

Does Tax Transparency Influence Wage Setting? Evidence on Gender Gaps

Cristiano C. Carvalho Trine Engh Vattø

May, 2025

Abstract

This paper examines the impact on gender wage disparities of an often-overlooked source of pay transparency: the public disclosure of tax information. We leverage a 2001 policy in Norway that made individual tax returns publicly accessible online. Using matched employer-employee data, we show that this abrupt increase in transparency—originally aimed at improving tax compliance—significantly narrowed gender wage gaps within firms. A difference-in-difference design reveals a 2.17 percentage point (8.7%) decline in the within-firm gender wage gap, driven entirely by rising female wages. The reduction was strongest in private sector firms and industries with initially larger gaps, and was more pronounced in municipalities previously lacking access to printed tax lists. We explore the underlying mechanisms and find that wage growth was concentrated among women who changed jobs following the increase in transparency, pointing to heightened job search and mobility as key drivers. Our findings highlight the potential of broad-based transparency policies to reduce gender wage inequality through shifts in both information access and labor market behavior.

1 Introduction:

Pay transparency has become an increasingly prominent topic in public discourse, recognized as a potential policy tool for promoting wage equity. In recent years, a variety of countries have adopted policies aimed at enhancing wage transparency. These range from relatively mild interventions—such as penalizing employers who prevent employees from discussing their salaries—to more stringent measures, including full wage disclosure requirements.¹ One of the most persistent and salient forms of wage inequality that pay transparency policies seek to address is the gender wage gap. Despite considerable convergence in qualifications over time—with women often surpassing men in educational attainment—the earnings gap between men and women remains substantial (Blau and Kahn, 2017). One potential explanation for the persistence of the gender wage gap, often raised by policymakers, is the limited information available to workers about firms’ pay practices.

Although these policies have gained momentum, empirical evidence on the impact of transparency on wage-setting behavior remains limited. This article contributes to the literature by examining the effect of pay transparency on gender wage disparities. We focus on an underexplored source of transparency: the public disclosure of individual tax information. Several countries—including Norway, Sweden, Finland, and Pakistan—have implemented such policies primarily to improve tax compliance, especially among business owners and the self-employed. However, little attention has been given to wage earners, as the potential compliance impact for this group may be limited due to their income being reported by employers. Despite this, public disclosure of tax information could act as a pay transparency mechanism, influencing firms’ wage-setting practices and potentially leading to significant impacts on wage dynamics.

Our identification strategy hinges on a significant national policy change in Norway that enhanced the transparency of individual tax returns. Historically, Norwegians’ tax records have been publicly accessible since the mid-nineteenth century, primarily through in-person visits to local tax offices or city halls. However, in the fall of 2001, this information was digitized and made available online for the first time, allowing anyone with internet access to search the records. The search results typically included the taxpayer’s full name, birth year, city, and postal code, along with details about their net income, total taxes paid, and net worth. Although the search does not provide specific information about taxpayers’ job characteristics, we leverage this sudden shift in public disclosure to identify the effects of increased pay transparency among coworkers.²

¹According to Cullen and Pakzad-Hurson (2021), U.S. policies range from sanctions on employers that restrict salary discussions to mandates requiring employers to inform prospective employees of salary ranges. In the EU, policies vary from full disclosure requirements for firms exceeding a certain size, as in Germany, to the publication of summary statistics—such as the mean, median, and gender pay gap—as in the U.K. and Denmark. Many countries also mandate full wage disclosure for public sector employees.

²Note that searches would not allow individuals to get information on people in the same profession or firm unless they know their peers’ name.

We integrate two primary data sources to construct workers' comparison groups. The first dataset comprises the universe of individual income tax returns for the Norwegian population from 1997 to 2006. It provides detailed information on various income sources—including wages, business and capital income, and government transfers. Importantly, it includes key variables from the publicly accessible tax records, such as taxable income, wealth, and total taxes paid. The second data source is a matched employer-employee dataset for the same period, which allows us to link each worker to their respective firm and, by extension, to their workplace peers. This linkage is crucial for constructing nuanced measures of income distribution—such as a worker's relative position within the firm—and for assessing the extent to which wages account for the income disclosed in the public records.

Although the disclosed information reflects aggregated taxable income from all sources, we begin by showing that it closely aligns with wage levels—except at the very top and bottom of the income distribution—while preserving individuals' relative rankings across the spectrum. Consequently, workers could plausibly use this transparency shock to reassess their position relative to peers. We further document that the previously increasing trend in national wage variance reversed abruptly in 2001, the year tax returns became publicly accessible online. This reversal was primarily driven by a reduction in within-firm variance, while between-firm variance remained largely unchanged. We also observe a marked shift in within-firm gender gaps. Gender disparities in Norway were sizable during the period under study: based on the average gender gap across firms, women earned 26.6% less than men prior to 2001. Following the increase in tax transparency, the gender gap declined sharply to 25.6% in 2001 and continued to narrow, reaching 23.5% by 2006.

As searches only provided information about peers if their names were known, we assume the information shock to be greatest among coworkers and thus analyze within-firm variation in gender gaps. We estimate the effects of increased transparency on the differential outcomes for male and female in their reference groups, using an approach that allows for a difference-in-difference interpretation. In our preferred estimation, we define the reference groups as full-time workers within the same firm. We then focus on the within-reference group variation in the gender gaps by controlling for reference group fixed effects interacted with time and reference group fixed effects interacted with gender. This approach isolates the effect of transparency by accounting for economic shocks that may affect reference groups differently and by controlling for workers' movements between reference groups with different gender gaps. We find that both groups experience similar wage trends prior to the change. However, beginning in 2001, female workers experience a relatively higher increase of 2.17% in their annualized wages compared to male workers, representing a reduction of 8.7% in the gender wage gap, given a baseline gap of 24.9%.

We explore the richness of our data and the nationwide disclosure in tax returns to examine how the reduction in gender wage gaps varies across economic sectors and different industries. We find significant variation in the effects along both dimensions. Regarding the

economic sector, tax transparency is associated with a reduction in gender gaps across all sectors, with the strongest effects observed in the private sector. In this sector, gender gaps decreased by 2.6 percentage points. Firms owned by households experienced a reduction of 1.1 percentage point, while public sector firms saw a reduction of 1.5 percentage points. Regarding the different industries, we find a significant reduction in the gender gaps for 13 out of the 15 aggregated classifications. The reduction in the gender gap in percentage points is positively correlated with the size of the baseline gender gap. For example, the mining and quarrying industry, which had the largest baseline gender gap of 37.4%, saw a reduction of 4.6 percentage points, while the education industry, with the smallest baseline gender gap of 13.2%, experienced a reduction of 1.3 percentage points. An overall correlation reveals that for every 10% increase in the baseline gender gap, there is a corresponding 1 percentage point reduction in the gender gap after the tax transparency is in place.

We leverage two sources of heterogeneity and conduct a triple-difference analysis to provide additional evidence that the transparency shock is driving the observed reduction in gender wage gaps. First, we compare municipalities with and without access to printed versions of the tax lists prior to 2001. Our findings indicate that municipalities lacking prior access experienced a stronger transparency shock, with female workers in these areas receiving an additional 1.2% relative wage increase compared to those in municipalities where the tax lists were already available. Second, using a two-way fixed effects decomposition based on Abowd, Kramarz, and Margolis (1999) (AKM), we estimate the within-firm wage premium gap and examine the effects across industries with high and low average gaps in 2000. We find that the impact of tax transparency is significantly larger in high-gap industries, where female wages increased by an additional 1.8% relative to those in industries with lower baseline disparities. We also show that the parallel trends assumption holds for each gender across the two groups of municipalities over time. This variation in access to printed tax catalogs allows us to estimate the separate effects of tax transparency for male and female workers and demonstrate that the observed reduction in gender wage gaps is driven by rising female wages, rather than declining male wages.

To investigate the mechanisms behind the observed reduction in the gender gap, we fix workers in their 2000 reference groups and follow the same individuals over the subsequent three years. To account for potential differences in male and female career trajectories unrelated to the tax transparency policy, we perform a placebo analysis using 1997 data to adjust all estimated effects for the 2000 cohort based on prevailing labor market trends. This approach isolates within-group variation over time. The results show that, relative to men, women experience a 1.61 percentage point increase in termination rates. However, this increase translates only partially into higher job-to-job transitions, which rise by 0.77 percentage points for women. Consistent with greater job search activity, women also experience a relative wage increase of 0.56% after three years.

We examine whether the observed wage increases are driven by improved offers within

firms or by workers transitioning to better matches in other firms, and the evidence points to the latter as the main driver. To begin, we restrict the sample to individuals who remained with their baseline employer for all three years following the policy implementation. Among this group, female workers exhibit a relatively higher wage growth of 0.36%, though the effect is not statistically significant. In contrast, the effects are significantly larger for workers who changed firms at any point during the subsequent three years. For these movers, female workers experience an additional wage growth of 0.6% relative to males, with these results benchmarked against placebo estimates from the pre-policy period.

An important question is whether the observed reduction in the gender wage gap following increased transparency is merely a consequence of female workers being disproportionately concentrated at the lower end of the wage distribution within their comparison groups. To assess this, we divide workers into deciles based on their position within their firm’s wage distribution and compare their wage growth to those in the fifth decile. We find that only workers in the bottom three deciles experience a statistically significant increase in wage growth relative to the reference group. To further isolate the effect, we re-estimate the reduction in the gender wage gap using alternative specifications that control for wage decile interacted with time and wage decile interacted with time and comparison group fixed effects. In both cases, the magnitude of the estimated gender gap reduction remains virtually unchanged. These findings suggest that while tax transparency does contribute to overall wage compression—reflected in greater wage growth for lower-wage workers, even after adjusting for placebo trends—this effect alone does not fully account for the observed reduction in gender wage gaps. Instead, the evidence points to distinct gender-specific impacts of transparency beyond general changes in the wage structure.

We also exploit variation in access to published tax catalogs and baseline industry wage premium gaps to validate our findings using this identification strategy. Although the coefficients are not statistically significant, due to reduced sample sizes in these subgroup analyses, the point estimates closely match the magnitude of the main effects and align with the hypothesis that greater transparency drives the observed wage changes. Specifically, gender wage gaps appear to decline more in municipalities without prior access to the tax catalog and in industries with higher baseline wage premium gaps.

The paper is organized as follows. Section 2 reviews the related literature. Sections 3, 4, and 5 present the institutional background on the public disclosure of tax information, describe the data, and provide descriptive statistics. Section 6 outlines the empirical strategy for the causal analysis. Section 7 presents the main results and explores the underlying mechanisms. Finally, Section 8 concludes.

2 Related Literature:

This paper contributes to several strands of literature. First, it provides evidence that workers care about how their wages compare to those of their peers. While recent studies have identified negative effects of pay disparities on morale, job satisfaction, and effort, this evidence has largely been limited to small-scale experiments (Bracha, Gneezy, and Loewenstein, 2015; Breza, Kaur, and Shamdasani, 2018), data from specific firms (Cullen and Perez-Truglia, 2022; Dube, Giuliano, and Leonard, 2019), or universities (Card et al., 2012). A key contribution of this study is introducing a large-scale, non-experimental intervention that increases pay transparency nationwide. We offer new insights into workers' and firms' labor market responses in equilibrium, demonstrating the important role of information on wage growth and job turnover.

Second, this paper contributes to the literature on the aggregate effects of pay transparency on wages. Consistent with existing evidence, we find that transparency leads to greater wage compression and reduced gender gaps within firms. Our findings extend prior work in two key ways. First, much of the existing literature focuses on narrow forms of transparency, such as the disclosure of aggregate wage statistics by gender (Böheim and Gust, 2021; Bennedsen et al., 2020; Gulyas, Seitz, and Sinha, *Forthcoming*). Second, studies that examine individual-level disclosures often limit their scope to specific sectors, such as universities (Baker et al., 2021; Card et al., 2012) or public employment (Mas, 2017). By contrast, we study a broad-based transparency shock involving individual-level disclosures and document its heterogeneous effects across sectors and industries—finding the largest reductions in gender gaps in private sector firms and in industries with larger initial disparities.

Closely related to our study is Baker et al. (2021), who also examine an individual-level transparency shock, focusing on university faculty salaries in Canada. They find that transparency laws significantly reduced the gender salary gap by 1.2–2 percentage points, corresponding to a 20–30 percent decline relative to the baseline conditional gap. We find similar point estimates: gender gaps declined by 2.17 percentage points in the overall economy and by 1.3 percentage points in the education industry. However, these reductions represent a smaller share of the initial gender gap in our setting, partly due to the more limited set of covariates available in our empirical design. Consistent with its lower baseline disparities, the education sector experienced smaller reductions in gender gaps compared to other industries. Notably, we show that the narrowing of gender wage gaps is driven by rising earnings among female workers, rather than wage reductions for male workers—contrasting with previous studies that find adverse wage effects associated with tax transparency (Cullen and Pakzad-Hurson, 2021; Baker et al., 2021; Bennedsen et al., 2020).

Third, our results offer insights into the extent of workers' misperceptions about their relative position among peers. While we do not directly observe these misperceptions, we provide a compelling test by showing that the tax transparency policy triggered significant

labor market adjustments—likely stemming from corrected beliefs. This interpretation is consistent with growing evidence that individuals often have substantial errors in assessing their rank in the wage distribution (Cullen and Perez-Truglia, 2022; Kristoffer B, Kreiner, and Stantcheva, 2023; Jäger et al., 2024). One contextual factor for our findings on wage compression is that workers tend to anchor their expectations of peer earnings and outside options on their own wages. As a result, lower-paid workers often overestimate, while higher-paid workers underestimate, their relative position—amplifying the potential for behavioral responses when transparent information becomes available.

Fourth, our results demonstrate that the public disclosure of tax and income information can meaningfully affect labor market inequality. Several countries—including Norway, Finland, Sweden, Iceland, and Pakistan—have adopted tax transparency policies, aiming to deter evasion by increasing reputational costs or the threat of whistle-blowing. Supporting this deterrence channel, Bø, Slemrod, and Thoresen (2015) show that Norway’s 2001 online tax disclosure policy reduced tax evasion among business owners by 2.7 percent. Our findings suggest that detailed tax disclosures can also act as a proxy for pay transparency, producing substantial effects on labor market outcomes. Policymakers should therefore consider not only the tax compliance benefits of such policies, but also their broader implications for wage setting and inequality.

Finally, other studies also use the effects of Norway’s tax disclosure policy. Perez-Truglia (2020) shows that posting tax information on the internet increased the gap in happiness between richer and poorer individuals by 29 percent, and it increased the life satisfaction gap by 21 percent. These results are even more striking, given that we also find evidence of compression in the wage distribution after the returns are posted online. Reck, Slemrod, and Vattø (2022) use the microdata of searches after the 2014 reform and find that social comparisons motivate the bulk of searches rather than tax compliance. Importantly, they also document that 15% of the searches target people in the same employment network even after the search anonymity is not in place anymore. Their results find no evidence that individuals increase their reported income after being targeted.

Closely related to our study is Rege and Solli (2015), who also examine the transparency shock generated by the 2001 online publication of tax returns. They find that the policy led to higher job separation rates and faster wage growth among lower-earning workers relative to higher earners. Our study brings new evidence by showing that, beyond overall wage compression, the transparency shock significantly reduced the gender wage gap. Furthermore, we demonstrate that this reduction cannot be fully explained by wage compression alone, suggesting that transparency had distinct effects on gender dynamics in the labor market.

3 Public Disclosure of Tax Information:

Although Norway has a long tradition of disclosing tax information, de facto access was limited before 2001. Individuals who wanted to view someone else’s tax records had to visit a local tax office or city hall during regular hours, and only within three weeks after the tax assessment was finished. This process was drastically simplified in the fall of 2001 when the government allowed and a national newspaper published all tax reports on its web page. After this, anyone with access to the internet could search for other taxpayers’ information. The websites listed full names and their respective income, wealth and total taxes paid. The searches could be done by multiple fields, such as last names, postal codes, and others.

Importantly, the tax searches do not include information about individuals’ occupations or employers. As a result, individuals can only access information about peers if they already know their names, limiting the ability to compare wages within a specific profession or firm. This restriction also reduces the potential for firms to use the data to poach workers from other employers, as doing so would require them to identify targets in advance. Furthermore, the reported income reflects an aggregate of all taxable income sources after applying standard deductions. While this includes all forms of wage income—such as fixed pay, bonuses, tips, and commissions—these components are not separately identifiable in the public records.

In the years following the initial release, the Norwegian government began introducing regulations on tax searches. In 2004, it prohibited the publication of tax data through channels other than the official search tool provided by the tax agency. Searches were also restricted to a three-week window following the data release, although this limitation was lifted after 2006. In 2011, access was further tightened when users were required to log in using a personal identification number (PIN) and password. The most significant change came in October 2014, when search anonymity was removed—taxpayers could now see who had viewed their information.

Making tax returns available online significantly increased their accessibility for a large share of the Norwegian population. According to the Norwegian Monitor Survey, 66% of respondents aged 25 to 65 reported having internet access at home in 2001—a figure that steadily increased to 98% by 2011. While no data exist on the early popularity of the online tax lists, a 2007 survey by Synovate found that 40% of respondents had used online tools to search for tax information. Notably, the survey revealed that many individuals used the searches to learn about their employment networks: 26% reported searching for a work colleague, making it the fourth most common reason after searching for a close relative (61%), oneself (53%), and friends (42%) (Perez-Truglia, 2020). The tendency to search within one’s employment network is further supported by administrative search data. Reck, Slemrod, and Vattø (2022) report that even after search anonymity was removed, 15% of non-self searches still targeted individuals within the same employment network.

4 Data:

We explore two main data sources in this project. The first set includes the universe of individual income tax returns for the Norwegian population from 1997 to 2006, which allows us to observe different sources of income, such as wage, business and capital income, and government transfers. Importantly, this data also contains the variables from the income tax return posted online: ordinary income, wealth, and total taxes paid. Even though this data set allows us to follow individuals over time, it does not bring additional information on employees' labor market outcomes apart from the total wage income in a specific year.

To complement the tax returns information and get a full picture of the labor market, we use the matched employer-employee information from 1997 to 2006. This data contains the universe of workers matched with their firms, which allows us to link each worker to their peers. It includes information about the job, such as wages, firms' industry, and employees' type of contracts (part or full-time), and about the workers, such as gender, education level, municipality of residence, and age.

We combine both datasets by unique individual identification numbers, resulting in a final sample of workers for whom we observe both tax records and employer information. Our final sample focuses on all full-time workers aged 25 to 60 years old. We focus on workers who were employed in a single firm in December of each year and normalize their wages to the monthly level by dividing the total wage received in the year by the number of days the contract was active and multiplying by 30.³ The normalized monthly wage is used to determine a worker's relative position within their firm in our wage compression analyses.

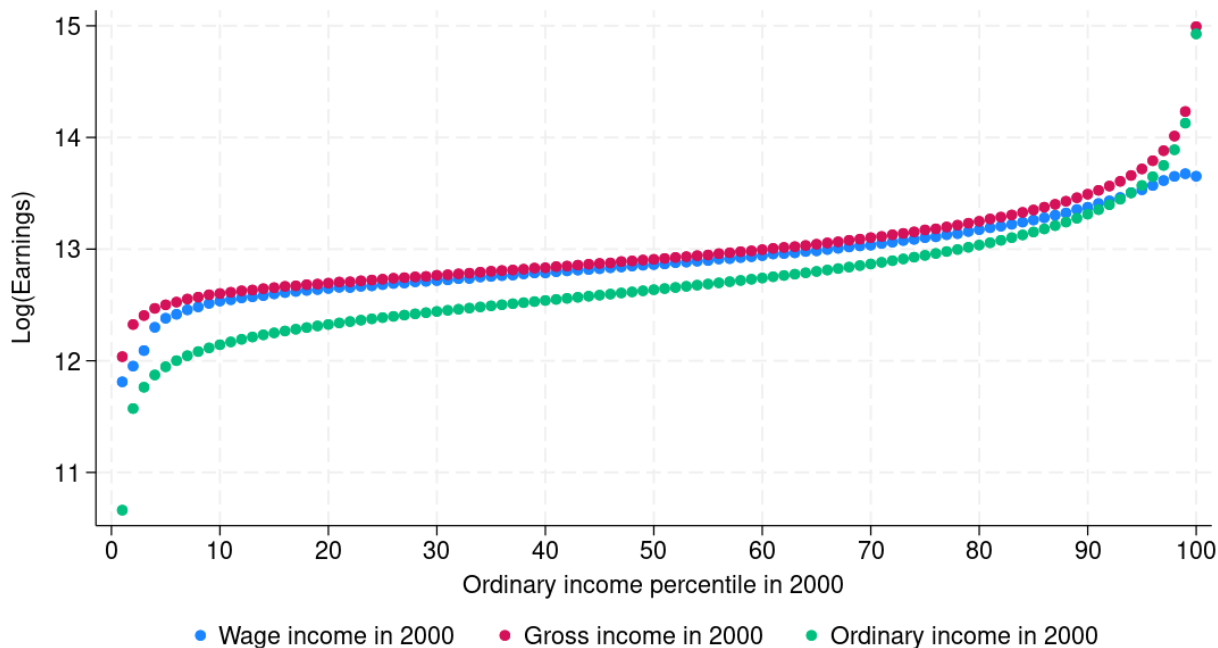
5 Descriptive Statistics:

5.1 Tax Information as a Proxy for Wage:

One important aspect of our context is that the magnitude of the transparency shock depends on how well the income variable posted online is associated with workers' wages, especially because the ordinary income measure includes all sources of income and tax deductions. As we show in figure 1, the ordinary income measure preserves the same ranking property as the wage income, which means that even though individuals cannot precisely know the wage income of other people, they can still use the measure to compare their relative position. Additionally, except for people on the percentile extremes, the wage income constitutes almost entirely the gross income, which implies that the role of other sources of income is limited for our sample. More precisely, from the percentile 5 to the percentile 91, the wage income represents more than 90% of the gross income.

³We impose the single firm requirement to make the ordinary income a better approximation for wages in a specific job.

Figure 1: Wage Variables by Ordinary Income Percentiles in 2000:



Note: First, we use the selected sample to order workers in percentiles of the national distribution using their ordinary income in 2000. Second, we plot the average ordinary income, gross income, and wage income for each percentile bin to check the ranking order is maintained in the other variables.

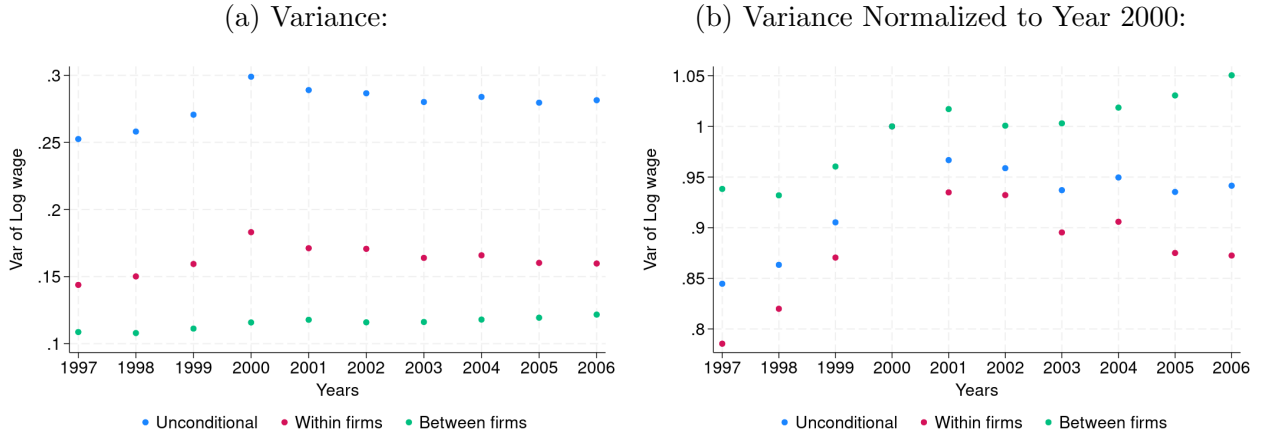
5.2 Trends in Wage Variance and Gender Gaps:

One potential concern is that within-firm wage variation would be of least importance since Norway, like other Nordic countries, has a high prevalence of union agreements. In figure 2a, we show the time trend for the wage variance in Norway, decomposing it into within and between firms variance. First, from 1997 to 2006, the within-firm variance is more important in explaining the total variance of wages (around two-thirds), which indicates that firms have the discretion to set their workers' wages. Second, the wage variance had a positive trend up to 2000; in 2001, when taxes were posted on the internet, the variance interrupted the positive trend and dropped. Importantly, the decrease in the variance that occurred after the tax transparency shock was driven by the within-firm variation as presented in 2b.

Gender gaps in Norway were also relatively high in the period considered in our analysis. Figure 3 presents the average gender gaps in each comparison group that had both men and women workers weighted by the group size. This measure is close to the within-comparison group results presented in the following sections since it considers the general gender gaps if men and women were equally distributed among different comparison bins.⁴ On average, women earn around 26.6% less than men before 2001. With an increase in tax transparency in 2001, the gender gap sharply decreases to 25.6% and keeps decreasing until 2006, when it reaches 23.5%.


⁴The comparison groups are defined as all workers in the same firm, same education level, and same tenure bracket.

Figure 2: Wage Variance Within and Between Firms:



Note: In each year, we decompose the total variance of the wages into within and between firms components in the following way: $\frac{1}{N-1} \sum_{i=1}^N (w_{ift} - \bar{w}_t)^2 = \frac{1}{N-1} \sum_{i=1}^N (w_{ift} - \bar{w}_{ft})^2 + \frac{1}{N-1} \sum_{i=1}^N (\bar{w}_{ft} - \bar{w}_t)^2$. Where N is the total number of workers in a specific year, w_{ift} is the individual wage, \bar{w}_{ft} is the firms' average wage, and \bar{w}_t is the unconditional average wage in a specific year.

Figure 3: Gender Gaps by Comparison Group:



Note: We first calculate the gender gaps for each specific comparison group, which is defined as all workers in the same firm with the same education level and tenure bracket. We focus on comparison groups with both men and women and use the total number of workers in each comparison group as a weight to calculate the yearly average gender gap.

6 Empirical Strategy: Tax Transparency and Gender Gaps:

To assess the causal effect of tax transparency on gender gaps, we exploit the 2001 reform to measure its differential impact on male and female workers. We use two different strategies to assess the impact of greater transparency on wage compression and gender gaps. First, we measure the general differential effects in the labor market by gender within the current comparison group, allowing workers to change their reference groups over time. This approach helps us document how inequality within the comparison group evolves under the analyzed policy. Second, we fix the worker’s comparison group as of 2000 and follow the same workers over three years to examine whether they exhibit differential responses after learning about their relative position due to the increased transparency.⁵ This approach helps us identify the mechanisms at play.

We focus on measuring how the policy affected within-firm gender gaps in the different estimation strategies.

The first specification relies on:

$$\log(w_{i,g,t}) = \sum_{j \neq 2000} \delta_j \cdot 1\{t = j\} \cdot \text{Female}_i + \mu_{t,g} + \mu_g \cdot \text{Female}_i + \varepsilon_{i,g,t} \quad (1)$$

where $\log(w_{i,g,t})$ denotes the logarithm of the annualized earnings of worker i in comparison group g in year t . The variable Female_i is a dummy equal to one for female workers. While we explore different definitions of comparison groups throughout the paper, our main specification defines the comparison group as all workers within the same firm. The term $\mu_{t,g}$ captures comparison group-by-year fixed effects and controls for macroeconomic shocks that differentially affect groups over time. The interaction term $\mu_g \cdot \text{Female}_i$ introduces group fixed effects by gender, accounting for changes in the gender gap that arise from worker movements across comparison groups. The regression has a difference-in-differences interpretation, where the term $\sum_{j \neq 2000} \delta_j \cdot 1\{t = j\} \cdot \text{Female}_i$ captures the leads and lags of the gender gap change relative to the baseline year 2000. The coefficients δ_j are the parameters of interest. Since tax reports became available online in 2001, the coefficients estimated for the years prior to the policy change serve as a placebo test and should be statistically insignificant if the identification strategy is valid. After the policy change, a positive coefficient indicates that increased tax transparency led to higher wage growth for females relative to males, thereby narrowing the gender gap. With $\mu_{t,g}$ included in the specification, δ_j is identified from within-group variation in gender gaps over time. As a result we effectively compare men and women as if they had the same fixed distribution across comparison groups.

To further assess whether the observed effects are driven by increased pay transparency, we exploit two sources of heterogeneity. First, we classify industries into high and low gender wage premium gap groups using 2000 estimates from a time-varying version of an additive

⁵One year before taxes were posted on the internet.

two-way fixed effects model for workers and firms (a time-varying version of Abowd, Kramarz, and Margolis (1999)) as described in the Appendix, and compare changes in gender gaps across these groups. Second, we estimate differential effects for municipalities with and without widespread access to printed versions of the tax list prior to its online release in 2001.⁶ We expect larger effects in municipalities without prior printed access—where the transparency shock was greater—and in industries with higher initial gender wage disparities, where the newly available information revealed greater inequality.

The results are estimated using the following triple difference style equation:

$$\begin{aligned} \log(w_{i,g,h,t}) = & \sum_{j \neq 2000} \delta_j \cdot 1\{t = j\} \cdot \text{Female}_i \cdot \text{het}_{i,h,t} + \sum_{j \neq 2000} \gamma_j \cdot 1\{t = j\} \cdot \text{Female}_i + \\ & + \mu_{t,g} + \mu_g \cdot \text{Female}_i + \varepsilon_{i,g,h,t} \end{aligned} \quad (2)$$

where $\text{het}_{i,h,t}$ represents a dummy variable for the different heterogeneity groups. In the first exercise, $\text{het}_{i,h,t}$ is a dummy variable that equals 1 for peer groups in industries with wage premium gap above the median before the legislation. In the second exercise, $\text{het}_{i,h,t}$ is a dummy variable that equals 1 for municipalities without printed versions of the tax catalog before the legislation. Notice that the comparison group fixed effects for each year are more granular than the two sources of heterogeneities, which allows us to omit some of the triple difference terms.⁷

We implement an alternative strategy to analyze changes in gender wage gaps by fixing comparison groups in the year 2000 and tracking the same individuals over time. To further address concerns about mean reversion, we use data from the first year of our sample (1997) to assign individuals' treatment status and comparison groups as a placebo cohort for additional control. In this analysis, we restrict the sample to workers with more than one year of tenure at their firms, as the publicly released information—based on annual income—would be more accurate and informative for these individuals.

The specification is presented below:

$$\begin{aligned} y_{i,(g,t_0),c,t} = & \sum_{j=1}^3 \delta_j \cdot 1\{t - t_0 = j\} \cdot \text{Female}_i \cdot 1\{c=2000\} + \sum_{j=1}^3 \gamma_j \cdot 1\{t - t_0 = j\} \cdot \text{Female}_i + \\ & \beta_1 \cdot \text{Female}_i \cdot 1\{c=2000\} + \beta_2 \cdot \text{Female}_i + \mu_{t,(g,t_0),c} + \varepsilon_{i,(g,t_0),c,t} \end{aligned} \quad (3)$$

⁶We follow the classification used by Bø, Slemrod, and Thoresen (2015), who document that in some municipalities, local groups distributed printed tax records before they were made publicly available online.

⁷The least granular comparison group is defined as all workers in a firm. Since in our sample, we only keep firms that did not change industries and municipalities over time, industry and municipality fixed effects will be colinear with firm fixed effects.

where $y_{i,g,t}$ represents the logarithm of the earnings or a dummy for separation from the firm for worker i in comparison group (g, t_0) , in cohort c , and in year t . The comparison groups are fixed in 1997 or 2000, and the same groups are followed irrespective of their separation decisions. In the main specification, the comparison group is defined as all workers in the same firm in the baseline year t_0 of their cohort c . $\mu_{t,(g,t_0),c}$ represents comparison group fixed effects for each year in the sample. We include individual fixed effects in our main wage specifications, whereas separation effects are estimated using cross-sectional variation. The coefficients δ_j are the parameters of interest and have a triple-difference interpretation, as they capture the change in the gender gap for each year relative to the baseline, net of the corresponding change observed in the placebo cohort. The estimation relies on the assumption that the 1997 cohort is representative of the counterfactual evolution of gender gaps in the absence of the policy change.

We also estimate fourth difference style coefficients by allowing different effects by the heterogeneous groups. We estimate different effects for switchers and stayers, high and low wage premium gap firms in the baseline period, and firms in municipalities with and without printed versions of the tax catalog before the legislation.

7 Causal Results:

7.1 General Effects of the Public Disclosure:

We begin our analysis by presenting the event-study coefficients, δ_j , estimated from the regression model in Equation 1, which illustrate the effects of tax transparency on the gender gap over time.

Figure 4 shows the coefficients from models with varying sets of controls. In the specification without controls, the overall gender gap in the Norwegian economy begins to decline around 2000, showing some evidence of a pre-trend likely driven by compositional effects. However, once we account for firm composition by controlling for reference group fixed effects and focusing on within-firm changes in gender gaps, the results indicate that the gender gap was stable in the pre-period. After tax returns were posted online in 2001, we observe a sharp and sustained increase in wage growth for females relative to males, which persists throughout the sample period. The event-study patterns are very similar whether we define the reference group as all coworkers or restrict it to coworkers with the same level of education and tenure.⁸ The qualitative results remain similar when we further account for worker composition by adding individual fixed effects on top of coworker fixed effects and accounting for differential trends in the pre-period. In Figure 4, we present the results without accounting for different wage growth pre-trends, and the results show that the gender gap was increasing

⁸We group education into three levels: (i) up to completed upper secondary education; (ii) college or equivalent, including supplementary courses; and (iii) master's degree or higher. Tenure is categorized into two groups: workers with less than one year and those with more than one year in their current job.

until 2000. However, this trend was interrupted in 2001 when the tax returns were posted online.

The regression results are presented in Table 1. The difference-in-differences average estimates across all periods, based on models with varying sets of controls, are shown across columns. Column (1) reports results from the specification without any additional controls. Columns (2) to (4) introduce reference group-by-time and reference group-by-female fixed effects. In Columns (2) and (4), the reference group consists of all coworkers within the same firm, while in Column (3), it consists of coworkers within the same firm who share the same education level and tenure. Additionally, Column (4) incorporates individual fixed effects and adjusts for differential linear trends using pre-period data. Across all specifications, the results consistently show that tax transparency is associated with a reduction in gender gaps. We begin by discussing the results from specifications (2) to (4), as their corresponding event-study graphs support the parallel trends assumption, allowing for a causal interpretation. The estimates are statistically significant at the 1% level and indicate that higher transparency reduced the gender gap by 1.91 to 2.68 percentage points, relative to a conditional gender gap that ranges from 22.3% to 24.9%. This effect implies that the gender gap decreased by approximately 8.5 to 10.7 percent relative to its initial level.

In the following analysis of the paper, we focus on specification (2), which defines the reference group as all coworkers within the same firm. We adopt this parsimonious approach without individual fixed effects, as in Section 6.3 we follow the same individuals over time and directly control for individual composition. All subsequent analyses include firm-by-time and firm-by-female fixed effects. Leveraging the richness of our data and the nationwide implementation of the policy, we next explore the differential effects across sectors and industries. The results by sector are presented in Table 2. Column (1) reports the baseline results using the full sample, while columns (2), (3), and (4) present results separately for private sector firms, firms owned by households, and public sector firms, respectively. The results show that tax transparency is associated with a reduction in gender gaps across all sectors, with the strongest effects observed in the private sector. In this sector, gender gaps decreased by 2.6 percentage points, corresponding to an 8.8 percent reduction relative to the baseline conditional gender gap of 29.3%. Firms owned by households experienced a reduction of 1.1 percentage point, or 5.7%, while public sector firms saw a reduction of 1.5 percentage points, or 8.2%.⁹

We further present the results for each of the 15 industry-aggregated classifications. In Figure 6 and Table 3, we report the difference-in-differences point estimates, stratified by the size of the baseline conditional gender gap for each industry in our sample. The point estimates are positive and statistically significant in 13 out of the 15 industry categories.

⁹Household firms is a classification used by Statistics Norway that broadly includes associations of households producing goods and services for their own use—such as housing cooperatives and condominium associations—as well as self-employed individuals whose business activities are inseparable from their household’s other economic activities. In this context, the category of self-employed is broad and includes firms where the owner is not necessarily the sole worker.

The reduction in the gender gap in percentage points is positively correlated with the size of the baseline gender gap. For example, the industry with the largest baseline gender gap of 37.4% — mining and quarrying — experienced a reduction of 4.6 percentage points, while the industry with the smallest baseline gender gap of 13.9% — education — saw a reduction of 1.3 percentage points. A simple correlation shows that a 10% increase in the baseline gender gap is associated with a 1 percentage point decrease in the gender gap after the policy implementation.

7.2 Heterogeneity by Municipality and Industry Gender Gap:

We further explore the heterogeneous effects of the policy by identifying groups potentially more exposed to pay transparency, using the regression model outlined in Equation 2. First, we leverage pre-period geographic variation in tax transparency levels, as documented by Bø, Slemrod, and Thoresen (2015). Prior to the online publication of tax reports in 2001, some municipalities distributed printed versions of tax lists. In certain areas, local groups—such as football club members or community bands—would sell copies of residents’ tax records door-to-door. We exploit this variation to compare changes in gender wage gaps across municipalities with and without such printed disclosures. Second, we classify industries into high and low gender wage premium gaps categories based on a time-varying AKM model estimates described in the Appendix using the 2000 premiums data, and then compare gender gaps changes in these groups. We choose to conduct this classification at the industry level rather than the firm level to avoid potential spurious effects driven by mean reversion.

We present the aggregated results for both heterogeneity analyses in Table 4. For each exercise, we estimate the model both with (Columns (2) and (4)) and without (Columns (1) and (3)) firm-by-time fixed effects. These heterogeneity specifications introduce a triple-differences framework, allowing us to identify differential difference-in-differences effects by gender. However, the individual effects for each gender can only be separately identified in the models that omit firm-by-time fixed effects. Provided that the parallel trends assumption holds, identifying the difference-in-differences effects separately for males and females is crucial, as it enables us to determine whether the overall reduction in the gender wage gap is primarily driven by wage gains among women or wage declines among men following the implementation of tax transparency.

We begin by examining the differential effects between municipalities with and without preexisting tax catalogs, displayed in columns (1) and (2). The triple-difference coefficients reported in the first row of Table 4 indicate that municipalities without prior access to the tax catalog experienced a significantly larger increase in wages for female workers relative to male workers, compared to municipalities where the catalog had already been available before the policy implementation. This pattern is consistent with the information shock being more pronounced in areas not previously exposed to the printed versions. The effects are similar across specifications, with relative wages increasing by 1% and 1.2% in the models without

and with firm-by-time fixed effects, respectively.

The second row of the table shows that the transparency shock also reduced gender wage gaps in municipalities that had already had access to the printed tax catalog, though the magnitude of the effect was smaller. In these municipalities, female wages increased by 1% and 0.9% in the models without and with firm-by-time fixed effects, respectively. The third row reports the estimated change in male wages across municipalities with and without prior access to the tax catalog. The coefficients are close to zero and not statistically significant, suggesting that the reductions in gender wage gaps are not driven by declining wages for male workers following the increase in transparency. The validity of these results depends on the parallel trends assumption holding for each gender in their wage differences across the two groups of municipalities over time. Figure 7a presents the difference-in-difference coefficients for each gender, showing that the wage differences between the groups of cities were constant for each gender prior to the posting of tax returns. The differential wage gap between cities with and without tax catalogs only increased for females after the tax transparency increases. In figures 7b and 7c, we further demonstrate that the parallel trends assumption holds for the triple-difference coefficients in both specifications.

We further investigate whether the information shock has a differential impact based on the average size of the gender wage premium gap within each industry, as shown in Columns (3) and (4) of Table 4. Industries are classified into two groups—those with high and low average within-firm gender wage premium gaps—using the median gap across industries as the cutoff. This classification is based on wage premium gaps estimated by a time-varying version of the additive two-way fixed effects model for workers and firms for the year 2000 as described in the Appendix. The results are also robust to using 1999 estimates. The triple-difference coefficients reported in the first row indicate that industries with larger gender wage gaps experienced significantly greater increases in wages for female workers relative to male workers than industries with smaller gaps, consistent with the transparency shock having a stronger corrective effect where disparities were initially more pronounced. The effects positive and significant in both specifications, with relative wages increasing by 0.8% and 1.8% in the models without and with firm-by-time fixed effects, respectively.

The transparency shock also reduced gender wage gaps in industries with initially smaller disparities, though the magnitude of the effect was more modest, as shown in the second row. In this group, female wages increased by 1% and 1.2% in the models without and with firm-by-time fixed effects, respectively. The third row reports the estimated change in male wages across high- and low-gender wage premium gap industries. The coefficients are significant and negative, which suggests that part of the reduction in gender wage gaps is driven by declining wages for male workers following the increase in transparency, in case the parallel trends assumption was respected. However, as displayed in Figure 8a, the parallel trends assumption is violated for the difference-in-difference coefficients for each gender: average wages for both males and females were already declining in high-gap industries relative to low-

gap industries even prior to the policy. Even though the difference-in-difference estimates by gender are not reliable, the triple-difference coefficients are, as the parallel trends assumption holds for the triple-difference estimates in both specifications, as shown in Figures 8b and 8c.

7.3 Mechanisms:

7.3.1 General Effects:

To explore the mechanisms underlying the observed within-group wage compression and reduction in gender gaps, we fix workers' comparison groups and follow them over time. This approach allows us to assess whether the observed gains stem primarily from improved matching—i.e., workers switching firms in response to new information—or from better offers within the same firm. We begin by examining worker mobility: specifically, how likely individuals are to leave their baseline firms and whether they secure employment at new firms. Figure 9 displays the relative probability that female workers leave their baseline firms compared to male workers in the same comparison group, for both the 1997 and 2000 cohorts, along with the difference between the two. As shown in Figure 9a, women are generally more likely to leave their firms in all three years following the baseline year for both cohorts. Assuming that the 1997 cohort reflects the counterfactual trend for the 2000 cohort in the absence of increased tax transparency, we estimate that transparency increases women's probability of leaving their baseline firms in any period by 1.61 percentage points.

We investigate if the higher probability of leaving is due to a higher job-to-job transition in figure 9b. While female workers are more likely to leave their firms, this behavior only partially results in higher re-employment rates. In both cohorts, female workers have a low probability of being found in a firm different from the one they were working in the baseline period. This implies they have higher chances of experiencing gaps in their employment history. After taking the difference from the job-to-job transition probabilities observed for the placebo cohort, we find an increase in the probability of finding a new firm in any given period of just 0.77 percentage points - about half the size of the effect on leaving the baseline firm. This effects emerge one year after the publication of tax returns. Concurrently, this group also faces an increase in employment gaps.

We continue the analysis by examining the effects on wages over the three years following the increase in tax transparency. Although the publication of tax returns led to some widening of employment gaps, workers may eventually transition into better job matches upon re-entry into the labor market. Figure 10 presents the evolution of gender wage gaps over the three-year period after the baseline. To account for changes in workforce composition, all estimations include individual fixed effects and year-specific comparison group fixed effects, where comparison groups are defined as workers employed in the same firm during the baseline year. This approach also helps to isolate the results from the effects of workers who quit and

never return.¹⁰

Figure 10 reports the relative changes in the gender wage gaps following the increase in tax transparency. First, both sets of estimates show that female workers experience slower wage growth compared to males in their reference group in the years following the baseline period, for both the 1997 and 2000 cohorts. This highlights the importance of controlling for the fact that females and males would have different career trajectories, regardless of changes in tax transparency. This is accounted for by using the 1997 cohort as a placebo group, which allows us to net out the effects estimated for the 2000 cohort. Second, the triple-difference estimates indicate that tax transparency increases the relative wages for females who remained in the labor market in any of the subsequent years, thereby reducing the gender wage gap by making the wage paths of males and females more similar. The average relative wage growth for female workers following increased transparency is 0.44% across all periods, with more pronounced gains of 0.67% in the second year and 0.56% in the third year.

To further understand the underlying dynamics, we examine the wage effects separately for workers who switch firms and those who remain in their baseline firms throughout the analysis period. Specifically, we differentiate between switchers—those who changed firms at any point after the baseline year—and stayers—those who remained at the same firm for all subsequent years. The results are presented in Table 5. As a reference, Columns 1 and 3 display the baseline triple-difference estimates of wage growth for female relative to male workers, capturing effects across all post-baseline years and in the third year alone, respectively—corresponding to the findings discussed in the previous paragraph. Columns 2 and 4 further decompose these effects by incorporating additional interaction terms for switchers in the regression model, allowing us to isolate the effects separately for switchers and stayers across all post-baseline years and in the third year, respectively.

The first-row coefficients show that the average relative wage growth observed for female workers among stayers is -0.15% across all periods and 0.36% in the third year, with neither estimate being statistically significant. The second row indicates that the wage effects relative to males are 1.19% and 0.60% higher for female switchers compared to female stayers over the full post-baseline period and in the third year, respectively. The third row reports the wage effects of switching for male workers: -0.62% across all years and 0.2% in the third year, with only the latter being statistically significant. These results suggest that the benefits of switching increased exclusively for female workers following the disclosure of tax returns and are the primary driver of the observed reduction in the gender wage gap.

¹⁰Estimating with individual fixed effects excludes the variation stemming from workers who exit the labor market permanently. This allows the diff-in-diff coefficients to reflect the wage change for workers who remain in the labor market after the publication of tax returns.

7.3.2 Are Gender Gap Reductions Driven by Wage Compression?

One possible explanation for the reduction in the gender wage gap following increased transparency is that female workers are typically located at the lower end of the wage distribution within their comparison group. In this section, we investigate whether differences in their relative position during the pre-period can account for the entire change in the gender wage gaps. We begin by examining whether tax transparency leads to wage compression among workers within the same firm during the baseline period. To do this, we divide workers into two groups: those earning below the median wage and those earning above it within their baseline firm. We then measure the differential wage effects for these two groups, using a similar approach to the gender gap analysis.

Figure 12 reports the relative wage increases of workers below the median compared to those above the median in their comparison group. First, in both sets of estimates, individuals in the lower position of their comparison group experience a higher wage growth trajectory in the subsequent years for both cohorts. This underscores the importance of controlling for mean reversion, which is addressed by using the 1997 cohort as a placebo to net out the effects estimated for the 2000 cohort. Second, the triple-difference estimates indicate that tax transparency increases wages for workers earning less than their peers, who remained in the labor market during the subsequent years. Although the average relative wage increase across all periods is not statistically significant, workers below the median experience a higher relative wage growth of 0.5% in the third year after the policy.

We further investigate how increased tax transparency affects wage growth across different deciles of the within-comparison group wage distribution. Since dividing relative positions into deciles requires at least ten workers per bin, our analysis is restricted to sufficiently large groups. Figure 13 presents the overall impact of the tax transparency policy on wage growth three years after implementation, comparing workers across different within-firm wage deciles, with those in the fifth decile serving as the reference group. First, we observe that for both cohorts, wage growth tends to decrease with a worker's relative position within the firm—those at the bottom of the distribution naturally experience faster wage growth, consistent with mean reversion. However, this pattern is significantly more pronounced for the 2000 cohort, which was subject to increased tax transparency. The effects are concentrated at the bottom of the wage distribution. After accounting for placebo effects using the 1997 cohort, we find that only workers in the lowest three deciles exhibit a positive and statistically significant wage response. For example, workers in the first decile experience a relative wage increase of 3.6% and workers in the third experience a salary increase of 0.78%.

We formally test whether the observed reduction in gender wage gaps can be explained by wage compression resulting from the tax transparency policy. Table 7 reports estimates of changes in the gender wage gap after accounting for workers' relative wage position in the baseline year. Column (1) shows that, without any control for wage rank, female workers experience a 0.78% higher wage growth relative to male workers. Columns (2) and (3)

sequentially add controls for wage deciles interacted with time and with comparison group fixed effects. In both cases, the estimated effect remains virtually unchanged, suggesting that relative wage position alone does not account for the observed gender gap reduction. Columns (4) and (5) test for heterogeneity in the gender gap reduction across the wage distribution by estimating separate effects for workers in the bottom and top halves of the wage distribution. The results reveal no statistically significant differences, indicating that the narrowing of the gender wage gap occurs broadly across the wage distribution rather than being concentrated among low- or high-wage workers.

7.3.3 Heterogeneous Effects:

To further assess whether the observed effects are driven by increased transparency, we examine whether the estimates differ across firms located in municipalities with and without a published tax catalog, as well as across industries with high and low baseline wage premium gaps. These heterogeneity analyses, reported in Table 6, focus on wage effects three years after the implementation of the tax transparency policy. While none of the coefficients are statistically significant, due to the reduced sample size in these subgroup analyses, the point estimates align with the hypothesis that greater transparency drives the observed wage effects. Specifically, the first-row coefficients show that the average relative wage growth for female workers is 0.13% in municipalities with a tax catalog (Column (1)) and 0.34% in industries with a low baseline wage premium gap (Column (2)). The second-row coefficients indicate that the wage effects for women relative to men are 0.52% and 0.53% higher in municipalities without the tax catalog and in industries with a high baseline wage premium gap, respectively. Although not statistically significant, these estimates are close in magnitude to the main effect of 0.56% reported in Section 6.3.1, lending further support to the interpretation that transparency contributed to the reduction in gender wage gaps.

8 Conclusion:

Pay transparency has gained increasing attention as a potential tool for promoting wage equity, with various countries recently adopting different forms of these policies. One important form of wage inequality that these initiatives aim to address is the persistent gender wage gap in the labor market. Despite significant gains in women’s educational attainment and labor force participation, the earnings gap between men and women remains substantial. When implementing pay transparency laws, policymakers often argue that a key reason for the persistence of gender wage disparities is workers’ limited access to information about firms’ pay practices. Yet, despite the growing popularity of these policies, empirical evidence on how transparency affects wage-setting behavior remains limited.

This paper contributes to the literature by examining the effect of pay transparency on gender wage disparities, focusing on an underexplored form of transparency: the public disclosure of individual tax information. Although such policies are primarily intended to enhance tax compliance—particularly among business owners and the self-employed—they may also produce unintended consequences for wage earners. We analyze a national policy shift in Norway, where in 2001, individual tax returns were made publicly accessible online. This abrupt and broad increase in transparency provides a unique setting to assess how enhanced access to pay information affects gender wage gaps among coworkers.

Using linked employer-employee data, we estimate the differential effects between female and male workers using a difference-in-difference approach that focuses solely on within-firm variation in gender wage gaps. Gender gaps remained stable until 2001, after which they declined by 2.17 percentage points—an 8.7% reduction relative to a baseline gap of 24.9%. The decline occurs across all economic sectors, with the strongest effects in the private sector, and across all industries, with a stronger reduction in sectors that had higher initial gaps. To bolster our identification, we conduct two triple-difference analyses showing stronger effects in municipalities without prior access to printed tax lists and in industries with higher baseline gender wage premium gaps. Importantly, this variation reveals that the reduction in gender gaps is driven by rising female wages rather than falling male wages.

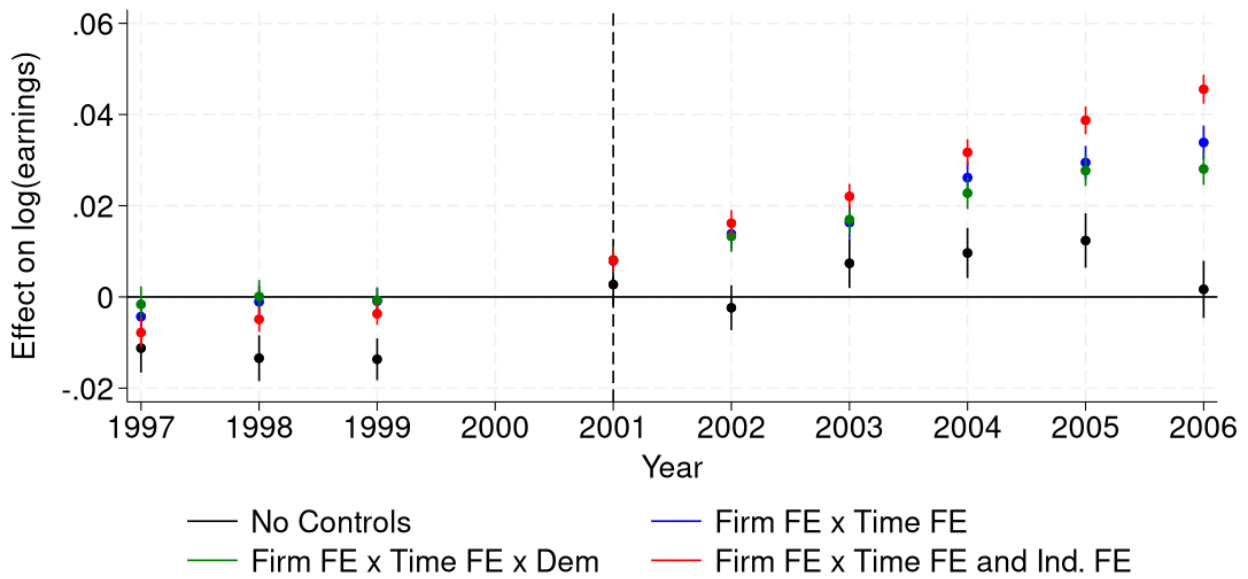
We explore the mechanism behind the gender gap reduction by fixing workers in their 2000 reference groups and tracking them for three years after the policy change, adjusting for typical labor market dynamics. Our findings show that increased transparency led to higher termination rates for female workers, followed by more frequent job-to-job transitions, ultimately resulting in higher wage growth. This wage increase is entirely concentrated among workers who switch firms. We also find that transparency boosted wages for workers at the bottom of the within-firm wage distribution in the pre-period, although this channel does not fully explain the observed reduction in gender wage gaps.

References

- Abowd, John M, Kramarz, Francis, and Margolis, David N (1999). “High wage workers and high wage firms”. *Econometrica* 67.2, pp. 251–333.
- Baker, Michael, Messacar, Derek, Han, Paul, Grogan, Jared, Madrazo, Chester, Tino, Stephen, Thornton, Annabel, and Zhu, Ruizhi (2021). Pay Transparency and the Gender Gap 1. URL: <https://www.bloomberg.com/news/articles/2017-04-13/tech-companies-tout-gender-pay-equity-but-balk-at->.
- Bennedsen, Morten, Simintzi, Elena, Tsoutsoura, Margarita, and Wolfenzon, Daniel (2020). Do Firms Respond to Gender Pay Gap Transparency? *. URL: <http://blob.capitoltrack.com/17blobs/e3526ab2-1360-4461-a1d3-b0580abe6172>.
- Blau, Francine D. and Kahn, Lawrence M. (2017). “The Gender Wage Gap: Extent, Trends, and Explanations”. *Journal of Economic Literature* 55, pp. 789–865.
- Bracha, Anat, Gneezy, Uri, and Loewenstein, George (Apr. 2015). “Relative pay and labor supply”. *Journal of Labor Economics* 33 (2), pp. 297–315.
- Breza, Emily, Kaur, Supreet, and Shamdasani, Yogita (May 2018). “The morale effects of pay inequality”. *Quarterly Journal of Economics* 133 (2), pp. 611–663.
- Böheim, René and Gust, Sarah (2021). The Austrian Pay Transparency Law and the Gender Wage Gap. URL: www.iza.org.
- Bø, Erlend E., Slemrod, Joel, and Thoresen, Thor O. (2015). “Taxes on the internet: Deterrence effects of public disclosure”. *American Economic Journal: Economic Policy* 7 (1), pp. 36–62.
- Card, David, Mas, Alexandre, Moretti, Enrico, and Saez, Emmanuel (2012). “Inequality at Work: The Effect of Peer Salaries on Job Satisfaction”. 102 (6), pp. 2981–3003. DOI: [10.1257/aer.102A2981](https://doi.org/10.1257/aer.102A2981).
- Cullen, Zoe B and Pakzad-Hurson, Bobak (2021). Equilibrium Effects of Pay Transparency. URL: <http://www.nber.org/data-appendix/w28903>.
- Cullen, Zoë and Perez-Truglia, Ricardo (Mar. 2022). “How Much Does Your Boss Make? The Effects of Salary Comparisons”. *Journal of Political Economy* 130 (3), pp. 766–822.
- Dube, Arindrajit, Giuliano, Laura, and Leonard, Jonathan (2019). “American Economic Association Fairness and Frictions”. *Source: The American Economic Review* 109 (2), pp. 620–663. DOI: [10.2307/26603155](https://doi.org/10.2307/26603155).
- Gulyas, Andreas, Seitz, Sebastian, and Sinha, Sourav (Forthcoming). “Does Pay Transparency Affect the Gender Wage Gap? Evidence from Austria”. *American Economic Journal: Economic Policy*.
- Jäger, S., Roth, C., Roussille, N., and Schoefer, B. (2024). “Worker beliefs about outside options.” *Quarterly Journal of Economics*.
- Kristoffer B, Hvidberg, Kreiner, Claus T., and Stantcheva, Stefanie (2023). “Social positions and fairness views on inequality.” *Review of Economic Studies* 90.6, pp. 3083–3118.

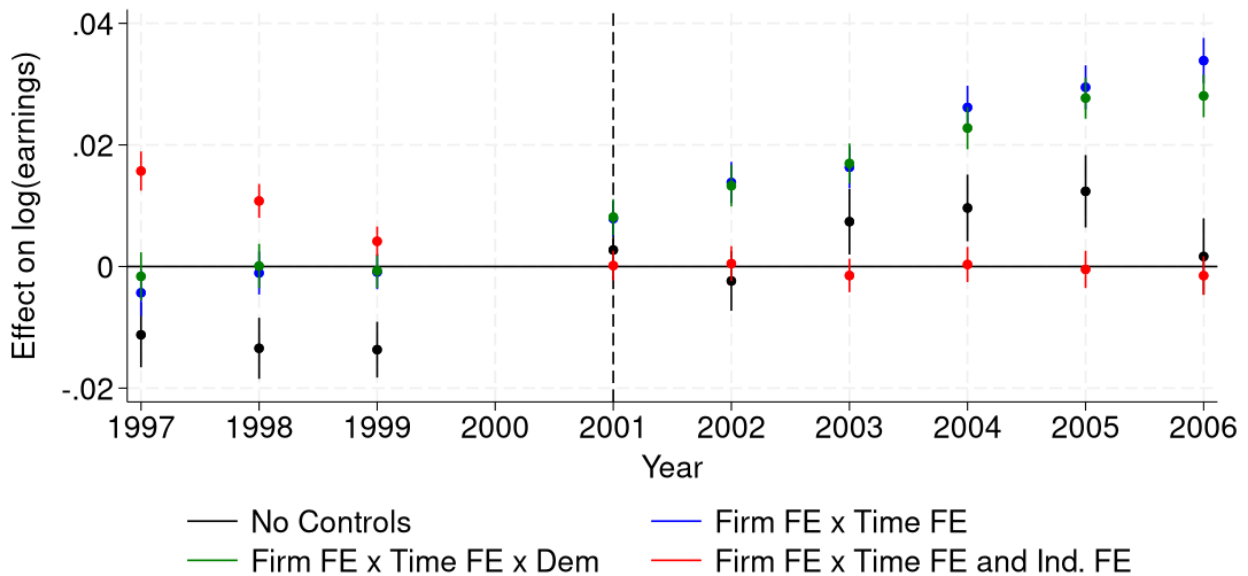
- Lachowska, Marta, Mas, Alexandre, Saggio, Raffaele, and Woodbury, Stephen A (2023). “Do firm effects drift? Evidence from Washington administrative data”. *Journal of Econometrics* 233.2, pp. 375–395.
- Mas, Alexandre (2017). “Does Transparency Lead to Pay Compression?” *Journal of Political Economy* 125 (5).
- Perez-Truglia, Ricardo (Apr. 2020). “The effects of income transparency on well-being: Evidence from a natural experiment†”. *American Economic Review* 110 (4), pp. 1019–1054. ISSN: 19447981. DOI: [10.1257/aer.20160256](https://doi.org/10.1257/aer.20160256).
- Reck, Daniel, Slemrod, Joel, and Vattø, Trine Engh (Aug. 2022). “Public disclosure of tax information: Compliance tool or social network?” *Journal of Public Economics* 212, p. 104708. ISSN: 00472727. DOI: [10.1016/j.jpubeco.2022.104708](https://doi.org/10.1016/j.jpubeco.2022.104708).
- Rege, Mari and Solli, Ingeborg F. (2015). “Lagging Behind the Joneses: The Impact of Relative Earnings on Job Separation”. *Unpublished*.

Figure 4: Difference-in-Differences Estimates of the Effect of Tax Transparency on Female Wages Relative to Male Wages



Note: The results are clustered by firms, and all regressions include year fixed effects. Confidence intervals are generated using a 5% significance level. The vertical line marks the year when tax returns were first posted online. We plot the interactions between the "female" variable and each year dummy, varying the set of fixed effects as described in the legend. The red coefficients are estimated using the detrended wage variable, which is generated by estimating the model in the pre-period, substituting the interaction of female and year dummies with a specific time trend for females, and then subtracting the predicted trend from the original wage variable.

Figure 5: Difference-in-Differences Estimates of the Effect of Tax Transparency on Female Wages Relative to Male Wages



Note: The results are clustered by firms, and all regressions include year fixed effects. Confidence intervals are generated using a 5% significance level. The vertical line marks the year when tax returns were first posted online. We plot the interactions between the "female" variable and each year dummy, varying the set of fixed effects as described in the legend.

Table 1: Effect of Pay Transparency on the Gender Gap:

Dep. Var: Log(Wages)	Group specification			
	(1)	(2)	(3)	(4)
After x Female	0.0146*** (0.0021)	0.0217*** (0.0015)	0.0191*** (0.0014)	0.0268*** (0.0011)
Observations	7,726,147	7,632,529	7,265,699	7,389,843
Reference Group	Everyone	Coworkers	Coworkers with Same Education, and Tenure	Coworkers
Reference Group - Time FEs	NO	YES	YES	YES
Reference Group - Female FEs	NO	YES	YES	YES
Individual FE	NO	NO	NO	YES
Adjusted for Linear Pre-trend	NO	NO	NO	YES
Conditional Gender Gap - Baseline	-0.275	-0.249	-0.223	-0.249

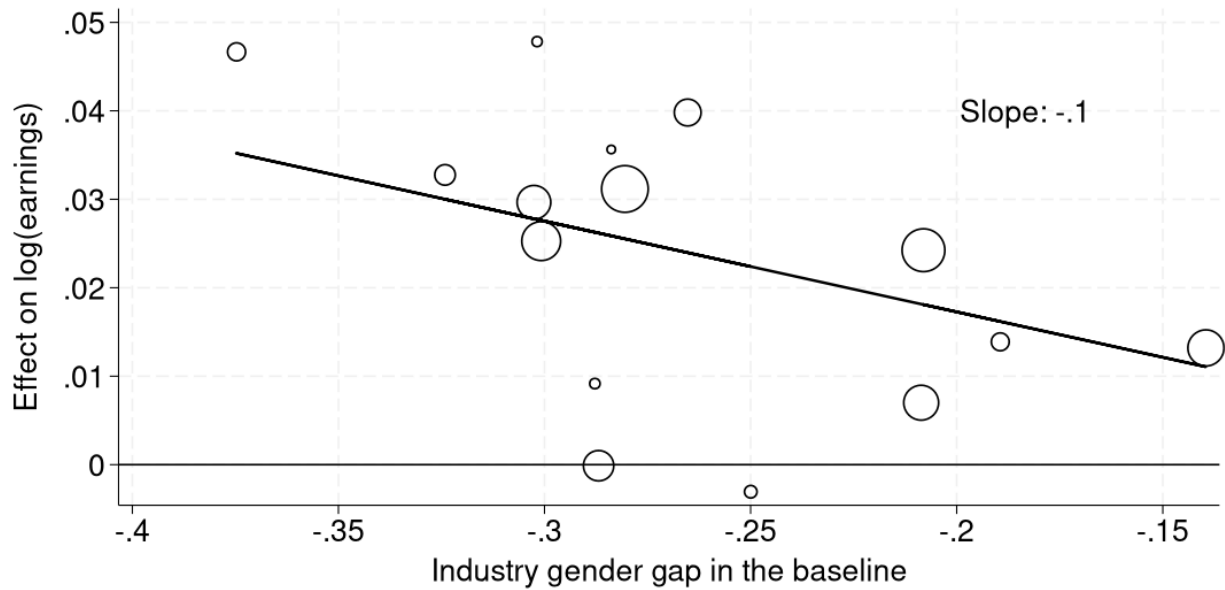
Note: Significance levels: * 10%, ** 5%, ***1%. The results are clustered by the firms. The conditional gender gaps are estimated by re-estimating the same model, replacing the reference group with fixed effects interacted with the female indicator. We then report the coefficient on the female dummy, which represents the conditional gender wage gap in the pre-period.

Table 2: Effect of Pay Transparency on the Gender Gap by Sector Baseline Gender Gap:

Dep. Var: Log(Wages)	Different Sector			
	All Sectors	Private	Household	Public
	(1)	(2)	(3)	(4)
After x Female	0.0217*** (0.0015)	0.0265*** (0.0021)	0.0118** (0.0058)	0.0150*** (0.0021)
Observations	7,632,529	4,785,178	154,803	2,691,756
Conditional Gender Gap - Baseline	-0.249	-0.293	-0.190	-0.182

Note: Significance levels: * 10%, ** 5%, ***1%. The results are clustered by the firms. The results are estimated in separate regression for each economics sector and all models include firm interacted with time fixed effects and firm interacted with gender fixed effects. The conditional gender gaps are estimated by re-estimating the same model, replacing the reference group with fixed effects interacted with the female indicator. We then report the coefficient on the female dummy, which represents the conditional gender wage gap in the pre-period.

Figure 6: Effect of Pay Transparency on the Gender Gap by Industry Baseline Gender Gap



Note: The figure displays the estimated relative wage increase for females across 15 industry classifications, plotted against each industry's baseline gender wage gap. Each point represents an industry, and the scatter plot is weighted by the number of observations in that industry. The regression line is estimated using a weighted regression of the female relative wage increase on the baseline gender gap, with weights equal to the number of observations per industry.

Table 3: Effect of Pay Transparency on the Gender Gap by Industry Baseline Gender Gap:

Dep. Var: Log(Wages)	Different Industries							
	Agriculture and Fishing	Mining and Quarrying	Manufacturing	Utilities	Construction	Wholesale and Retail	Hotels	Restaurants
After x Female	0.0357* (0.0210)	0.0467** (0.0187)	0.0312*** (0.0030)	-0.0030 (0.0082)	0.0398*** (0.0082)	0.0253*** (0.0030)	0.0479*** (0.0156)	0.0092 (0.0157)
Observations	39,048	197,622	1,372,204	93,096	446,317	920,619	58,626	57,895
Conditional Gender Gap - Baseline	-0.283	-0.374	-0.280	-0.250	-0.265	-0.300	-0.301	-0.287
Dep. Var: Log(Wages)	Different Industries							
	Transport, Storage and Communication	Financial Intermediation	Real Estate	Public Adm	Education	Health Social Work	Community, Social and Personal Activities	
After x Female	-0.0001 (0.0112)	0.0328*** (0.0061)	0.0297*** (0.0040)	0.0070*** (0.0023)	0.0132*** (0.0024)	0.0243*** (0.0049)	0.0139*** (0.0052)	
Observations	574,223	259,754	696,323	759,675	818,559	1,149,162	188,815	
Conditional Gender Gap - Baseline	-0.286	-0.324	-0.302	-0.208	-0.139	-0.208	-0.189	

Note: Significance levels: * 10%, ** 5%, ***1%. The results are clustered by the firms. The results are estimated in separate regression for each industry and all models include firm interacted with time fixed effects and firm interacted with gender fixed effects. The conditional gender gaps are estimated by re-estimating the same model, replacing the reference group with fixed effects interacted with the female indicator. We then report the coefficient on the female dummy, which represents the conditional gender wage gap in the pre-period.

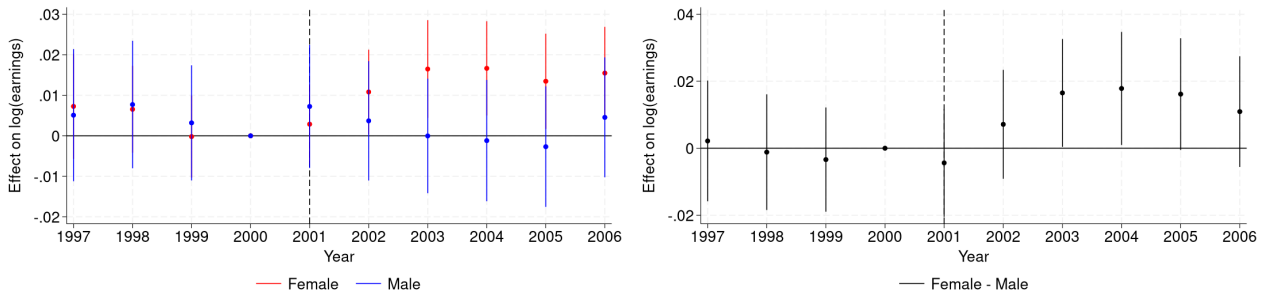
Table 4: Effect of Pay Transparency on the Gender Gap by Different Heterogeneities

Dep. Var. Log(Wage)	Different Heterogeneity Dummies			
	No Tax Catalog		High Gap Industry	
	(1)	(2)	(3)	(4)
After x Female x Het.	0.0107** (0.0046)	0.0120** (0.0049)	0.0082*** (0.0028)	0.0181*** (0.0029)
After x Female	0.0105*** (0.0038)	0.0091** (0.0041)	0.0107*** (0.0020)	0.0127*** (0.0023)
After x Het.	-0.0016 (0.0038)		-0.0229*** (0.0025)	
Observations	1,573,139	1,554,567	7,710,852	7,632,529
Firm - Time FEs	NO	YES	NO	YES
Firm - Female FEs	YES	YES	YES	YES


Note: Significance levels: * 10%, ** 5%, ***1%. Standard errors are clustered at the firm level. Models in columns (1) and (3) include education-by-year-by a dummy for tenure less than one year, a continuous tenure variable interacted with year, and firm-by-gender fixed effects. Models in columns (2) and (4) include firm-by-year fixed effects and firm-by-gender fixed effects. *Het* is a dummy variable that equals 1 for municipalities without printed versions of the tax catalog prior to the legislation in the columns labeled “No Tax Catalog”, and equals 1 for peer groups in industries with a wage premium gap above the median prior to the legislation in the columns labeled “High Gap Industry”.

Figure 7: Effect of Pay Transparency on the Gender Gap by Different Heterogeneities - Difference in Difference Graphs - Municipalities With and Without Tax Catalogs

(a) DiD Effects for Each Gender - Without Including Firm x Time FE (b) Third Difference Effects - Without Including Firm x Time FE:



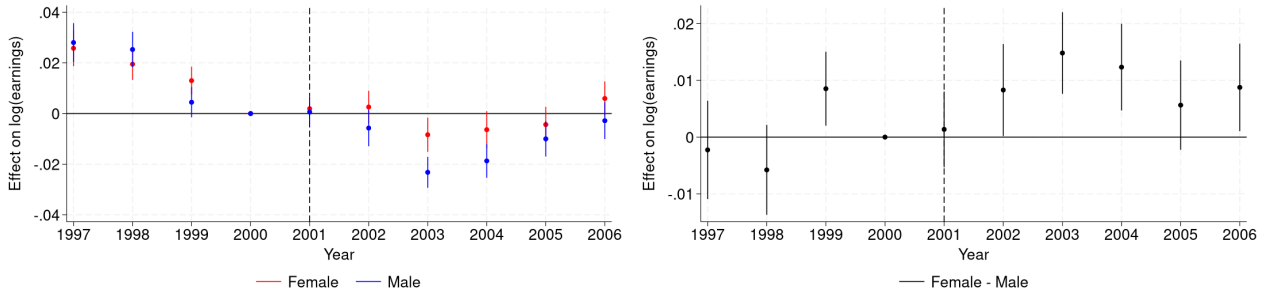
(c) Third Difference Effects - With Including Firm x Time FE:




Note: Standard errors are clustered at the firm level, and all regressions include year fixed effects. Confidence intervals are constructed at the 5% significance level. The vertical line indicates the year when tax returns were first published online. Models in panels (a) and (b) include education-by-year interactions, a dummy for tenure less than one year, a continuous tenure variable interacted with year, and firm-by-gender fixed effects. The model in panel (c) includes firm-by-year and firm-by-gender fixed effects. In panel (a), we estimate a triple-difference model and report coefficients equivalent to difference-in-difference estimates for males and females, using municipalities without access to the tax catalog as the treated group. Panels (b) and (c) report the triple-difference coefficients directly.

Figure 8: Effect of Pay Transparency on the Gender Gap by Different Heterogeneities - Difference in Difference Graphs - Industries With High and Low Gender Gap

(a) DiD Effects for Each Gender - Without Including Firm x Time FE: (b) Third Difference Effects - Without Including Firm x Time FE:



(c) Third Difference Effects - With Including Firm x Time FE:




Note: Standard errors are clustered at the firm level, and all regressions include year fixed effects. Confidence intervals are constructed at the 5% significance level. The vertical line indicates the year when tax returns were first published online. Models in panels (a) and (b) include education-by-year interactions, a dummy for tenure less than one year, a continuous tenure variable interacted with year, and firm-by-gender fixed effects. The model in panel (c) includes firm-by-year and firm-by-gender fixed effects. In panel (a), we estimate a triple-difference model and report coefficients equivalent to difference-in-difference estimates for males and females, using with high wage premium gap as the treated group. Panels (b) and (c) report the triple-difference coefficients directly.

Figure 9: Effect of Pay Transparency on Job Terminations by Gender

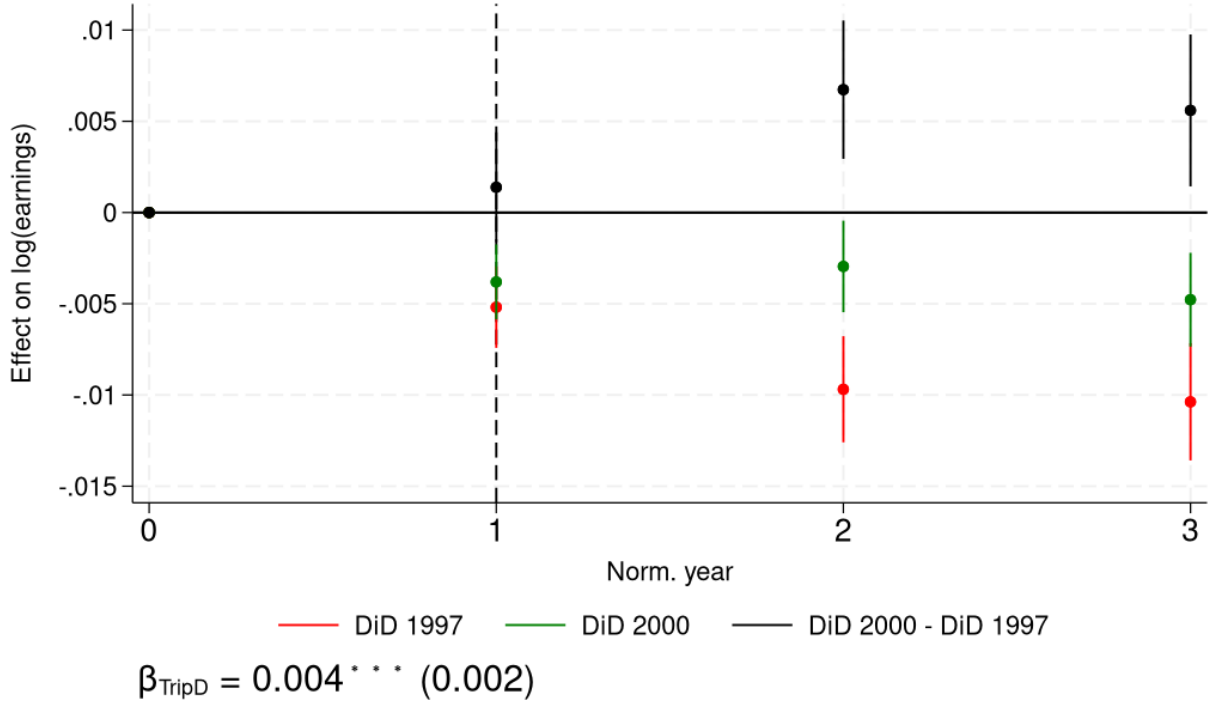
(a) Any Termination:

(b) Switch to Another Job:



Note: Standard errors are clustered at the baseline firm level, and all regressions include year fixed effects. Confidence intervals are constructed at the 5% significance level. The models are estimated using cross-sectional variation, with the dependent variable equal to one if there was a termination or job switch during the period indicated on the x-axis. It includes firm interacted with cohort fixed effects. Any Period is a dummy variable equal to 1 if the worker is not employed at the baseline firm in any period after the baseline year. All Period is a dummy equal to 1 if the worker is not employed at the baseline firm in all three years following the baseline year. The 1997 and 2000 cohorts correspond to samples where the baseline year is set to 1997 and 2000, respectively. The coefficients are reported separately for each cohort. The coefficients displayed in black correspond to the interaction between the female indicator and the dummy for the 2000 cohort.

Figure 10: Effect of Pay Transparency on Gender Gaps Over the Years



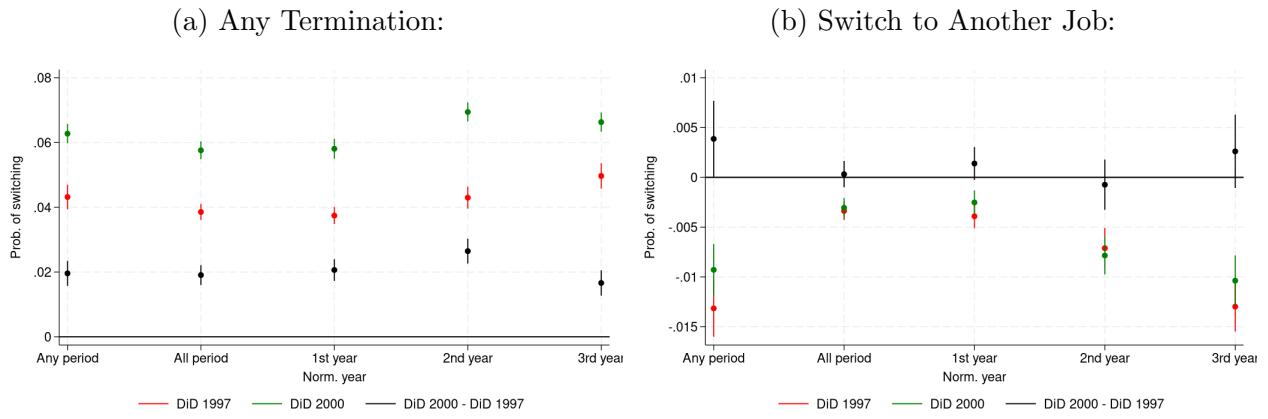
Note: Standard errors are clustered at the baseline firm level, and all regressions include year fixed effects. Confidence intervals are constructed at the 5% significance level. The model is estimated using individual fixed effects and firm-by-year-by-cohort fixed effects. The 1997 and 2000 cohorts correspond to samples where the baseline year is set to 1997 and 2000, respectively. The coefficients are reported separately for each cohort. The coefficients displayed in black correspond to the interaction between the female indicator, a dummy for the 2000 cohort, the a dummy for each normalized year.

Table 5: Effect of Pay Transparency on Wage for Switchers and Stayers:

Dep. Var. Log (Wage)	Gender Effects			
	All Years		3 Years After	
	(1)	(2)	(3)	(4)
After x Female x $1\{c = 2000\}$	0.0044*** (0.0016)	-0.0015 (0.0017)	0.0058*** (0.0021)	0.0036 (0.0024)
After x Female x $1\{c = 2000\}$ x Switch		0.0119*** (0.0024)		0.0060* (0.0036)
After x $1\{c = 2000\}$ x Switch		-0.0062*** (0.0017)		0.0020 (0.0026)
Observations	4,334,351	4,334,351	1,870,040	1,870,040


Note: Significance levels: * 10%, ** 5%, ***1%. Standard errors are clustered at the firm level (baseline). The model includes individual fixed effects and firm-by-year-by-cohort fixed effects. Switch is a dummy variable equal to one for workers who moved from their baseline firm within any of the three years following the baseline period. Workers who leave their baseline firm after the baseline year and never return do not contribute to the identification of any coefficients, and are virtually dropped out of the sample.

Figure 11: Effect of Pay Transparency on Job Terminations by Relative Position



Note: Standard errors are clustered at the baseline firm level, and all regressions include year fixed effects. Confidence intervals are constructed at the 5% significance level. The models are estimated using cross-sectional variation, with the dependent variable equal to one if there was a termination or job switch during the period indicated on the x-axis. It includes firm interacted with cohort fixed effects. The reported coefficients correspond to the interaction between the an indicator for workers below the median in their baseline firm and a dummy for the 2000 cohort.

Figure 12: Effect of Pay Transparency on Wages by Relative Position



Note: Standard errors are clustered at the baseline firm level, and all regression include year fixed effects. Confidence intervals are constructed at the 5% significance level. The model is estimated using individual fixed effects and firm-by-year-by-cohort fixed effects. The reported coefficients correspond to the interaction between the an indicator for workers below the median in their baseline firm, a dummy for the 2000 cohort, the a dummy for each normalized year.

Table 6: Heterogeneous Effects of Pay Transparency on Gender Gaps:

Dep. Var: Log(Wage)	Gender Effects	
	Catalog (1)	High Gap Industry (2)
After x Female x $1\{c=2000\}$	0.0013 (0.0084)	0.0034 (0.0024)
After x Female x $1\{c=2000\} \cdot het_i$	0.0052 (0.0096)	0.0053 (0.0040)
Observations	401,924	1,870,040

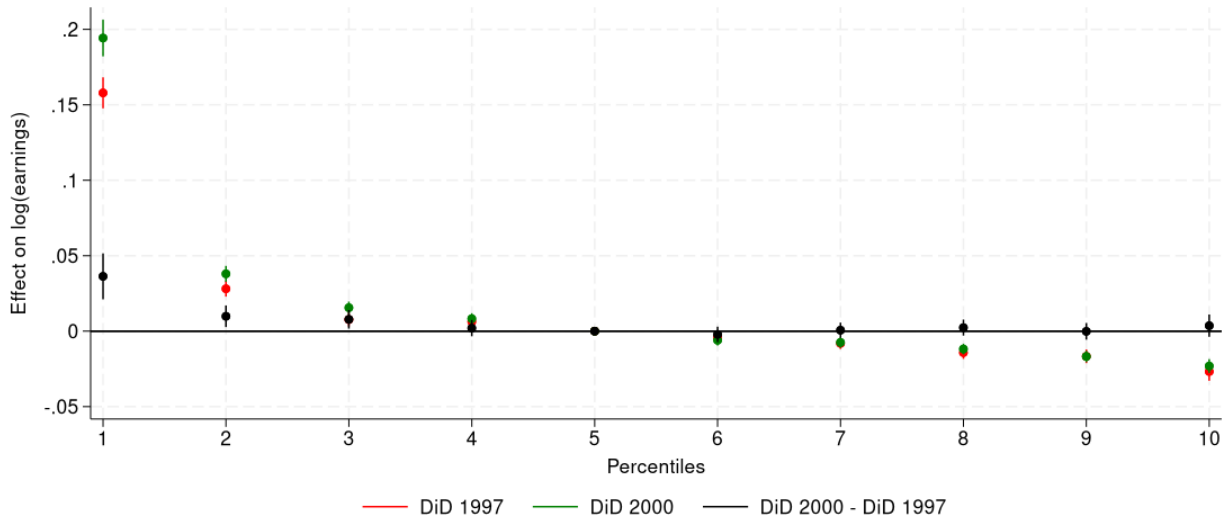
Note: Significance levels: * 10%, ** 5%, ***1%. The model includes individual fixed effects and firm-by-year-by-cohort fixed effects. Het_i is a dummy variable equal to one if the individual was in a city with a printed version of the tax catalog in the baseline period (Column 1), or if the individual was in an industry with a high baseline wage premium gap (Column 2).

Table 7: Effect of Pay Transparency on Gender Gaps After Controlling for Relative Position:

Dep. Var: Log(Wage)	Relative position interacted with gender				
	(1)	(2)	(3)	(4)	(5)
After x Female x $1\{c=2000\}$	0.0078*** (0.0023)	0.0059** (0.0026)	0.0088*** (0.0028)	0.0105*** (0.0026)	0.0070** (0.0031)
After x Female x $1\{c=2000\} \cdot Below_{i,(g,t_0),c}$				-0.0081* (0.0043)	0.0032 (0.0053)
Baseline wage deciles interacted with time and cohort FE	NO	YES	NO	YES	NO
Baseline wage deciles interacted with comparison group, time and cohort FE	NO	NO	YES	NO	YES
Observations	1,455,742	1,455,742	1,220,298	1,455,742	1,220,298

Note: Significance levels: * 10%, ** 5%, ***1%. Standard errors are clustered at the firm level (baseline). The model includes individual fixed effects and firm-by-year-by-cohort fixed effects. Below is a dummy variable equal to one if the individual was below the median wage in their firm during the baseline period.

Figure 13: Difference-in-Differences Coefficients for the Effect of Tax Transparency on Wages Across Within-Firm Wage Deciles



Note: Standard errors are clustered at the baseline firm level, and all regressions include year fixed effects. Confidence intervals are constructed at the 5% significance level. The model includes individual fixed effects and firm-by-year-by-cohort fixed effects. The reported coefficients correspond to the quadruple interaction of (i) an indicator for each decile of the wage distribution within the worker's baseline firm, (ii) a dummy for the 2000 cohort, (iii) a dummy for the third normalized year, and (iv) a female indicator.

Appendix - Supplementary tables and figures:

Appendix - Two-way Fixed Effects Model:

We estimate firm-specific pay premiums for men and women separately for each year in our sample (1997–2006). In Equation 4, we implement a time-varying version of the additive two-way fixed effects model for worker i and firm j , inspired by Abowd, Kramarz, and Margolis (1999) and extended by Lachowska et al. (2023), stratifying the estimation by gender $G(i)$. The sample is restricted to firms that employed at least five workers and had at least one male and one female employee in any given period.

$$\log(w_{it}) = \alpha_i + \Psi_{J(i,t),t}^{G(i)} + X'_{it}\beta^{G(i)} + r_{i,t} \quad (4)$$

$\log(w_{it})$ represents the logarithm of the earnings variable for worker i in year t , α_i captures individual fixed effects, $\Psi_{J(i,t),t}^{G(i)}$ incorporates gender-specific firm effects by year, X'_{it} accounts for gender-specific returns to covariates, and includes year fixed-effects for each education level, and $r_{i,t}$ is the error term.

We estimate the model separately for males and females, using the largest set of firms connected through worker transitions for each group independently (connected set). In this specification, only differences in firm fixed effects are identified, as the model must be es-

timated by holding one firm constant as the reference in the baseline period. A similar normalization is required for interpreting firm-specific wage premiums and firm-specific gender wage premiums.

To avoid spurious comparisons when estimating gender wage premiums, we restrict the sample to firms that are double-connected through both male and female job switchers. We then compute each firm's wage premium gap as the difference between the firm fixed effects estimated separately for men and women in the year prior to the tax transparency shock (2000). We further calculate the average within-firm gender gap premium for each industry by averaging the firm-premium gap weighted by the total number of workers in each firm, and classify the industries with premium gaps below and above the median.