2020; Gulyas et al., 2020; Mas, 2017; Obloj and Zenger, 2020). These papers are focused on wage compression, but as a secondary outcome, some studies report a significant decrease in wages (Baker et al., 2019; Bennedsen et al., 2020; Blundell, 2021; Duchini et al., 2020; Obloj and Zenger, 2020), while others find a tight null effect on wages (Baker et al., 2019; Böheim and Gust, 2021; Gulyas et al., 2020). We contextualize these findings by proxying the degree of individual bargaining power workers have in each study based on the share of workers covered by a union or collective bargaining agreements. We find that in markets with high individual bargaining power (low unionization rates), pay transparency mandates serve to lower average wages. In markets with low individual bargaining power (high unionization rates), pay transparency mandates have little to no effect on wages of either men or women. We conduct a mixed-effects meta-regression analysis and find that a 10 percentage point decrease in share of labor market unionization is associated with a 0.17 percentage point larger decrease in wages. The differential effects persist for at least four years.

One alternative framework for evaluating pay transparency centers on the psychic costs of discovering that one earns less than her peers (Bracha et al., 2015; Breza et al., 2018; Card et al., 2012; Cohn et al., 2014; Cullen and Perez-Truglia, 2019; Dube et al., 2019; Perez-Truglia, 2020). These papers generally find that upon learning peers receive higher wages, workers lower effort and intend to leave the firm if wages are not adjusted, as suggested in Akerlof and Yellen (1990). These findings are not at odds with our bargaining model; indeed, the presence of a morale cost upon learning a higher-paid peer’s wage may microfound a worker’s ability to renegotiate her wage—there is a credible threat to the firm’s profit if it does not raise wages. Even if workers could not instigate renegotiations, a profit-maximizing firm would unilaterally raise wages to avoid this morale cost (Eliaz and Spiegler, 2013), leading to a similar equilibrium effect. In principal, an employer could be constrained or naive, hence allowing for declining morale to translate into lower productivity that subsequently results in lower wages. Empirically there is limited evidence of productivity declines when employers have time to adjust wages (Bennedsen et al. (2020); Cullen and Pakzad-Hurson (2021); Duchini et al. (2020)).

Our finding that pay transparency lowers worker bargaining power raises the question of why we do not observe more firms voluntarily selecting high levels of transparency. Indeed, well-cited studies find that the majority of firms attempt to limit pay transparency (Hegewisch et al., 2011; McCarthy, 2018). This finding is consistent with our model predic-

tions due to challenges in committing to a particular level of pay transparency. Consider a firm that promises a worker to post all wages on a company message board, to be updated whenever wages adjust, in order to reach a favorable initial wage agreement. After the negotiation, the firm would have a profitable deviation to simply neglect updating the message board, or worse, systematically under-report wages. In equilibrium, our model predicts that when the firm cannot contract on the level of transparency, it will select full secrecy regardless of the profit-maximizing exogenous level of transparency.¹⁶ This points to an important role for transparency legislation, and in particular, policies that promote credible information, such as protections for co-workers to circulate salary information.¹⁷

When considering who will advocate for pay transparency, it is important to keep in mind that not all firms are predicted to benefit from transparency. Increasing transparency raises firm profits in expectation, but is not beneficial for every firm type in our model, nor is it harmful for every worker. A firm with a low value for labor can benefit from transparency because under pay privacy, workers demand a larger premium over their outside options (supply effect), and these demands exceed the value of this firm, leading to low employment. Workers with low outside options stand to gain relatively more from an increase in transparency because they benefit more than others from renegotiation, which can result in higher wages for these workers.

The fact that low-outside option workers benefit relatively more from transparency leads to within-firm wage compression. But this should not be mistaken for a statement about overall wage inequality in an economy. Our model predicts that the magnitude of the wage decline will depend on the difference between the expected value of labor drawn by the firm and the expected outside option drawn by a worker – these relative values likely differ across industries based on labor market tightness and the extent of competition between firms. Moreover, our bargaining theory is fundamentally premised on the notion that workers are comparable and learn about their value to the firm upon seeing coworkers wages. Our predictions *will* change for specialized jobs: in settings where the firm’s value for a worker is specific to the individual, either because the worker is highly specialized or essential to retain, the bargaining effects from transparency do not apply.¹⁸ Therefore, transparency likely has

¹⁶As we discuss in Section II.E, workers suffer a collective action problem and each worker will always seek out wage information to the extent allowed by the firm.

¹⁷Third-party websites such as Glassdoor and PayScale have also arisen to solve this commitment problem. Similarly, online labor markets have the advantage of preserving a history of posted wages, hence providing the commitment device. Some companies have independently developed tools to achieve credible salary transparency. For example, Buffer posts a salary spreadsheet and a salary calculator on its website as a function of observables so that employees can test that their own salary is correctly computed and are able to verify the pay of others. Of course, technologies like this may also come with other costs, such as rigidity in the wage formula. We extend our model to accommodate the possibility that a firm has access to such technology.

¹⁸Other channels such as social comparisons may become first-order in these cases, such as in the market

the largest effect in industries when roles are standardized and workers are substitutable.

The remainder of the paper is organized as follows. Section II lays out our model and presents our main theoretical findings. Section III describes our event-study analysis of US state ROWTT laws. Section IV presents a formal analysis of transparency policies investigated in other studies, and demonstrates how wage effects are predicted in these settings by the level of individual bargaining power. Section V concludes. Omitted proofs, additional empirical results, and extensions are contained in the Appendix.

II. MODEL

II.A. Setup

Time is continuous, and is indexed by $t \in \mathbb{R}_+$. There is a single firm in the economy, and a unit measure of workers I . Each worker $i \in I$ has a private outside option $\theta_i \stackrel{iid}{\sim} G[0, 1]$, which is the flow payment i receives when unemployed.¹⁹ The firm has a constant-returns-to-scale production function. We assume for now that productivity of labor is common across all workers: $v \sim F[0, 1]$, and is known only to the firm. (We discuss the case in which workers are heterogeneously productive in Section II.D.) All agents exponentially discount the future at rate δ , are risk neutral, and seek to maximize discounted expected flow payments. We assume that F and G are twice continuously differentiable with densities f and g , respectively. We also assume agents have strictly increasing virtual reservation values, i.e. $\theta + \frac{G(\theta)}{g(\theta)}$ is strictly increasing in θ and $v - \frac{1-F(v)}{f(v)}$ is strictly increasing in v .

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

We model transparency as the random arrival of information about current wages. At time $t \geq 0$ each matched worker observes the set of wages the firm pays to employed workers, W_t , according to an independent Poisson arrival process with (commonly known) rate $\lambda \in [0, \infty) \cup \{\infty\}$, where we take $\lambda = \infty$ to mean that the process arrives at every time $t \geq 0$.

for high-earning CEO and executive pay (Faulkender and Yang, 2013; Mas, 2016; Schmidt, 2012a).

¹⁹There is a known measurability issue with the assumption of a continuum of i.i.d. random variables (Judd, 1985). A solution is to assume that worker outside options are drawn “almost” i.i.d. in the sense of Sun (2006). This solves the measurability issue and has the intuitive and intended property that the distribution of realized outside options is given by the same function $G(\cdot)$.

For convenience, we assume that $W_0 = \{\bar{w}\}$.²⁰ Therefore, higher λ corresponds to more transparency.

When wage information arrives to worker i at time t , i and the firm renegotiate i 's wage using the same bargaining protocol: i submits a new offer $w_{i,t}$ and she remains employed and receives a flow wage of $(1-k)w_{i,t} + k\bar{w}$ for all $t' \geq t$ if and only if $w_{i,t} \leq \bar{w}$. The interpretation is that observing the exact wages of coworkers allows a worker to force the firm back to the bargaining table, perhaps due to the (unmodeled) threat of legal action that is only credible if the worker can provide direct proof that a peer is being paid more.

The timing of the stage game is as follows for every worker i who has not yet been permanently unmatched from the firm: First, at each time $t \geq 0$ worker i learns W_t independently with arrival rate λ . Second, if either $t = 0$ or $t > 0$ and worker i observes W_t , then i bargains with the firm according to the protocol laid out above.

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

In Appendix G.1 we expand our model to allow workers to search for work across multiple firms, and show that many results are robust to this extension. We also consider extensions in which: the firm can accept or reject offers individually instead of picking a single \bar{w} , the firm can reject worker renegotiation offers at cost but without permanently unmatched, and the firm negotiates with workers under different bargaining protocols. Our main results extend to these settings, and further details are available upon request.

II.B. Equilibrium

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

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

²⁰Without this assumption, all workers under full transparency (and a measure zero set of workers for any $\lambda > 0$) face an openness issue of wanting to renegotiate wages at the earliest time $t > 0$. It is possible to deal with this issue as in [Simon and Stinchcombe \(1989\)](#): suppose workers can only renegotiate every $\frac{1}{N}$ periods, $N > 1$. Define a worker's payoff in continuous time as the limiting value as $N \rightarrow \infty$. Using this definition, even if a worker observes nothing at $t = 0$, her payoff under full transparency is equivalent to the case in which she receives a wage of \bar{w} for all $t \geq 0$. For ease of notation, we continue with the simplifying, if unrealistic, assumption that $W_0 = \{\bar{w}\}$.

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

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

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

A3 assists in tractability. It also removes equilibria in which workers and the firm pool on a predetermined wage from consideration.²¹

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

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

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

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

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

²¹Leininger et al. (1989) suggest similarities between the set of continuous equilibria and a set of discontinuous equilibria in static double auctions, and so we do not believe this to be a conceptually limiting constraint. We discuss the connection of our game to static double auctions below.

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 F that each worker i and the firm respectively solve

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

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

where

$$\Omega = (1 - \Lambda)k + \Lambda \quad (3)$$

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

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

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

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

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

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

Ω is submodular in Λ and k , implying that a fixed increase in transparency is more impactful the smaller k , the nominal bargaining power of the firm. When $\Lambda = 1$, $\Omega = 1$ implying that at full transparency, the nominal bargaining power k does not affect the

equilibrium outcome. Similarly, when the firm has all of the nominal bargaining power, i.e. $k = 1$, the equilibrium outcome is constant in the level of transparency Λ . This matches our earlier, intuitive descriptions: under full transparency, all workers learn \bar{w} immediately and secure this wage if it is higher than their outside option, regardless of k . The firm therefore “posts” \bar{w} knowing that all employed workers will receive this wage. When $k = 1$ the firm makes an initial TIOLI offer \bar{w} to each worker, and all workers with a lower outside option will be employed by the firm at this wage. Wages are “transparent” to workers in that all workers know the firm pays a common wage.

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

Relaxing Bargaining Assumptions

Collective bargaining

We have assumed that workers have individual bargaining power, but in many labor markets, wages are set through collective bargaining. We augment our model to allow for a union to negotiate wages, and workers choose to work at the negotiated wage or to consume their outside options. In this case, workers have no individual bargaining power. As a result, we show transparency will not affect the equilibrium outcome.

Consider our base model, but instead suppose that \bar{w} is a prevailing wage negotiated by a union and the firm. We model the bargaining process to arrive at \bar{w} as follows: prior to the hiring of individual workers at $t = 0$ a positive-measure set of workers $I' \subset I$ is exogenously selected to constitute union representatives. The union seeks to maximize the average surplus of union representatives; therefore, the union can represent the interests of a diverse set of workers if I' contains workers with a wide range of outside options θ_i .

The union and the firm negotiate \bar{w} prior to the hiring of workers using the same double-auction protocol we have studied; the firm offers \bar{w}' and the union offers w' . $\bar{w} = (1 - K)w' + K\bar{w}'$ for $K \in [0, 1]$ if and only if $w' \leq \bar{w}'$. All union representatives $i \in I'$ work at flow wage \bar{w} if an agreement is reached ($w' \leq \bar{w}'$). Note that when $K = 0$, our protocol reflects Leontief’s standard “monopoly union” model, and when $K > 0$, our protocol is more akin to the standard efficient bargaining model (McDonald and Solow, 1981).

Without observing \bar{w} , individual workers $i \in I \setminus I'$ bargain as before, with $k = 1$. That is, each worker i who makes a wage offer $w_{i,t}$ at time t is employed if and only if $w_{i,t} \leq \bar{w}$. If she is employed, she receives a flow wage of $(1 - k)w_{i,t} + k\bar{w} = \bar{w}$. All workers must make an initial wage offer $w_{i,0}$. We refer to this as the *union bargaining model*.

Similarly to in our base model with $k = 1$, all workers $i \in I \setminus I'$ have a weakly dominant strategy to set $w_{i,0} = \theta_i$, implying that all such workers i with $\theta_i \leq \bar{w}$ are hired and receive

flow wage \bar{w} . As workers lack individual bargaining power (i.e. $k = 1$) there is no scope for renegotiation. The following result follows from this observation:

REMARK 1. *For any $I' \subset I$ and any $K \in [0, 1)$, the set of equilibrium outcomes of the union bargaining model is constant in λ .*

At-will renegotiation

Our base model also assumes that workers cannot rebargain until pay information arrives. If workers can rebargain at will, any $k > 0$ will have the same equilibrium outcome as $\Omega = 1$ since all employed workers back out \bar{w} at $t = 0$ and immediately renegotiate by offering this amount (sidestepping the technical issue as in Footnote 20). When $k = 0$, workers do not learn \bar{w} until the arrival of pay information. Interestingly, we show that even when $k = 0$, no employed worker will ever renegotiate her wage in equilibrium until the first moment information arrives. The key step in proving this result is showing that a worker does not learn exploitable information about \bar{w} if her initial offer is accepted. Any worker strategy that says “offer w when initially hired at time 0 and offer $w' > w$ at time $t > 0$ if I have not learned the wages of my coworkers” is not optimal, because if offering w' at time t improves the expected utility of the worker, she would be even better off offering w' at time 0.²² We formalize and prove the following result in the appendix.

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

II.C. Main Results

The equilibrium bargaining strategies of workers and the firm are interdependent for any $\Omega \in (0, 1)$. Workers decide how aggressively to make initial offers depending on how the firm sets \bar{w} , while the firm sets \bar{w} as a function of how aggressively the workers make initial offers. While there exists a unique equilibrium when $\Omega \in \{0, 1\}$, Satterthwaite and Williams (1989) show that there exists a continuum of equilibria satisfying Equations 4 and 5 for $\Omega \in (0, 1)$. This set lacks natural ordering, limiting the possibility for general claims about the entire set of equilibria. However, experimental evidence in Radner and Schotter (1989) suggests that equilibria in which w_i^* and \bar{w} are linear functions of θ_i and v , are focal and most likely to be played in practice. We similarly observe linear worker wage offers in an experimental setting (Cullen and Pakzad-Hurson, 2021).

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

To focus our attention on linear equilibria, we restrict attention to a two-parameter family of power law distributions of worker outside options and firm values. We show that this family—given in Equation 6—admits a unique linear equilibrium for any $\Omega \in [0, 1]$.²³ In this section we study the properties of the linear equilibrium, and analyze the effects of increasing transparency.

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

As r increases, v is on average lower and as s increases, θ is on average higher. Therefore, increasing r or s reduces the average surplus from employment. We define a linear equilibrium below and show that distributions of this type admit a unique linear equilibrium. This result extends work by Chatterjee and Samuelson (1983) who show existence of a linear equilibrium when F and G are uniform, corresponding to the case in which $r = s = 1$. We note again that these distributional restrictions are unnecessary for our analysis if one is only interested in comparing $\Omega = 0$ to $\Omega = 1$ (i.e. full secrecy with $k = 0$ to full transparency).

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

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

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

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

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

²³The approach of making parametric assumptions to ensure linear equilibrium is common. One recent example on CEO pay is Edmans et al. (2012). Power law distributions are commonly observed in economic situations such as ours, including worker income and firm productivities. See Gabaix (2009, 2016) for details.

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.

Fix $k < 1$. We compare the earnings of workers i and j who are hired in equilibrium under both of two transparency levels, $\Lambda' < \Lambda''$, so we do not confound employment effects of increasing transparency.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 . Second, higher transparency decreases the expected time it takes before both workers renegotiate to \bar{w} , reducing dispersion of their earnings as $\bar{w} - w_j^* > \bar{w} - w_i^*$. The first effect increases the initial wage gap between i and j , however, the latter effect dominates in the long run, leading to more equalized expected earnings, regardless of δ .

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

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

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

Division of Surplus Increasing pay transparency increases the expected profits of the firm, decreases average worker surplus, and lowers average discounted wages. The **demand**

²⁴The restriction that workers be hired under both transparency levels is necessary, as we show in Theorem 3, because increasing transparency can increase the hiring rate. A previously unemployed, high outside option worker may find employment only when transparency is increased. To make this point concrete, take some small $\epsilon > 0$ and consider increasing transparency from Λ' to $\Lambda'' = \Lambda' + \epsilon$, such that more workers are employed in equilibrium under Λ'' . In Appendix F we show that w_i^* and \bar{w} are continuous in Λ and so the expected earnings of any worker j hired under both transparency regimes is barely affected by an ϵ increase in transparency. However, a worker i who over-negotiates at level Λ' receives her outside option θ_i for her entire duration in the market, while if she manages to find employment at the firm under Λ'' her average total earnings will be greater than, and bounded away from, θ_i (as she always asks for a premium $w_i^* - \theta_i > 0$). But note that $\theta_i > \theta_j$, so the total earnings of i and j are not equalized by increased transparency.

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

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

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

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

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

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

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

Ω^* is decreasing in both $\mathbb{E}(v)$ and $\mathbb{E}(\theta)$.²⁵ As $\mathbb{E}(v)$ converges to 0 full transparency becomes close to employment maximizing, and as $\mathbb{E}(\theta)$ converges to 1 full privacy becomes close to employment maximizing. For intuition, we return to Proposition 5. As $\mathbb{E}(v)$ decreases, the firm's markdown $v - \bar{w}$ is likely to be small regardless of Ω . Therefore, increasing transparency does not greatly reduce the number of workers with $\theta_i < \bar{w}$. But by increasing transparency, workers will shade down their initial offers w_i^* , reducing the number of workers who over-negotiate. Similarly, as $\mathbb{E}(\theta)$ increases, most workers offer small premia $w_i^* - \theta_i$ regardless of Ω . Increasing transparency has little effect on these premia, but instead discourages the firm from setting a large markdown.

An increase in transparency increases hiring only for sufficiently low value firm types. Consider an increase in transparency leading to an increase in Ω' to $\Omega'' > \Omega'$. The submodularity of the ex-post hiring rate in v and Ω means that the firm hires more workers under Ω'' than Ω' if and only if v is below a particular threshold. It also implies that the value of Ω that maximizes the ex-post hiring rate is weakly decreasing in v (Topkis, 1998). These comparative statics on ex-post hiring also hold for ex-post social surplus. In fact, the ex-post maximizer of the hiring rate also maximizes ex-post social surplus. Because each employed worker earns a wage weakly greater than her outside option, in equilibrium each employed worker increases social surplus by $v - \theta_i > 0$, implying that social surplus is strictly increasing in the hiring rate. Therefore, Ω'' increases ex-post social surplus if and only if v is below some threshold.

II.D. Other Transparency Processes

Other pay transparency policies may not directly promote word-of-mouth communication of individual wages. Instead, these policies could reveal average wages, average wage gaps across groups, or salary ranges.

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

We show that such policies have similar equilibrium effects as increasing word-of-mouth transparency, as studied in our base model. The central insight is that all of these policies, including word-of-mouth information sharing, are equilibrium objects that reveal information about the firm’s willingness to pay for a position. As we show below, the policies we study increase workers’ information about the maximum wage they can receive, which in turn affects (re)negotiations, triggering the supply and demand effects.

Average Salary and Gender Pay Gap Disclosure

We make the following changes to our model: first, suppose the information arrival process does not reveal individual wages, but rather reveals average wages. Second, suppose that information reaches all workers at the same $t \geq 0$. As before, let λ be the arrival rate of wage information. This process reflects the passage of a policy requiring the firm to disclose its average wage to the workforce.

By **A3**, both \bar{w} and w_i^* are strictly increasing in v and θ_i , respectively. As w_i^* is independent of \bar{w} , there is a one-to-one relationship between average wage (prior to information arrival) and \bar{w} . Therefore, upon observing the average wage, workers learn \bar{w} in equilibrium.

Workers similarly learn \bar{w} in equilibrium if the arrival process reveals the wage gap across groups. Suppose there are two types of workers, m (male) and f (female), and each worker i belongs to exactly one type. Let G_ℓ represent the distribution of outside options for type $\ell \in \{m, f\}$, and let $G(x) := qG_m(x) + (1 - q)G_f(x)$ for all $x \in [0, 1]$, where $q \in [0, 1]$ is the proportion of men in the market. If G_m dominates G_f in the likelihood ratio order, that is $\frac{g_m(x)}{g_f(x)}$ is strictly increasing in x , then as \bar{w} increases, the average wage of m types increases by more than the average wage of f types. Therefore, there is again a one-to-one relationship between the size of the wage gap and \bar{w} , implying that workers learn \bar{w} in equilibrium following the arrival of the information process.

The following result summarizes both of these results.

PROPOSITION 6. *Suppose the information process arrives to all workers simultaneously, with rate λ .*

1. *For any $\lambda \geq 0$, if the arrival process reveals the average wages of all workers, then the set of equilibrium outcomes satisfying **A1-A3** is identical to that in our base game.*
2. *For any $\lambda \geq 0$, if the arrival process reveals the gap between the average wages of m types and f types, then the set of equilibrium outcomes satisfying **A1-A3** is identical to that in our base game if G_m dominates G_f in the likelihood ratio order.*

Heterogeneous Worker Qualities

Until now we have assumed that all workers are equally productive. Here we discuss our findings in contexts where there may be significant heterogeneity in worker productivities. Information arrival about wages reveals the range of salaries offered to all workers. Therefore, we refer to transparency in this context as salary range revelation.

As a useful benchmark, suppose each worker $i \in I$ has a publicly observable type $\tau \in \mathcal{T}$ where \mathcal{T} is some countable set, each containing a positive measure of workers. Let $v_\tau \sim F_\tau[0, 1]$ be the productivity of type τ workers, which is known only to the firm. Each worker i of type τ also has a private outside option $\theta_i \stackrel{iid}{\sim} G_\tau[0, 1]$, which is the flow payment i receives when not matched to the firm. As before, F_τ and G_τ are twice continuously differentiable distributions with full support over $[0, 1]$ for all $\tau \in \mathcal{T}$. Our base model is a special case in which $|\mathcal{T}| = 1$, that is, all workers are equally productive with productivity $v \sim F[0, 1]$ and outside option distribution G .

As before, each worker i of type τ makes an initial wage offer $w_{i,\tau}^*$, and then an additional wage offer after observing the wages of peers. The firm picks a maximum wage $\bar{w}_\tau(v_\tau)$ for each type τ .

If all workers' types τ are known, and the information process additionally reveals whether $v_\tau \leq v_{\tau'}$ or $v_\tau \geq v_{\tau'}$ for any τ and τ' then the results of our paper go through within type. That is, each τ forms a different market. On equilibrium path, the firm picks the maximum wage for type τ workers $\bar{w}_\tau(v_\tau)$ as in the base model given distributions F_τ and G_τ , and each worker i of type τ picks an initial offer $w_{i,\tau}^*$ as in the base model given distributions F_τ and G_τ . Upon observing the salary range, each worker θ identifies the maximum wage associated with her productivity type, and offers that amount to the firm.

We therefore focus on the more interesting case in which workers are differentially productive, but do not know their own productivity type. Suppose there are two types of workers, with productivities v and V , respectively. v and V are drawn independently from the same distribution F . Each worker is equally likely to have productivity type v or V . The firm knows each worker's productivity type, but no worker observes her own productivity. To highlight mechanisms at play, we study the extreme case in which outside options are distributed independently of productivity so that workers do not receive a signal of their relative productivities.

Under full privacy ($\lambda = 0$), the equilibrium outcome mirrors that of the base model. Therefore, firm profits, the expected hiring rate, and wage dispersion are the same as before.

For tractability, we consider only the effects of full transparency ($\lambda = \infty$). At $t = 0$, each worker will observe the maximum wage the firm is willing to pay each type of worker, which we denote by \bar{W}_v and \bar{W}_V , respectively in the setting of unknown worker types. Without loss of generality, we assume that $v < V$ so that $\bar{W}_v \leq \bar{W}_V$. We note that if $k > 0$, this model

is again not meaningfully different than before. All workers i will initially offer $w_i^* = \bar{W}_v$. Based on this, each i will perfectly learn her productivity type: if her initial wage is equal to \bar{W}_v , she knows she has productivity type v and will not renegotiate, however, if her initial wage is greater than \bar{W}_v then she knows she has productivity type V and will (almost, see Footnote 20) immediately renegotiate to wage \bar{W}_V .

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

When a worker demands \bar{W}_V the firm will reject her offer with probability $\frac{1}{2}$, which clearly reduces the hiring rate and firm profits compared to the baseline model where there is no uncertainty about worker productivity. On the other hand, low outside option, productivity V workers will offer \bar{W}_v , meaning that the firm is able to hire some high productivity workers at low wages, increasing profits. We show that, because of this latter effect, the firm sets \bar{W}_v higher than it would have for the same v with known worker productivities.

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

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

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

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

II.E. Endogenous transparency and the need for legislation

The model thus far has assumed transparency is exogenously set at a common level for all firm types. In reality, a firm may have the ability to select its own level of transparency. In this section, we discuss the equilibrium outcome of a game in which the firm selects the level of transparency after observing its value v .

Formally, the *endogenous transparency* game proceeds as follows. Prior to workers arriving at $t = 0$, the firm observes v and simultaneously selects its maximum wage $\bar{w} \in [0, 1]$ and its level of transparency $\lambda \in [\underline{\lambda}, \infty)$. $\underline{\lambda} \in [0, \infty)$ is the minimum allowable level of transparency by law, and is common knowledge across the firm and all workers. We do not allow the firm to pick $\lambda = \infty$. We later relax this assumption and allow the firm to have access to a technology to set $\lambda = \infty$ with some probability $p > 0$.

Workers do not observe the selected level of λ , and the game continues as in our base model. The unobservability of the choice of λ to workers reflects that the firm is unable to contract with workers on the level of transparency. For example, a firm cannot promise that it will not discourage worker-to-worker communication about wages in the future, or that it will regularly inform workers of salary information that leads to costly renegotiations.

PROPOSITION 9. *Let $k < 1$. $\lambda = \underline{\lambda}$ in any equilibrium of the endogenous transparency game, regardless of the value of the firm.*

Despite the fact that an exogenously set, high level of transparency maximizes the profit of some firm types (see Example 1 in the Appendix), no firm type will pick a level of transparency that is higher than the minimum level allowed by law, $\underline{\lambda}$. The main intuition is that the unobservability of the selected level of transparency removes commitment power a firm obtains from higher transparency in our base model. In any candidate equilibrium in which the firm selects (λ, \bar{w}) where $\lambda \in (\underline{\lambda}, \infty)$, the firm has a profitable deviation to $(\underline{\lambda}, \bar{w})$ —in both cases, a zero measure of workers observes \bar{w} at $t = 0$, and therefore, all others will make the same initial offers. The firm will therefore employ the same set of workers at the same initial wages. However, by selecting $\underline{\lambda}$, the firm defers costly renegotiations to farther into the future, on average.

This result therefore implies that a law increasing the minimum-allowable transparency from $\underline{\lambda}$ to $\underline{\lambda}'$ will have the same effect as increasing the exogenous level of transparency from $\underline{\lambda}$ to $\underline{\lambda}'$.

Our results change in a continuous manner if the firm has, with probability p and independently of its value v , the technology to set $\lambda = \infty$. One real-world example is Buffer, a company that has built a salary-formula app into its website, that allows employees to “test” their own salaries and discover those of others as a function of observables.²⁶ Workers do not directly observe whether a firm has access to this technology. Our result changes, in that each firm will select $\lambda \in \{\underline{\lambda}, \infty\}$ in equilibrium, with a subset of firm types with access to the technology selecting $\lambda = \infty$. Only a firm that selects $\lambda = \infty$ can reveal to workers both its choice of transparency and \bar{w} , as all workers observe \bar{w} at $t = 0$. Therefore, firm

²⁶This particular technology also commits the firm to a rigid wage structure, and may therefore come with other costs, which we do not model.

types that maximize profit under a sufficiently high level of exogenous transparency in our base model will select $\lambda = \infty$ in equilibrium. An important message, however, is that the selection of $\lambda = \infty$ is a knife-edge result; a firm will not select an arbitrarily high, interior level of transparency regardless of its value.

For sufficiently small p , the firm will, with arbitrarily high ex-ante probability, select $\lambda = \underline{\lambda}$ in equilibrium. Therefore, the expected discounted surplus of each worker is approximately unchanged between our base game in which transparency is exogenously fixed at $\underline{\lambda}$ and one in which the firm selects the level of transparency, but there is legislation preventing the firm from selecting any transparency level strictly less than $\underline{\lambda}$.

We have not formally modeled the choice of workers to “bury their heads in the sand” and ignore wage information. Nevertheless, a richer model that allows each worker to ignore information would lead each worker to seek out wage information to the fullest extent allowed by the firm: for fixed \bar{w} higher transparency helps workers at the point of (re)negotiation. Because each worker has zero measure, no single worker will affect the equilibrium payoff, and therefore actions, of the firm.

III. EMPIRICAL EVIDENCE FROM “RIGHT OF WORKERS TO TALK” U.S. STATE LAWS

III.A. “Right of Workers to Talk” Legislation History

We study the enactment of legislation that commits employers to a minimum level of transparency through strong protections for coworkers who discuss pay with each other, which we refer to as “Right of Workers to Talk” (ROWTT) laws. As early as 1935, a clause in the National Labor Relations Act (NLRA) established worker rights to discuss pay in the spirit of facilitating collective bargaining; however, these protections were described in very general terms, “protecting concerted activity,” and violators did not face punitive damages, which led to the critique of the NLRA as a “toothless tiger” (Green, 2014).²⁷ More recently, individual U.S. states have enacted ROWTT laws, purportedly designed to combat discriminatory pay. These laws prohibit employers from retaliating against workers who discuss or inquire about coworker pay, with a clear enforcement mechanism: financial penalties to employers who violate the provisions. All state ROWTTs newly allow workers to sue employers in violation of ROWTT,²⁸ and all ROWTT laws provide prescribe punitive damages and/or fines for employers found in violation. In New Hampshire, ROWTT makes a manager herself criminally liable for violating the terms of the policy. These laws received

²⁷The National Labor Relations Act specifically excludes anyone in a supervisory role, government workers, and workers in certain travel occupations from its purview.

²⁸D.C.’s ROWTT is the only one that requires complaints be adjudicated by the labor board.

bipartisan support; both Republican and Democratic governors signed ROWTT bills into law.²⁹ In 2016, 25% of U.S. workers—all those working for federal contractors—received the same protections.³⁰

Anecdotal and survey evidence suggests workers have made use of these protections. The Philadelphia Inquirer writes that the “Google spreadsheet was the most powerful labor tool in 2019,” in reference to the use of shared spreadsheets as a way of spreading salary information. One employee who created such a spreadsheet told the reporter, “We were nervous enough that we checked online to figure out whether we were doing anything illegal by talking about what we were paid. We weren’t.” One journalist tracked down salary spreadsheets that have been created among museum professionals, baristas, journalists, ad agency staffers, and public interest lawyers (Reyes, 2019). Hegewisch et al. (2011) and Sun et al. (2021) survey workers in 2010 and 2017, respectively, about the prevalence of employer policies banning employers from discussing wages with one another. Between the dates of these two survey waves, 11 out of 13 ROWTT policies in our analysis were enacted. On our behalf the authors linked the outcome data of both surveys to state enactment of ROWTTs, revealing that the prevalence of non-disclosure policies declined twice as fast in states that enact an ROWTT during this time window than in those that do not. Among states that enact ROWTT legislation between 2010 and 2017, the share of workers subject to non-disclosure mandates falls 42%, from 16.2% of workers in 2010 to 9.4% in 2017.³¹

In Figure I we provide a timeline and geographical depiction of the enactment of each state law. There are 13 such policies in our study window (2000-2016), spanning the West, Northeast, Mid-Atlantic, and Midwest regions, but excluding the South. While ROWTT laws vary in their precise language, all legislation included in our analysis protects the right of workers to disclose their own salary and inquire about the salaries of others, and applies to all workers, with a few exceptions such as HR representatives.

III.B. Outcome data and sources

We use data from the American Community Survey (ACS) to track wages and employment between 2000 and 2016. Starting in 2000, the ACS surveyed more than 3 million individuals annually, allowing us to identify 4,077,593 individuals working in states that en-

²⁹12 Democrats, 5 Republicans, and 1 Independent governor signed these policies into legislation.

³⁰See https://web.archive.org/web/20170508135445/https://www.dol.gov/ofccp/regs/compliance/factsheets/Workplace_Rights_JRF_QA_508c.pdf.

³¹We would not expect rates to drop to zero even after the passage of ROWTT for several reasons; first, employees responding to the survey may not be aware of recent changes their employer made; secondly, employers may not have responded to the law either due to ignorance or as part of a long tradition of using gag laws to intimidate employees. Rates of non-disclosure policies decline in states that have not enacted an ROWTT by 2017 fall 22.5% from 20% to 15.5% over this time span, which also includes the 2016 federal mandate extending ROWTT to all federal contractors.

act ROWTT laws during our window of analysis. The ACS contains information on hours worked per week, weeks worked per year, sector, occupation, industry, U.S. state of work, and demographic characteristics, in addition to annual earnings information with a cap of \$250,000.³² We complement the main ACS sample with a measure of union coverage from the Current Population Survey (CPS). We merge in unionization at the occupation level using the standardized 1990 occupation codes provided by Flood et al. (2020) and Ruggles et al. (2021).

III.C. Empirical Strategy

We carry out a multi-period difference-in-difference design, often referred to as an event-study analysis. We follow Donohue III and Heckman (1991) and use neither event data nor outcome data after 2016, when complementary federal policies extend ROWTT protections to all federal contractors, and additional transparency requirements are enacted.³³

Our key identifying assumption is that the precise timing of ROWTT enactment during this twelve-year period, among states that eventually pass ROWTT laws, is uncorrelated with underlying wage and employment dynamics. We empirically test this assumption by examining how wages, employment and additional labor market features evolve in each state leading up to the enactment of ROWTT. Under the presumption that the states that eventually pass ROWTT laws are more similar along unobservables than states that do not pass these laws, our baseline specification excludes states that never enact ROWTT from our analysis.³⁴

In our baseline specification, we also assume that the effect of transparency on our outcomes is homogeneous across cohorts. If cohort effects were not homogeneous, the multi-period difference-in-differences estimation strategy may not result in an unbiased estimate of the average treatment effect across cohorts. As a robustness check we relax our assumption about homogeneous treatment effects and estimate the effect of each cohort separately.

Across all specifications we restrict our sample to prime working-aged individuals, ages 24 to 54, employed full-time in the private sector.³⁵ The reason we focus exclusively on the

³²Unfortunately these data do not include firm identifiers, the reason we cannot test within-firm predictions.

³³In addition to extending ROWTT protections to contractors (25% of the workforce), the federal government began requiring all firms in the country with more than 100 employees to begin reporting average wages broken down by race and gender in December, 2016.

³⁴We recognize trade-offs in making this choice. Excluding untreated states reduces power and introduces potential for colinearity between year fixed effects and linear treatment effects. While we expect our effects are non-linear (discontinuous around the year of the event) and hence, our estimation less susceptible to this colinearity, we verify this by replicating the analysis using our final cohort as a pure set of control states as well as adding back in states that never pass ROWTT.

³⁵We consider a worker to be employed full-time if they self-report that they usually work at least 35 hours/week and work for at least 48 weeks in the last year.

private sector is that many local laws have made salaries public information for government workers; for example, in California two-thirds of cities independently chose to disclose the compensation of city employees prior to a 2010 mandate to disclose salaries of all municipal employees (Mas, 2017).³⁶ For this reason the ROWTT laws may or may not increase transparency for these workers. We discuss the effect of ROWTT in the public sector further in Appendix Section B.

We estimate the dynamic effect of ROWTT laws over the 3 years following their enactment. We also estimate the dynamic effect of ROWTT laws over the 6 years prior to their enactment as a test of whether enactment was precipitated by any underlying events that could co-move with our outcomes of interest, such as a rise in pro-business sentiments and related policies.

We estimate the following multi-period difference-in-differences specification:

$$y_{ist} = \alpha_s + \sum_{w=-6}^{-2} \beta_w \mathbf{1}\{t - E_s = w\} + \sum_{w=0}^3 \beta_w \mathbf{1}\{t - E_s = w\} + \gamma \mathbf{1}\{t - E_s < -6\} + \delta \mathbf{1}\{t - E_s > 3\} + \lambda \mathbf{X}_{ist} + \epsilon_{ist} \quad (7)$$

Where i indexes individuals, s indexes states and t indexes year. In our main specification, y_{ist} is the logarithm of annual wage income or the share of all people employed full-time in the private sector. α_s is a state fixed effect and E_s is the year when state s enacts the ROWTT policy. Thus, $t - E_s$ indexes years relative to the event. $w = -1$ is the omitted reference period, γ and δ are indicators for periods outside the event window. \mathbf{X}_{ist} is a vector of controls that include age (quadratic), education, year-by-industry (NAICS 3-digit) and year-by-occupation (SOC 3-digit) indicators. We allow for interactions between available demographic characteristics, namely marital status, race and gender, and we allow region-by-industry effects to differ by gender. Standard errors are two-way clustered by state and year.³⁷

In our baseline specification, we report the results from a balanced composition of states following the enactment of the law, we estimate the dynamic post period effects for states with events up through 2013 separately, and report the coefficients from the following interaction term added to Equation 7:

$$\sum_{w=0}^3 \beta_w^* \mathbf{1}\{t - E_s = w\} \times \mathbf{1}\{E_s \leq 2013\} \quad (8)$$

³⁶To identify all these local policies, Mas (2017) resorts to newspaper articles on the wayback machine. To carry this out nation-wide posed an insurmountable hurdle for us.

³⁷When a worker lives and works in different states, we use the state in which they are primarily employed.

In a series of robustness tests we include year-by-Census-division fixed effects α_{tr} , and we weight our sample to estimate a counterfactual where the composition of workers in each education-gender-state cell remains fixed throughout the post period.³⁸ To relax our assumption of homogeneous treatment effects across cohorts, we allow treatment effects to vary depending on the year ROWTT is enacted, and calculate the weighted average of each year-specific treatment effect. We implement this using the Sun and Abraham (2020) interaction-weighted estimator, designed to recover average treatment effects even in the presence of underlying heterogeneity across years.³⁹ We include the full specification and estimation details in Appendix Section A.

We estimate the heterogeneous effects of ROWTT separately for occupations in which workers have a relatively high versus low degree of individual bargaining power by splitting occupations according to whether they are above or below the median unionization rate of 7%. We add the following interaction terms to Equation 7.

$$\sum_{w=-6}^{-2} \beta_w \mathbf{1}\{t - E_s = w\} \times \mathbf{1}\{\text{above median}\} + \sum_{w=0}^3 \beta_w \mathbf{1}\{t - E_s = w\} \times \mathbf{1}\{\text{above median}\} + \gamma \mathbf{1}\{t - E_s < -6\} \times \mathbf{1}\{\text{above median}\} + \delta \mathbf{1}\{t - E_s > 3\} \times \mathbf{1}\{\text{above median}\} \quad (9)$$

III.D. Results

Figure II presents the results based on the econometric framework described in Section III.C. The event-study graph shows the evolution of log wages (Panel A) and private sector employment (Panel B) in each of the 6 years leading up to the enactment of ROWTT and 3 years after enactment. The year before the event (-1) corresponds to the omitted category, and thus the corresponding coefficient is always zero by construction.

When inspecting Figure II, Panel A note that the coefficients refer to wage differences relative to the period prior to ROWTT enactment, after controlling for calendar year fixed effects. As a result, a coefficient of zero does not imply that wages remain stagnant in nominal terms; rather, it indicates similar growth rates of wages in states leading up to the enactment of ROWTT. The range along the y-axis has been set to approximately +/- 1 standard deviation in average wages over time, within a state. In the six years leading up to the enactment of ROWTT, coefficients are precisely estimated and statistically indistinguishable

³⁸We take the year before the policy is enacted as the reference year and estimate the educational distribution of each state separately for men and women. Within each state, we then reweight the sample in every other year to match the education-by-gender distribution in the reference year.

³⁹This estimator is part of a large recent applied econometric literature investigating robust dynamic effect estimators and weaknesses of the standard OLS approach (Athey and Imbens, 2018; Borusyak et al., 2021; Callaway and Sant’Anna, 2020; de Chaisemartin and D’Haultfœuille, 2020; Goodman-Bacon, 2018).

from zero, suggesting that our assumption of parallel trends in wages holds.

By contrast to the period before enactment, after ROWTT is enacted, the evolution of wages diverges from the wage path of states which enact ROWTT in different years. One year after enactment, wages are 2.2% lower (p-value < 0.001), and by three years after enactment wages are 2.6% lower (p-value = 0.017) than the period prior to enactment. A 2.6% decline in average private sector wages, on this time scale, represents an economically sizeable and rapid shift, especially in light of evidence that downward wage rigidities are pervasive in the U.S. economy (Grigsby et al., 2021). To put this result in perspective, if firms had fixed nominal wages during the window we examine, year over year employees would have experienced a comparable decline in real wages on average, approximately 2.8%.⁴⁰ Most likely, the wage path diverged following ROWTT enactment through the slowing pace of wage increases and the downward revision of starting wages.

In Appendix Section B, we conduct the event study using unconditional quantile regressions (Firpo et al., 2009). Wages in the 50th percentile closely track average wages, falling to 1.8% one year after enactment (p-value=0.008), and 1.9% in the third year (p-value=0.083). Average post-treatment wage effects are negative and statistically different from zero across the income distribution, ranging from 3.8% to 1.4% between the 10th and 90th percentiles, but wage effects are largest for the lowest earners, consistent with our hypothesis that transparency’s effect will be largest when roles are standardized, workers are substitutable, and renegotiations occur frequently, such as in high-churn jobs.⁴¹

Wage declines could stem from a change in hiring practices resulting from ROWTT. Our theory predicts that a shift in bargaining power from workers to firms could both lower wages paid and, under some circumstances, reduce the employment of high outside option workers. Panel B of Figure II reports the estimated coefficients from the same event study specification as Panel A, replacing our dependent variable with the share of workers employed full-time in the private sector. The range of our y-axis is set to be approximately +/- 1 standard deviation of the average share employed in the private sector overtime within states. Our point estimates suggest that employment remains constant leading up to the ROWTT enactment, and continues on the same path after ROWTT. After one year, the coefficient is 0.25% (p-value= 0.165) and after three years the point estimate is 0.39% (p-

⁴⁰According to <https://www.ssa.gov/oact/cola/central.html>, nation-wide average wages rise by 2.8% each year from 2004-2016.

⁴¹Our bargaining theory is premised on the notion that workers can be compared to one another. When workers are specialized, the bargaining effects no longer hold. Our model also predicts that workers can renegotiate immediately upon learning the wage information of their peers. If workers must wait for a scheduled performance review to renegotiate, then λ represents the rate at which workers both learn the wages of peers *and* have the ability to renegotiate. Therefore, markets with high churn or more frequent wage negotiations, may have a larger increase in effective transparency following the enactment of ROWTT. Using nation-wide payroll data, Cullen and Perez-Truglia (2021) find that over 80% of new hires in a given year are hourly workers in lower paying positions.

value=0.165). We cannot reject zero impact on employment during the 3 years after ROWTT enactment.

While employment does not appear to be affected by ROWTT, in theory composition changes could be masked. Employees could, for example, be systematically leaving high value-of-labor employers for low value-of-labor employers. Alternatively highly skilled workers may be leaving the private sector at a faster rate and low-skilled workers hired at a faster rate despite constant aggregate employment levels. We carry out two additional exercises to test if the composition of employed workers responds to ROWTT and contributes to the wage declines. First, we estimate a counterfactual model of the wage effects whereby the share of employees in each detailed education-by-gender cell remains of equal size in every year subsequent to $t = -1$. Figure A.1 Panel C shows the results remain essentially unchanged. After the first year, wages decline by 2.0% (p-value = 0.023) and fall to 2.5% (p-value = 0.125) after three years. Second, we carry out a decomposition exercise in Appendix Section D using liberal estimates of composition changes. We conclude composition changes from employment could account for up to half of our estimated overall effect on wages.

In Table I we report the results of several alternative specifications to our baseline. In Col. 1, we present our baseline results, the multiperiod difference-in-differences estimator with a fixed composition of states in the post period. We report the dynamic effects around ROWTT in +/- 3 years around the event, and we also report the average of the post-treatment period: 1.8% (p-value < 0.001). In Col. 2, we allow all cohorts (2004-2016) to contribute to all periods for which the data are available. In the post period each coefficient is identified by a different set of states (eg. outcome data is not available in period +3 for the cohort with events in 2014). The post event coefficients average 1.5% (p-value < 0.001) in the post-treatment period, and fall by 2.6% three years after enactment, exhibiting nearly identical results as our baseline. In Col. 3 we include detailed region-by-year fixed effects using detailed Census divisions, effectively restricting comparisons between states to neighboring states.⁴² The average post-treatment effect is 1.7% (p-value = 0.074). This is statistically indistinguishable from our baseline specification. In Col. 4, we report the point estimates and standard errors for the education-by-gender reweighted specification. The average post-treatment effect is 1.5% (p-value = 0.053). Finally, we compute the Sun-Abraham interaction-weighted estimator which relaxes our assumption of homogeneous treatment effects across cohorts, as described in Appendix Section A. The average post-treatment point estimate is 2.2% (p-value < 0.001), which is again similar to our other estimates. These results of all these alternative specifications are discussed in further detail in Appendix Section and are presented graphically in Figure A.1. In Figure A.2, we carry out the same robustness

⁴²We pool together the “West North Central” and “East North Central” divisions to form the “Midwest” Census region to ensure that there are no divisions containing only a single treated state.

checks for our employment results. Across all specifications we consistently find employment effects close to zero and precisely estimated.⁴³

In Figure III, we report the wage effect for above and below median rates of unionization at the occupation level.⁴⁴ Our theory predicts that as individual bargaining power declines, the effect of transparency shrinks (Section II.B). In unionized firms, the firm may first negotiate with a union representative, so that individual workers face a more rigid wage schedule. Within the agreed-to wage schedule, workers may have more limited bargaining power to negotiate pay. We use the rates of unionization within an occupation to approximate individual bargaining power. In Figure III Panel A we separately plot the dynamic effects of ROWTT for occupations with above and below the median share of unionized workers, estimated jointly following Equations 7 and 9. In Panel B we plot the difference between the effects for occupations with low and high rates of unionization. Leading up to the enactment of ROWTT, wages in high and low unionized occupations follow the same trajectory. Starting the year that ROWTT policies are enacted, wages between these groups of occupations diverge. Among relatively unionized occupations, wages fall by 1.4% (p-value = 0.130) one year after enactment and remain at 1.6% (p-value = 0.052) three years after enactment. For occupations with relatively low rates of unionization, wages decline nearly twice as much, an additional 1.6% (p-value = 0.032) over the post-period window, and experience wage declines of 3.2% three years after enactment. In Appendix Section B we show the average post treatment effect grows systematically as unionization rates fall from 20% in the upper quartile down to 2% in the bottom quartile. These results offer suggestive evidence that indeed, collective bargaining agreements that reduce individual bargaining power also mitigate the effects of pay transparency on the bargaining position of workers.

When interpreting our heterogeneous treatment effects, it is important to consider alternative interpretations to the causal relationship we present. While our theory predicts a causal relationship, our empirical test does not rule out the possibility that occupations with higher rates of unionization are different along dimensions that could mute the effects of ROWTT but are orthogonal to individual bargaining power. For example, occupations with higher rates of unionization could be associated with institutions that do not adapt quickly to changes in the labor market and operate in less mobile labor markets; hence, are slow to adjust their wage trajectories and do not do so within our window. Alternatively, the policy itself could have been enforced differently within occupations that have higher rates of unionization. In Section IV, we strengthen this empirical test by exploring the relationship between unionization and transparency’s effect across a wide range of labor markets.

⁴³In unreported results, we replicate Table I including all U.S. states. Results are qualitatively the same, and we do not gain sufficient power to supplant our baseline specifications.

⁴⁴The median rate is 7% and there is a five-fold difference in unionization rates between occupations above and below the median.

III.E. Main Threats to Internal Validity

Inherent in our empirical strategy are several assumptions. The first is that ROWTT laws are enacted in isolation; in other words, these policies are not simultaneously coupled with additional legislation or timed around another noteworthy event. We have reason to believe this is the case. While nearly all ROWTT legislation are amendments to existing equal pay laws,⁴⁵ in only four cases is there arguably related legislation enacted around the same time.⁴⁶ Our results are robust to excluding these four events.

Relatedly, the decision to enact ROWTT cannot be itself driven by changes that are underway, in essence a story of reverse causality whereby declining wages leads to the enactment of ROWTT, rather than the other way around. Reverse causality is typically less of a concern when effects are discontinuous and occur after the policy is enacted; nevertheless, we also collect facts about the motivation for the passage of ROWTT laws. More than three-quarters of the ROWTT policies refer to pay discrimination in the title or preamble describing the the law (the partial equilibrium narrative), and nowhere is there mention of wage levels.⁴⁷ We also consider the possibility that a third factor leads to both declining wages and the enactment of ROWTT. We could, in principle, be detecting a rapid shift in sentiment in favor of businesses that, either through policies or atmosphere, effectively shifts bargaining power towards firms and simultaneously leads to the enactment of ROWTT. Or alternatively a shift in sentiment in favor of labor and related policies (including ROWTT) intended to combat inequality and yet, inadvertently, this movement also lowers average wages. To test whether there exist concurrent shifts in the business environment or pro-labor activities around ROWTT enactment, we study whether state corporate tax changes co-vary with the enactment of ROWTT policies. State corporate tax rates change with a high degree of frequency, providing an index of the state’s business environment year over year. We take advantage of a dataset constructed by [Slattery and Zidar \(2020\)](#), and replicate our event study analysis with state corporate tax rates as the dependent variable. In Appendix Figure A.8, we show that corporate tax rates do not change around ROWTT enactment. Coefficients are small, and the confidence interval always includes zero.

⁴⁵11 of 13 were passed as amendments to existing equal pay laws; D.C. and CT were not tied to previous protections.

⁴⁶VT enacts a new law about working mothers in the workplace, and new guidelines supporting flexible working arrangements. MN sets aside money for grants to create programs to hire women in different workplaces. NH creates additional anti-retaliation laws. DE creates new provisions and protections regarding reproductive health. Salary history bans, salary range posting mandates and wage gap disclosure policies are not coupled with ROWTT laws nor are they enacted within several years of any ROWTT law that we study in the window 2004 to 2016.

⁴⁷Commonly used phrases are “pay equity,” “equal pay,” “pay discrimination,” and “pay differential.”

IV. EXAMINATION OF OTHER PAY TRANSPARENCY MANDATES THROUGH OUR BARGAINING FRAMEWORK

Our model predicts similar equilibrium labor market outcomes of pay transparency policies that require employers to post individual salaries or report pay gaps between men and women to all workers (Section II.D). In this section, we gather empirical evidence for these claims.

Several of the most sweeping pay transparency laws enacted, including in Austria, Denmark, and the U.K., require employers to reveal wage gaps between men and women. Interventions in Canada and the U.S. require the online posting of wages for certain labor markets. We compile results from studies that evaluate these policies to test the equilibrium effects predicted by our bargaining framework in diverse settings. We refer to this empirical exercise as a meta-analysis, because we examine average wage declines across settings and combine the data from these studies in a mixed-effects meta-regression to test our comparative static prediction that wages decline less when unionization rates are higher.

IV.A. *Criteria for selecting pay transparency studies*

We seek to include the universe of pay transparency studies, subject to several criteria: first, the study evaluates the effect of a pay transparency policy in a real-world labor market. Second, the study evaluates the effect of pay transparency on the wages of all employees in that labor market. Third, the study refers to the policy evaluated as “pay transparency” or a related term.⁴⁸

In practice, our restrictions lead some high-quality studies (e.g. Breza et al. (2018); Gächter and Thöni (2010)) to be excluded as they investigate transparency in a lab context or lab in the field where researchers role play as employers. Some high-quality studies (e.g. Burn and Kettler (2019); Gamage et al. (2020); Kim (2015); Roussille (2021)) are excluded as they investigate the effect of transparency on subsets of workers.

While we can take steps to identify the universe of studies that meet these criteria, one concern with meta-analyses is that publication bias results in studies skewed toward finding a significant effect (Andrews and Kasy, 2019). In our case, this is likely a relatively minor concern. Overall wage levels are only a secondary outcome in the all of these studies; one study (Mas, 2017) primarily focuses on wage compression between high- and low-paid workers, and the remainder focus first and foremost on the gender wage gap, consistent with

⁴⁸The full list of search terms include “pay transparency,” “wage transparency,” “salary transparency,” “pay disclosure,” “wage disclosure,” and “salary disclosure.” We searched for papers on the Econ lit database, SSRN, arxiv, NBER working papers series, IZA working paper series, Google Scholar, and the works cited of other included studies. We performed this search several times, with the final search being conducted in May, 2021.

the stated goal of pay transparency policies to close the wage gap between men and women and other minorities.

IV.B. Overview of studies & results extraction

Our search results in eight independently-conducted papers, one of which provides two studies internally.⁴⁹ In total, these papers evaluate six distinct pay transparency mandates spanning five countries. In four of these studies, governments mandate full disclosure of employee salaries, and in the remaining four, wage gaps between men and women.⁵⁰ Appendix Table A.1 provides details about each of the studies.

We extract information about overall wage effects and labor market unionization from each study. We select the author’s preferred specification when clear, as is the case for six of the eight studies. When not specified, we select the specification closest to our theoretical framework, i.e. examining wage spillovers within position.⁵¹ For three of eight studies, the authors do not report a single post-treatment effect. To minimize assumptions about the covariance between estimates, we do not aggregate over annual estimates when authors do not report a single post-treatment effect; rather, we choose the final period in the window reported.⁵²

All but three studies specifically report the effect of transparency policies on men’s wages, and then provide the differential effect of the policy on women’s wages.⁵³ We impute the overall wage effect of transparency by weighing the changes in men’s and women’s wages by the share of men in the industry, and the pre-transparency ratio of female to male wages.⁵⁴

⁴⁹Baker et al. (2019); Bennedsen et al. (2020); Blundell (2021); Böheim and Gust (2021); Duchini et al. (2020); Gulyas et al. (2020); Mas (2017); Obloj and Zenger (2020) (two studies are included in Baker et al. (2019)).

⁵⁰Disclosure of employee salaries is facilitated by newspapers and other organizations who release salary information garnered through Freedom of Information Act requests in Mas (2017) and Obloj and Zenger (2020).

⁵¹Baker et al. (2019) and Obloj and Zenger (2020) present two preferred specifications each. In Baker et al. (2019), one specification considers a worker as treated if the wage of a coworker at the same department and institution is revealed. Another specification consider a worker as treated if the wage of a coworker at the same department, institution, and *rank* is revealed. We select the latter specification because our model’s predictions are in settings where wages of peers with the same value to the employer are revealed. The authors note on page 14 that this specification is the one that better captures “horizontal” rather than “vertical” comparisons. We apply the same reasoning to our choice of specification in Obloj and Zenger (2020).

⁵²Böheim and Gust (2021) estimate the effect of the transparency policy on different cohorts of firms which vary in the timing of treatment. We similarly select the final cohort in their analysis. All cohorts have wage effects that are statistically indistinguishable from zero. Weighing the average change in each cohort by number of workers leads to similar inferences.

⁵³Mas (2017) provides wage effects for managers and non-managers, Gulyas et al. (2020) use women as the base category and provide the differential wage effect for men, and Böheim and Gust (2021) present the wage effects separately for men and women.

⁵⁴Bennedsen et al. (2020) also report an increase in firm profits per worker. Given this outcome is only

IV.C. Results

In Figure IV, we graphically present the relationship between the share of the workforce covered by a collective bargaining agreement in each study (x-axis), and the estimated effect of the pay transparency mandate on log wages (y-axis). For each study we include two points. The first is an effect size directly reported in the paper and refers to the effect of pay transparency on men’s wages.⁵⁵ For these point estimates we also plot the reported 95% confidence interval. We include a second point, lighter in color, to indicate our imputed estimate of transparency’s effect on the overall population.

The results of these studies match our theoretical predictions. Observations generally fall below the x-axis, indicating a negative impact of pay transparency on wages, and follow an upward-sloping line, indicating that the effect on wages is smaller in magnitude as a higher share of the workforce has wages set by a collective bargaining agreement. The resulting slope on the effect of transparency on men’s wages is 0.018 (p-value=0.008), implying that a 10 percentage point reduction in the share of workforce under a collective bargaining agreement results in a 0.18 percentage point larger decrease in men’s wages following a transparency intervention. Studies with nearly full coverage by a collective bargaining agreement see no statistically significant change in wages following the transparency intervention.

Comparing the point estimate for men only, and our imputed point estimate for the entire population (the lighter point), reveals that the effect for the whole population follows the same patterns as it does for men only. However, the effect on women’s wages is smaller than the effect on men’s wages. As these studies all investigate the effect of wages *within* firm, our Theorem 1 predicts that transparency results in a smaller wage gap when the distribution of men’s outside options first order stochastically dominates that of women.⁵⁶ In Appendix Section E, we offer more details about the reported gender wage gaps featured in these studies.

Two of the studies included in our meta-analysis offer an explanation for the average wages they observe after the transparency mandate.⁵⁷ They offer two distinct explanations,

tracked in one study, we are not able to formally include it in our meta analysis, although we note that it is consistent with our theoretical predictions. They find a point estimate of 0.014, with standard error 0.042.

⁵⁵As we note earlier, we are unable to include the effect of pay transparency on men’s wages for Mas (2017) and Böheim and Gust (2021).

⁵⁶Formally, suppose there are two types of workers, m (male) and f (female), and each worker i belongs to exactly one type. Let G_ℓ represent the distribution of outside options for type $\ell \in \{m, f\}$, and let $G(x) := qG_m(x) + (1 - q)G_f(x)$ for all $x \in [0, 1]$, where $q \in [0, 1]$ is the proportion of men in the market. Denote the average equilibrium expected earnings of an employed worker of type $\ell \in \{m, f\}$ as $T(\Omega, v, \theta_i, \ell)$. $G_m(\cdot)$ first-order stochastically dominates $G_f(\cdot)$ then $\frac{\mathbb{E}_{G_f}[T(\Omega, v, \theta_i, f)]}{\mathbb{E}_{G_m}[T(\Omega, v, \theta_i, m)]}$ converges monotonically to 1 as Ω converges to 1 for all v . The reason is that 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, and Theorem 1 therefore implies wage equalization across genders.

⁵⁷With one exception, these papers do not offer a formal theory of wage adjustments, but propose high

each well-suited to their context. [Mas \(2017\)](#) focuses on public aversion to high pay, and the pressure that the public eye could place on employers to reign in high salaries. This channel may be especially important when salaries are available to the public, and when public tax dollars support wages, as in the case of public sector workers ([Baker et al., 2019](#); [Mas, 2017](#); [Obloj and Zenger, 2020](#)). We expect this channel to play a smaller role in cases where employers or employees can only share salaries internally ([Bennedsen et al., 2020](#); [Böheim and Gust, 2021](#); [Gulyas et al., 2020](#)).⁵⁸ Moreover, we might expect the public eye to similarly constrain the wages of unionized and non-unionized workers. [Bennedsen et al. \(2020\)](#) focus on the potential for worker morale to respond to visible pay inequality between peers, followed by declines in productivity, and in turn, average wage declines. Indeed, the partial equilibrium effect of lower morale and lower productivity in response to peer pay inequality have been shown in several settings where the employer did not have agency to adjust wages ([Breza et al., 2018](#); [Cullen and Pakzad-Hurson, 2021](#)). We view this channel as entirely consistent with our bargaining framework to the extent that the morale and productivity consequences are a threat that leads employers to equalize wages, as discussed in [Eliaz and Spiegler \(2013\)](#). In settings where employers can adjust wages, productivity is measured in a few cases; these cases broadly suggest that employers indeed make adjustments to off-set productivity declines on the order of magnitude observed in partial equilibrium settings. [Bennedsen et al. \(2020\)](#) find log sales per employee falls minimally and insignificantly the first year after the mandate, though wages fall by 1.4%. Sales dip in the second year leaving room for the possibility that average wage declines preceded (or even caused) the productivity decline rather than the other way around.⁵⁹ [Duchini et al. \(2020\)](#) measure firm output per worker and finds a statistically insignificant decline; stock returns dip three days after the policy but recover by the fourth day.⁶⁰ Wage declines are observed even after controlling for productivity in [Duchini et al. \(2020\)](#) and [Obloj and Zenger \(2020\)](#). In our own experiments, we find the net effect of morale pressures is met with the strategic response of employers to equalize wages, and hence, our equilibrium

quality verbal explanations. In Appendix F of his paper, [Blundell \(2021\)](#) models how worker and/or consumer preferences for equal pay affect firm's wage setting practices. In his model, men and women perform different jobs in the firm, and therefore, firms set a completely common wage for each job, with or without transparency.

⁵⁸[Cullen and Perez-Truglia \(2019\)](#) offer evidence that, internally, employees do not exhibit an aversion to vertical inequality.

⁵⁹Year 1 sales per employee fall by -0.008 (s.e.=0.012) one year after a transparency mandate, while log wages fall by -0.014 (s.e.=0.006). Log sales dips further in the following year (-0.040 log points (s.e.=0.015)) as do wages (-.033 (s.e.=0.007)),

⁶⁰[Duchini et al. \(2020\)](#) report the 3-day cumulative abnormal returns decrease by around 35 basis points following the publication of gender equality data but the effect fades away after four days. Firm output per worker falls by 0.029 (s.e.=0.035); however, the authors do not specify firm output units in the current draft to the best of our knowledge, so we struggle to interpret the magnitude of this effect.

predictions of lower average wages hold (Cullen and Pakzad-Hurson, 2021).

V. CONCLUSION

Although pay transparency has been in the political and popular spotlights due to its perceived benefits to workers, its effect on wages and hiring are not well understood. Our theoretical and empirical analyses of equilibrium wages and employment under greater transparency reveals consequences that run counter to popular wisdom.

Pay transparency reduces the bargaining power of employees in settings where workers start out with a degree of individual bargaining power. The intuition behind this result is that employers credibly refuse to pay high wages to workers with strong outside options to avoid costly renegotiations with others. When workers bargain collectively, individual outside options play a smaller role in wage negotiations and thus, transparency’s effect is muted. We find evidence consistent with this in the U.S. private sector following the rollout of “Right of Worker to Talk” U.S. state laws between 2004 and 2016. In the three years following enactment of the law, wages decline approximately 2% overall, but they declined half as much in occupations with above-median rates of unionization compared to occupations with below-median rates of unionization.

Our findings highlight the central role that bargaining plays in mediating wages. Without an equilibrium response through bargaining, we would expect wages to rise after transparency is introduced, as transparency’s direct effect of revealing pay disparities allows low-wage workers to negotiate higher pay. Our findings are consistent with prior research on the effects of pay transparency policies; a meta-analysis of studies evaluating a wide range of pay transparency mandates collectively corroborate the negative impact of transparency on average wages, and our bargaining model offers a unified framework to understand differences in the effect size across studies as a function of individual bargaining power in the local labor markets.

REFERENCES

- George A. Akerlof and Janet L. Yellen. The Fair Wage-Effort Hypothesis and Unemployment. *The Quarterly Journal of Economics*, 105(2), 1990.
- Isaiah Andrews and Maximilian Kasy. Identification of and correction for publication bias. *American Economic Review*, 109(8):2766–94, August 2019. doi: 10.1257/aer.20180310. URL <https://www.aeaweb.org/articles?id=10.1257/aer.20180310>.
- Susan Athey and Guido Imbens. Design-based analysis in difference-in-differences settings with staggered adoption. *arXiv:1808.05293 [cs, econ, math, stat]*, Sep 2018. URL <http://arxiv.org/abs/1808.05293>. arXiv: 1808.05293.
- Kyle Bagwell and Asher Wolinsky. Game theory and industrial organization. *Handbook of Game Theory with Economic Applications*, 3:1851–1895, 2002.
- Michael Baker, Yosh Halberstam, Kory Kroft, Alexandre Mas, and Derek Messacar. Pay Transparency and the Gender Gap. *NBER Working Paper No. 25834*, November 2019.

- Erling Barth, Alex Bryson, James C. Davis, and Richard Freeman. It's Where You Work: Increases in the Dispersion of Earnings across Establishments and Individuals in the United States . *Journal of Labor Economics*, 34(2):S67–S97, 2013.
- Morten Bønedal, Elena Simintzi, Margarita Tsoutsoura, and Daniel Wolfenzon. Do Firms Respond to Gender Pay Gap Transparency? *NBER Working Paper No. 25435*, April 2020.
- Dirk Bergemann and Johannes Hörner. Should First-Price Auctions Be Transparent? *American Economic Journal: Microeconomics*, 10(3):177–218, 2018.
- Francis Bloch and Harl Ryder. Two-Sided Search, Marriages, and Matchmakers. *International Economic Review*, 41(1):93–115, 2000.
- Jack Blundell. Wage responses to gender pay gap reporting requirements. *Centre for Economic Performance Discussion Paper 1750*, 2021.
- René Böheim and Sarah Gust. The Austrian Pay Transparency Law and the Gender Wage Gap. *IZA Discussion Paper 14206*, 2021.
- Gabriele Borg, Diego Gentile Passaro, and Santiago Hermo. Do Minimum Wages Increase Rents? Evidence of Spatial Spillovers Using Local High-Frequency Data. *mimeo*, 2021.
- Kirill Borusyak, Xavier Jaravel, and Jann Spiess. Revisiting event study designs: Robust and efficient estimation. *Working Paper*, page 48, May 2021.
- Anat Bracha, Uri Gneezy, and George Loewenstein. Relative pay and labor supply. *Journal of Labor Economics*, 33(2):297–315, 2015.
- Giulia Brancaccio, Dan Li, and Norman Schürhoff. Learning by Trading: The Case of the U.S. Market for Municipal Bonds. *Mimeo*, 2020.
- Hanna Brenzel, Hermann Gartner, and Claus Schnabel. Wage bargaining or wage posting? evidence from the employers' side. *Labour Economics*, 29:41–48, 2014.
- Emily Breza, Supreet Kaur, and Yogita Shamdasani. The Morale Effects of Pay Inequality. *Quarterly Journal of Economics*, 133(2):611–663, 2018.
- Kenneth Burdett and Melvyn G. Coles. Long-Term Partnership Formation: Marriage and Employment. *The Economic Journal*, 109(456):F307–F334, 1999.
- Ian Burn and Kyle Kettler. The more you know, the better you're paid? Evidence from pay secrecy bans for managers. *Labour Economics*, 59:92–109, 2019.
- David A. Butz. Durable-Good Monopoly and Best-Price Provisions. *American Economic Review*, 80(5):1062–1076, 1990.
- Sydnee Caldwell and Nikolaj Harmon. Outside Options, Bargaining, and Wages: Evidence from Coworker Networks. *Mimeo*, 2019.
- Brantly Callaway and Pedro H. C. Sant'Anna. Difference-in-differences with multiple time periods. *arXiv:1803.09015 [econ, math, stat]*, Dec 2020. URL <http://arxiv.org/abs/1803.09015>. arXiv:1803.09015.
- David Card, Alexandre Mas, Enrico Moretti, and Emmanuel Saez. Inequality at Work: The Effect of Peer Salaries on Job Satisfaction. *American Economic Review*, 102(6):2981–3003, 2012.
- Hector Chade. Matching with Noise and the Acceptance Curse. *Journal of Economic Theory*, 129(1):81–113, 2006.
- Kalyan Chatterjee and William Samuelson. Bargaining under Incomplete Information. *Operations Research*, 31(5):835–851, 1983.
- Alain Cohn, Ernst Fehr, Benedikt Herrmann, and Frédéric Schneider. Social Comparison and Effort Provision: Evidence from a Field Experiment. *Journal of the European Economic Association*, 12:877–898, 2014.
- Thomas E. Cooper and Timothy L. Fries. The most-favored-nation pricing policy and negotiated prices. *International Journal of Industrial Organization*, 9:209–223, 1991.

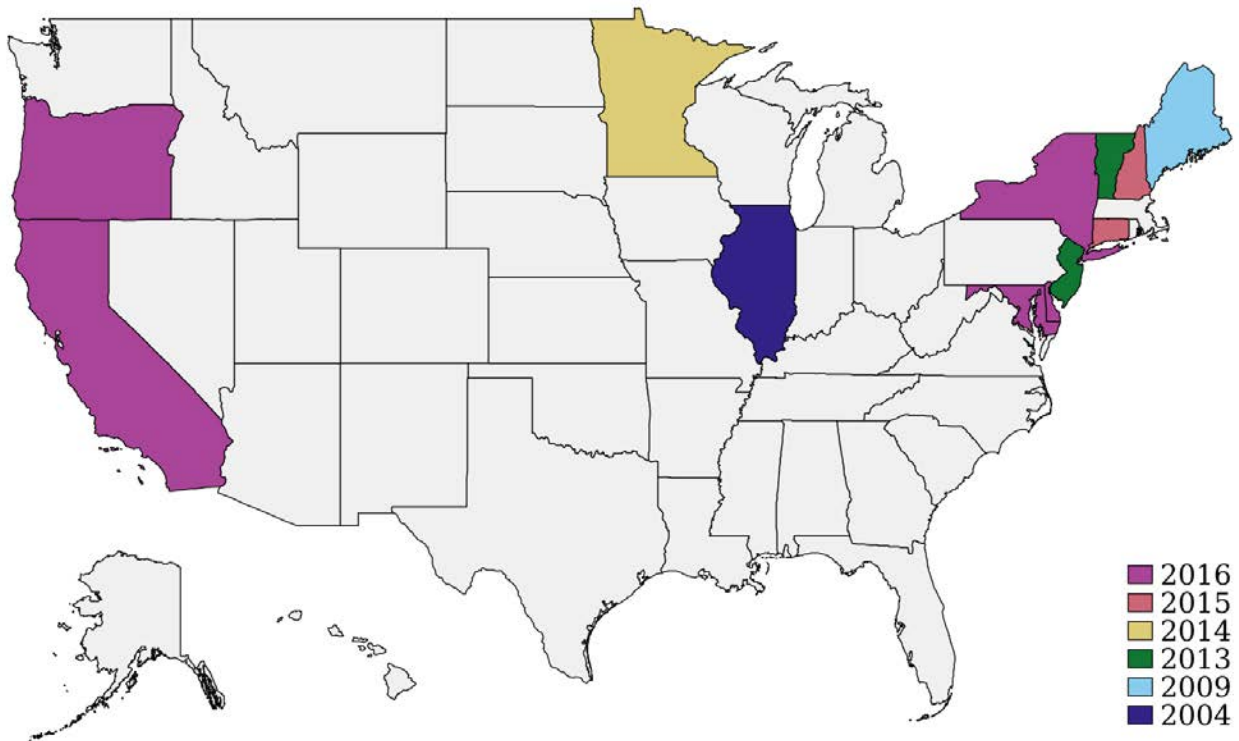
- Zoë Cullen and Bobak Pakzad-Hurson. Microfoundations of Pay Transparency: Morale and Bargaining. *mimeo*, 2021.
- Zoë Cullen and Ricardo Perez-Truglia. How Much Does Your Boss Make? The Effects of Salary Comparisons. *NBER Working Paper No. 24841*, 2019.
- Zoë Cullen and Ricardo Perez-Truglia. Hitting the Benchmark: Information Frictions, Pay, and Pay Equity. *Working Paper*, 2021.
- Clément de Chaisemartin and Xavier D’Haultfœuille. Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–2996, Sep 2020. ISSN 0002-8282. doi: 10/gg94xm.
- John J. Donohue III and James Heckman. Continuous versus episodic change: the impact of civil rights policy on the economic status of blacks. *Journal of Economic Literature*, 29:1603–1643, 1991.
- Arindrajit Dube, Laura Giuliano, and Jonathan Leonard. Fairness and Frictions: The Impact of Unequal Raises on Quit Behavior. *American Economic Review*, 109(2):620–663, 2019.
- Emma Duchini, Stefania Simion, and Arthur Turrell. Pay Transparency and Cracks in the Glass Ceiling. *Mimeo*, May 2020.
- Alex Edmans, Xavier Gabaix, Tomasz Sadzik, and Yuliy Sannikov. Dynamic CEO Compensation. *Journal of Finance*, 67(5):1603–1647, 2012.
- Kfir Eliaz and Ran Spiegler. Reference Dependence and Labor-Market Fluctuations. *NBER Macroeconomics Annual*, 28:159–200, 2013.
- Tore Ellingsen and Åsa Rosén. Fixed or Flexible? Wage-Setting in Search Equilibrium. *Economica*, 70(278): 233–250, 2003.
- European Commission. Pay transparency in the EU: A legal analysis of the situation in the EU Member States, Iceland, Liechtenstein and Norway. Technical report.
- 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.
- Sergio Firpo, Nicole M. Fortin, and Thomas Lemieux. Unconditional quantile regressions. *Econometrica*, 77(3):953–973, 2009. ISSN 1468-0262. doi: <https://doi.org/10.3982/ECTA6822>.
- Sarah Flood, Mariam King, Renae Rodgers, Steven Ruggles, and J.Robert Warren. Integrated public use microdata series, current population survey: Version 8.0 [dataset]. *Minneapolis, MN: IPUMS*, 11, 2020.
- Drew Fudenberg and David K. Levine. Maintaining a Reputation when Strategies are Imperfectly Observed. *The Review of Economic Studies*, 59(3):561–579, 1992.
- Xavier Gabaix. Power Laws in Economics and Finance. *Annual Review of Economics*, 1(1):255–294, 2009.
- Xavier Gabaix. Power Laws in Economics: An Introduction. *Journal of Economic Perspectives*, 30(1): 185–206, 2016.
- Simon Gächter and Christian Thöni. Social comparison and performance: Experimental evidence on the fair wage-effort hypothesis. *Journal of Economic Behavior & Organization*, 76(3):531–543, 2010.
- Danula K. Gamage, Georgios Kavetsos, Sushanta Mallick, and Almudena Sevilla. Pay Transparency Initiative and Gender Pay Gap: Evidence from Research-Intensive Universities in the UK. *IZA Discussion Paper 13635*, 2020.
- Andrew Goodman-Bacon. Difference-in-Differences with Variation in Treatment Timing. *NBER Working Paper 25018*, 2018.
- Michael Z. Green. How the NLRB’s Light Still Shines on Anti-Discrimination Law Fifty Years after Title VII. *Nevada Law Journal*, 14(3):754–769, 2014.
- John Grigsby, Erik Hurst, and Ahu Yildirmaz. Aggregate nominal wage adjustments: New evidence from administrative payroll data. *American Economic Review*, 111(2):428–471, Feb 2021. ISSN 0002-8282. doi: 10.1257/aer.20190318.
- Wulong Gu and Peter Kuhn. Centralization and Strikes. *Labour Economics*, 5(3):243–265, 1998.

- Andreas Gulyas, Sebastian Seitz, and Sourav Sinha. Does Pay Transparency Affect the Gender Wage Gap? Evidence from Austria. *Mimeo*, July 2020.
- Robert E. Hall and Paul R. Milgrom. The Limited Influence of Unemployment on the Wage Bargain. *American Economic Review*, 98(4):1653–1674, 2008.
- Ariane Hegewisch, Claudia Williams, and Robert Drago. Pay Secrecy and Wage Discrimination. *Institute for Women’s Policy Research*, 2011.
- Johannes Hörner and Nicolas Vieille. Public vs. Private Offers in the Market for Lemons. *Econometrica*, 77(1):29–69, 2009.
- International Labour Organization. Global wage report 2018/19: What lies behind gender pay gaps. Technical report.
- Kenneth L. Judd. The law of large numbers with a continuum of IID random variables. *Journal of Economic Theory*, 35(1):19–25, 1985.
- Ayça Kaya and Qingmin Liu. Transparency and price formation. *Theoretical Economics*, 10(2):341–383, 2015.
- Marlene Kim. Pay Secrecy and the Gender Wage Gap in the United States. *Industrial Relations*, 54(4):648–667, 2015.
- David M. Kreps and Robert Wilson. Reputation and imperfect information. *Journal of Economic Theory*, 27(2):253–279, 1982.
- Peter Kuhn and Wulong Gu. A Theory of Holdouts in Wage Bargaining. *American Economic Review*, 88(3):428–449, 1998.
- Peter Kuhn and Wulong Gu. Learning in Sequential Wage Negotiations: Theory and Evidence. *Journal of Labor Economics*, 17(1):109–140, 1999.
- Bradley Larsen. The Efficiency of Real-World Bargaining: Evidence from Wholesale Used-Auto Auctions. *Review of Economic Studies*, 88(2):851–882, 2021.
- Bradley J. Larsen and Anthony Lee Zhang. A Mechanism Design Approach to Identification and Estimation. *NBER Working Paper 24837*, 2018.
- Edward P. Lazear. Retail Pricing and Clearance Sales. *American Economic Review*, 76(1):14–32, 1986.
- Wolfgang Leininger, Peter B. Linhart, and Roy Radner. Equilibria of the Sealed-Bid Mechanism for Bargaining with Incomplete Information. *Journal of Economic Theory*, 48:63–106, 1989.
- 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, August 2017.
- Kyle McCarthy. Pay Secrecy? 60 Percent in Tech Told Not to Discuss Compensation. *Blind Workplace Insights*, Jul 2018. URL <https://www.teamblind.com/blog/index.php/2018/07/16/pay-secrecy-60-of-tech-workers-say-yes/>.
- Ian M. McDonald and Robert M. Solow. Wage Bargaining and Employment. *American Economic Review*, 71(5):896–908, 1981.
- Claudio Michelacci and Javier Suarez. Incomplete Wage Posting. *Journal of Political Economy*, 114(6):1098–1123, 2006.
- Paul Milgrom and John Roberts. Predation, reputation and entry deterrence. *Journal of Economic Theory*, 27(2):280–312, 1982.
- Tomasz Obloj and Todd Zenger. The influence of pay transparency on inequity, inequality, and the performance-basis of pay in organizations. *Mimeo*, 2020.

- 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. *American Economic Review*, 110(4):1019–54, 2020.
- 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.
- Phillip Reese. See how far union membership has declined in California. *The Sacramento Bee*, 2019.
- Juliana Feliciano Reyes. Why a Google spreadsheet was the most powerful tool for labor in 2019. *The Philadelphia Inquirer*, Dec 2019. URL <https://fusion.inquirer.com/business/google-spreadsheet-salary-organizing-baristas-journalists-lawyers-20191231.html>.
- Fernando Rios-Avila. Recentered influence functions (rifs) in stata: Rif regression and rif decomposition. *The Stata Journal*, 20(1):51–94, Mar 2020. ISSN 1536-867X. doi: 10/ghqmt6.
- Nina Roussille. The central role of the ask gap in gender pay inequality. *Mimeo*, 2021.
- Steven Ruggles, Sarah Flood, Sophia Foster, Ronald Goeken, Jose Pacas, and Matthew Schouweiler, Meganand Sobek. Ipums usa: Version 11.0 [dataset]. *Minneapolis, MN: IPUMS*, 11, 2021.
- 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*, 2012a.
- Peter Schmidt. Study Finds Continued Growth of Unions for Faculty Members and Graduate Students. *The Chronicle of Higher Education*, 2012b.
- Guido Schwarzer. meta: An r package for meta-analysis. 7:40–45, Dec 2007.
- Fiona Scott Morton. The Strategic Response by Pharmaceutical Firms to the Medicaid Most-Favored-Customer Rules. *The RAND Journal of Economics*, 28(2):269–290, 1997a.
- Fiona Scott Morton. The Interaction between a Most-Favored-Customer Clause and Price Dispersion: An Empirical Examination of the Medicaid Rebate Rules of 1990. *Journal of Economics & Management Strategy*, 6(1):151–174, 1997b.
- Leo K. Simon and Maxwell B. Stinchcombe. Extensive Form Games in Continuous Time: Pure Strategies. *Econometrica*, 57(5):1171–1214, 1989.
- Gary R. Siniscalco, Erin M. Connell, and Chad Smith. State Pay Equity Laws: Where a few go, many may follow. *Mimeo*, 2017.
- Cailin Slattery and Owen Zidar. Evaluating state and local business tax incentives. *Journal of Economic Perspectives*, 34.2(Spring):90–118, 2020.
- Liyang Sun and Sarah Abraham. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 2020.
- Shengwei Sun, Jake Rosenfeld, and Patrick Denice. On the Books, Off the Record: Examining the Effectiveness of Pay Secrecy Laws in the U.S. *Institute for Women’s Policy Research*, 2021.
- Yeneng Sun. The exact law of large numbers via Fubini extension and characterization of insurable risks. *Journal of Economic Theory*, 126(1):31–69, 2006.
- Jean Tirole. From Bottom of the Barrel to Cream of the Crop: Sequential Screening with Positive Selection. *Econometrica*, 84(4):1291–1343, 2016.
- Donald M. Topkis. *Supermodularity and Complementarity*. Princeton University Press, 1998.
- Wolfgang Viechtbauer. Conducting meta-analyses in r with the metafor package. *Journal of Statistical Software*, 36(11):1–48, Aug 2010. ISSN 1548-7660. doi: 10.18637/jss.v036.i03.
- Jelle Visser. ICTWSS Database, version 6.0. *Amsterdam Institute for Advanced Labour Studies (AIAS), University of Amsterdam*, 2019.

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

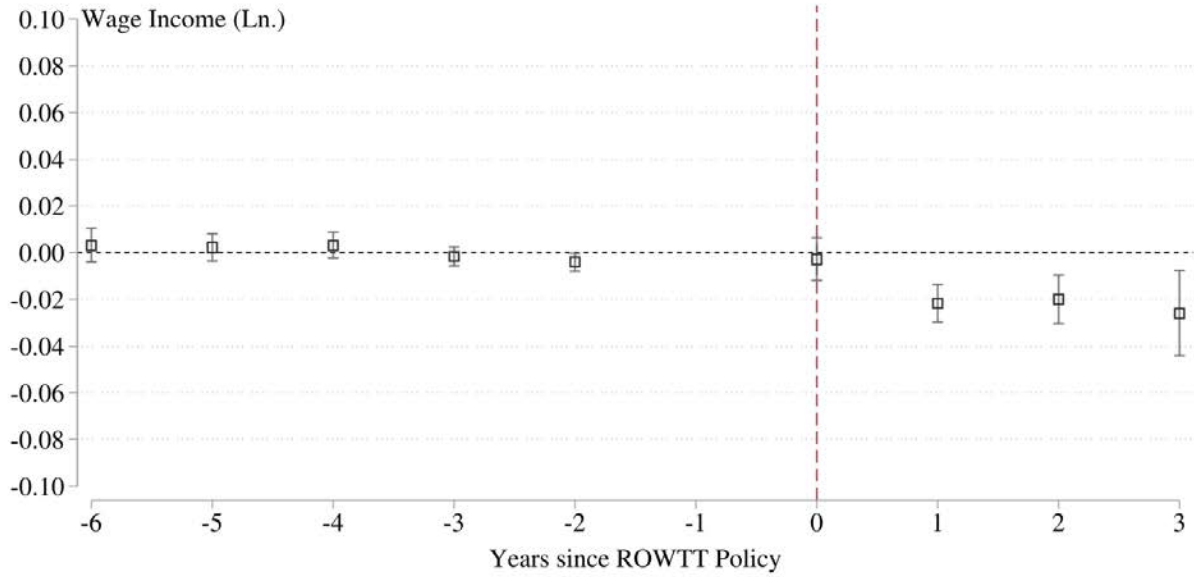
FIGURE I: YEAR RIGHT OF WORKERS TO TALK (ROWTT) LAW IS EFFECTIVE



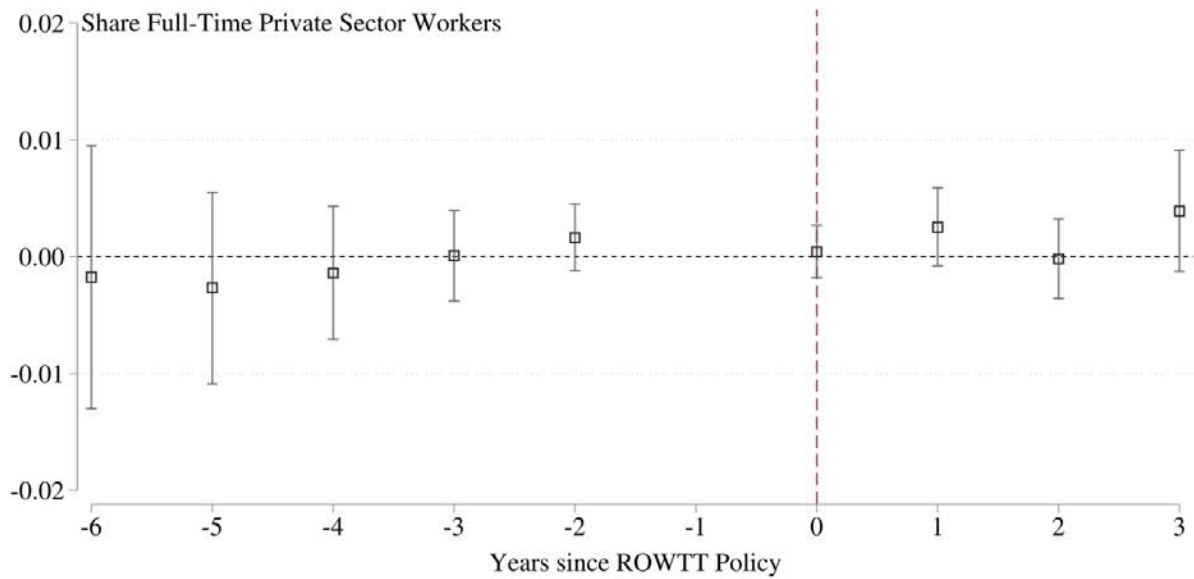
Note: This figure displays the set of states enacting Right of Workers to Talk (ROWTT) policies prior to and including 2016, when a federal ROWTT came into effect.

FIGURE II: EFFECT OF ROWTT POLICIES ON WAGE INCOME AND EMPLOYMENT

PANEL A: WAGE INCOME



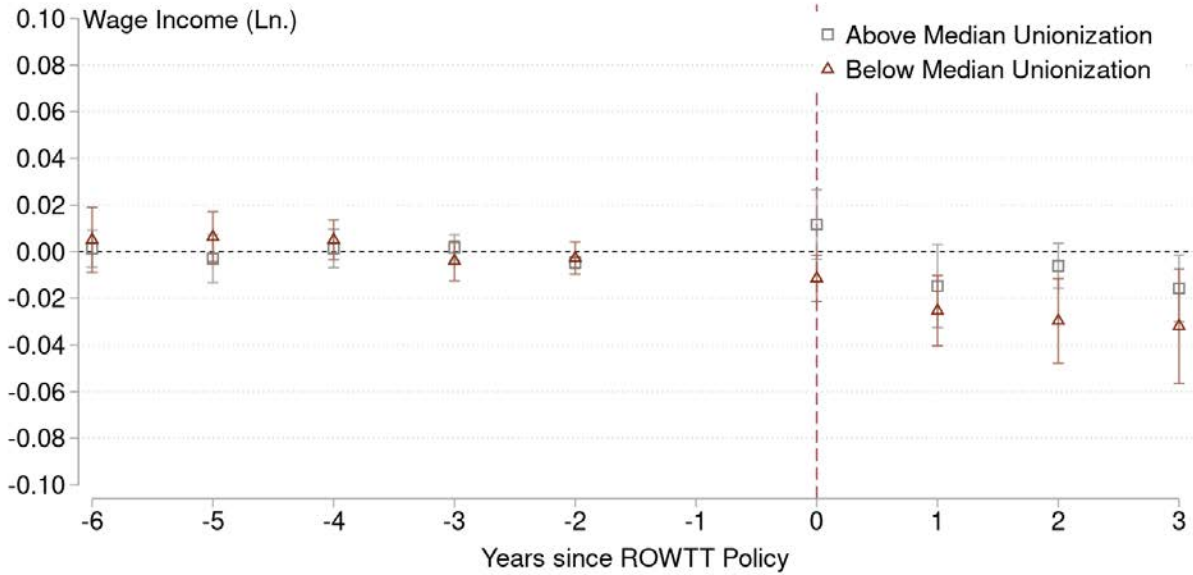
PANEL B: EMPLOYMENT



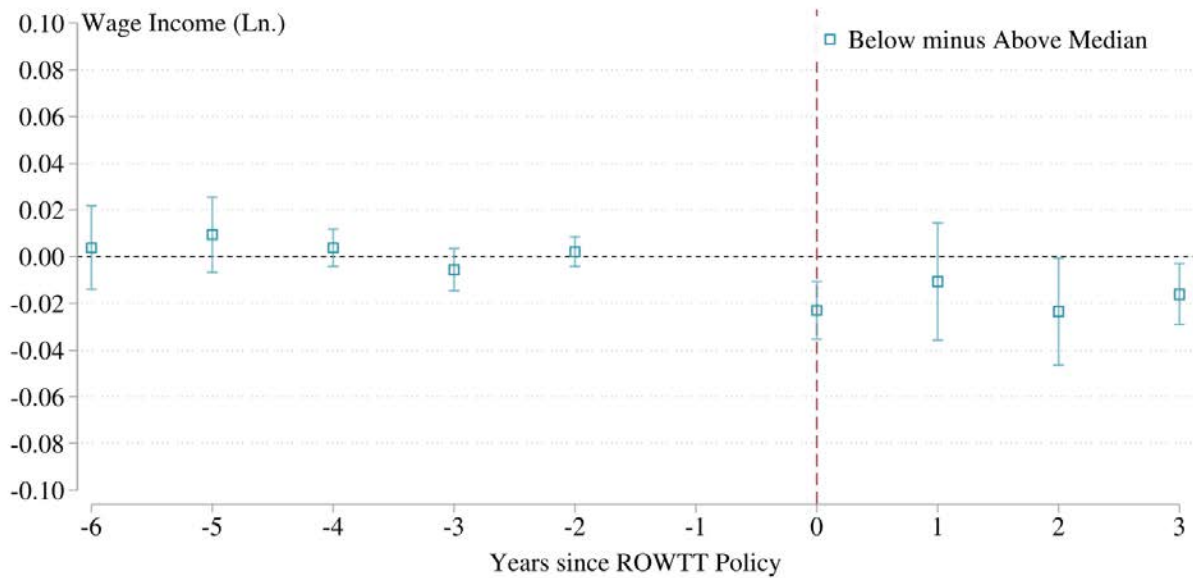
Note: In this figure, we present our baseline multiperiod difference-in-difference estimates. In this baseline specification, we report the results from a balanced composition of states following the enactment of the law. Thus, we estimate the dynamic post period effects for states with events through 2013 separately and report these in periods 0 to +3. See Equations 7 and 9 for more information on this specification. The standard deviation of the state-level mean from 2000 to 2016 is 0.103 for the natural logarithm of wage income and 0.016 for the share of full-time private sector workers.

FIGURE III: HETEROGENEOUS EFFECTS OF ROWTT POLICIES ON WAGE INCOME:
HIGH AND LOW UNIONIZATION

PANEL A: UNION SPLIT, BELOW- AND ABOVE-MEDIAN UNIONIZATION RATES

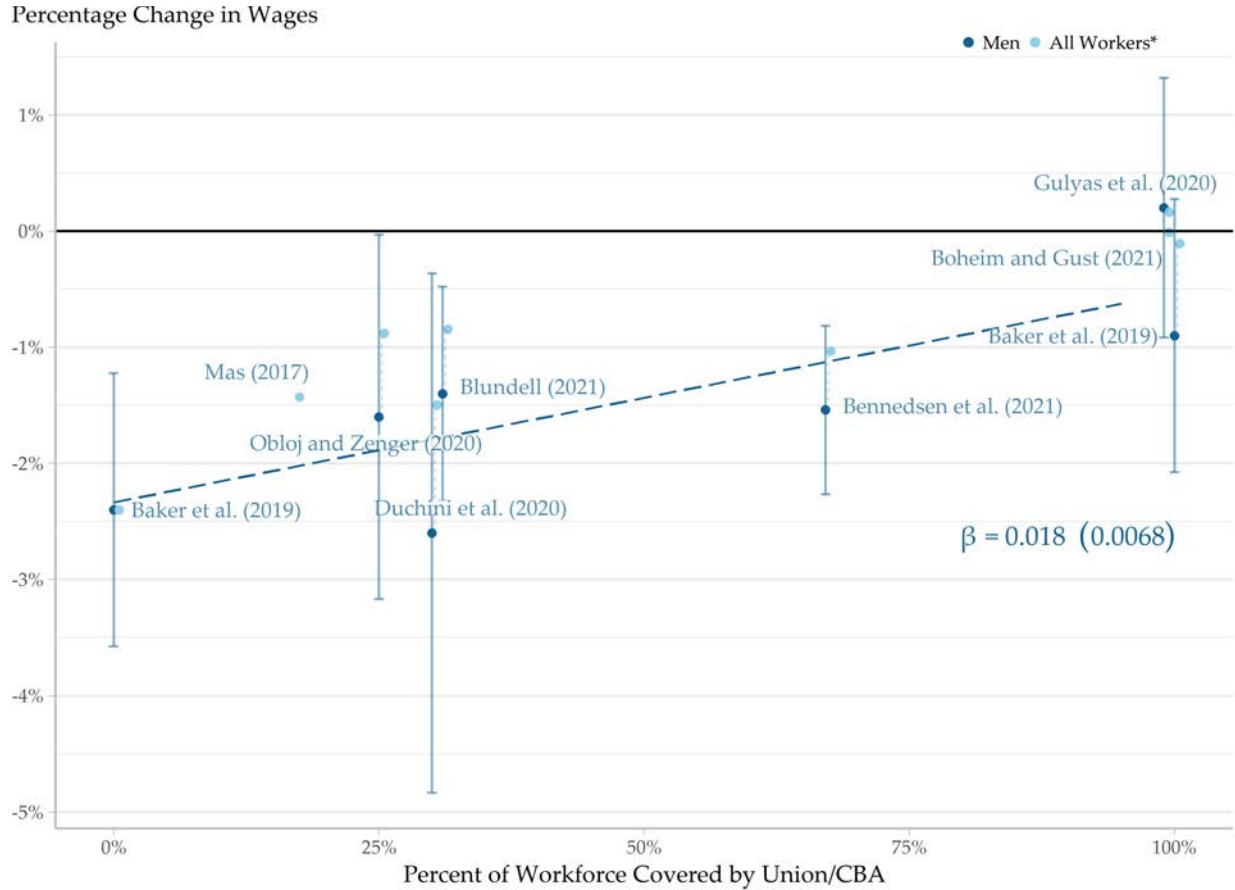


PANEL B: UNION DIFFERENCES, BELOW- vs. ABOVE-MEDIAN UNIONIZATION RATES



Note: In this figure, we present our baseline multi-period difference-in-difference estimates from a balanced composition of states following the enactment of the law. Thus, we estimate the dynamic post period effects for states with events through 2013 separately and report these in periods 0 to +3. See Equations 7 and 9 for more information on this specification. The standard deviation of the state-level mean from 2000 to 2016 is 0.103 for the natural logarithm of wage income and 0.016 for the share of full-time private sector workers. We use data from the Current Population Survey to estimate the share of workers covered by a union or collective bargaining agreement at the occupation level each year and split at the median occupation.

FIGURE IV: EFFECT OF TRANSPARENCY ON WAGES BY INDIVIDUAL BARGAINING POWER, EXISTING STUDIES



Note: In this figure, we graphically present estimates from the related literature. For the majority of studies, we plot two observations, one for the effect of transparency on the wages of men (dark blue series), and one for the imputed effect of transparency on the wages of all workers (light blue series). The x-axis represents the share of workers covered by a union/collective bargaining agreement, and the y-axis the percentage change in wages. We report the estimated effect of the unionization rate on the impact of pay transparency recovered from a mixed-effects meta-regression model (Schwarzer, 2007; Viechtbauer, 2010). Since the estimates for all workers are imputed for some studies, we only report the meta-regression results for the male series for which standard errors are known and displayed. In Section IV we describe the criteria for inclusion in our analysis, and provide the details of each study in Table A.1.

TABLE I: DYNAMIC EFFECT ESTIMATES: ALTERNATIVE SPECIFICATIONS

| | Multiperiod DID Estimator | | | | Sun-Abraham IW Estimator |
|---|---|-------------------|-------------------|-------------------|-----------------------------|
| | PANEL A: WAGE INCOME (LN.) | | | | |
| Mean Pre-Treatment Estimate | -0.001 (0.001) | -0.001 (0.000) | -0.001 (0.001) | -0.002 (0.000) | -0.003 (0.000) |
| Dynamic Post Treatment Effect Estimates | | | | | |
| $t = 0$ | -0.003 (0.005) | -0.005 (0.003) | -0.003 (0.006) | -0.007 (0.006) | -0.017 (0.001) |
| $t = 1$ | -0.022 (0.004) | -0.016 (0.006) | -0.020 (0.008) | -0.013 (0.006) | -0.022 (0.002) |
| $t = 2$ | -0.020 (0.005) | -0.014 (0.007) | -0.019 (0.008) | -0.015 (0.006) | -0.021 (0.002) |
| $t = 3$ | -0.026 (0.009) | -0.026 (0.011) | -0.025 (0.015) | -0.027 (0.012) | -0.028 (0.003) |
| Mean Effect, $t \geq 0$ | -0.018 (0.005) | -0.015 (0.006) | -0.017 (0.009) | -0.015 (0.007) | -0.022 (0.002) |
| Mean Difference: Post – Pre | -0.016 (0.004) | -0.014 (0.005) | -0.015 (0.009) | -0.014 (0.007) | -0.019 (0.002) |
| Balanced Post-Period | Yes | No | Yes | Yes | Yes |
| Year-by-Region FE | No | No | Yes | No | No |
| Fix Gender-by-Education Composition | No | No | No | Yes | No |
| N | 2,306,833 | 2,306,833 | 2,306,833 | 2,306,752 | 2,306,833 |
| Mean (within state SD) | 10.770 (0.103) | 10.770 (0.103) | 10.770 (0.103) | 10.770 (0.103) | 10.770 (0.103) |
| | PANEL B: SHARE FULL-TIME PRIVATE SECTOR WORKERS | | | | |
| Mean Pre-Treatment Estimate | 0.000 (0.001) | 0.000 (0.001) | -0.000 (0.001) | -0.000 (0.001) | 0.000 (0.000) |
| Dynamic Post Treatment Effect Estimates | | | | | |
| $t = 0$ | 0.000 (0.001) | -0.001 (0.002) | -0.000 (0.002) | -0.003 (0.002) | 0.002 (0.002) |
| $t = 1$ | 0.003 (0.002) | 0.002 (0.002) | -0.001 (0.002) | 0.002 (0.002) | 0.001 (0.001) |
| $t = 2$ | -0.000 (0.002) | 0.003 (0.004) | 0.000 (0.004) | 0.005 (0.002) | -0.002 (0.001) |
| $t = 3$ | 0.004 (0.003) | 0.003 (0.004) | 0.002 (0.003) | 0.001 (0.003) | 0.002 (0.001) |
| Mean Effect, $t \geq 0$ | 0.002 (0.001) | 0.002 (0.001) | 0.000 (0.002) | 0.001 (0.002) | 0.001 (0.001) |
| Mean Difference: Post – Pre | 0.001 (0.001) | 0.002 (0.002) | 0.001 (0.001) | 0.002 (0.002) | 0.001 (0.001) |
| Balanced Post-Period | Yes | No | Yes | Yes | Yes |
| Year-by-Region FE | No | No | Yes | No | No |
| Fix Gender-by-Education Composition | No | No | No | Yes | No |
| N | 4,072,277 | 4,072,277 | 4,072,277 | 4,072,225 | 4,072,277 |
| Mean (within state SD) | 0.566 (0.016) | 0.566 (0.016) | 0.566 (0.016) | 0.566 (0.016) | 0.566 (0.016) |

Note: In Cols. 1-4, we use the standard multiperiod DID estimator to recover the dynamic effect of state-level ROWTT legislation on wage income (Panel A) and on the share of workers in full-time private sector employment (Panel B). For the wage income analysis in Panel A, we explicitly restrict the sample to workers in full-time private sector employment. In Col. 1, we present the baseline model, balancing the set of states identifying the post-treatment dynamic effects by absorbing post-treatment dynamic effect estimates for cohorts with events after 2013. In Col. 2 our estimates includes all cohorts from 2000 to 2016. In Col. 3, we add year-by-region fixed effects to our baseline specification. In Col. 4, we reweight our sample by education-by-gender within each state. We take the year before the policy is enacted as the reference year and estimate the educational distribution of each state separately for men and women. Within each state, we then reweight the sample in each year to match the education-by-gender distribution in that state’s reference year. In Col. 5, we use the [Sun and Abraham \(2020\)](#) interaction-weighted (IW) estimator to allow for heterogeneous treatment effects across cohorts. The IW estimator requires that the last-treated cohort be used as a control group in the absence of never-treated units. Thus, in this specification, the 2016 cohort does not contribute to dynamic effect estimates. We balance the post-treatment estimate by estimating the full set of cohort-specific dynamic effects, but excluding the 2014 and 2015 cohorts from the post treatment interaction-weighted estimates.