Statistisk sentralbyrå
Statistics Norway

Does Tax Transparency Influence Wage Setting? Evidence on Gender Gaps

Cristiano C. Carvalho and Trine E. Vattø

TALL SOM FORTELLER

DISCUSSION PAPERS

1033

Discussion Papers: comprise research papers intended for international journals or books. A preprint of a Discussion Paper may be longer and more elaborate than a standard journal article, as it may include intermediate calculations and background material etc.

The Discussion Papers series presents results from ongoing research projects and other research and analysis by SSB staff. The views and conclusions in this document are those of the authors.

Published: February 2026

Abstracts with downloadable Discussion Papers
in PDF are available on the Internet:

https://www.ssb.no/discussion-papers
http://ideas.repec.org/s/ssb/dispap.html

ISSN 1892-753X (electronic)

Abstract

This paper examines how an often-overlooked source of pay transparency—the public disclosure of tax information—affects gender wage gaps. We exploit a 2001 change in Norway that made individual tax returns searchable online. Using matched employer–employee data and a difference-in-differences design, we find that within-firm gender wage gaps fell by 2.2 percentage points (8.7 percent), driven by rising female wages. Effects are strongest in private-sector firms, industries with initially larger gaps, and municipalities that previously lacked easy access to printed tax lists. Wage gains are concentrated among job-changing women, suggesting that broad-based transparency mainly operates through improved information for job search.

Keywords: Gender Wage Gap, Income Transparency, Public Disclosure of Tax Information

JEL classification: J16, J31, J38

Acknowledgements: We are especially grateful to Joel Slemrod, Mel Stephens, Jim Hines, and Ben Scuderi for their guidance and support throughout this project. We also thank Thor Olav Thoresen, Terje Skjerpen, John Bound, Charlie Brown, Ashley Craig, Amanda Kowalski, Mike Mueller-Smith, Ana Reynoso, Damian Vergara, and Basit Zafar for helpful comments and discussions.

Address: Cristiano C. Carvalho, Puc-Rio. Email: cristianocarvalho@econ.puc-rio.br, Trine Engh Vattø, Statistics Norway, Research Department. E-mail: trine.vatto@ssb.no

Sammendrag

Norge har lange tradisjoner for offentlig innsyn i skattelister, og slike ordninger har gjerne vært begrunnet med hensynet til etterlevelse og tillit i skattesystemet. Samtidig kan offentlig innsyn i skattepliktig inntekt endre informasjonsgrunnlaget arbeidstakere har om kollegers inntekter og egne alternative jobbmuligheter, og dermed også påvirke lønnsfastsettelse og fordeling i arbeidsmarkedet. Fram til begynnelsen av 2000-tallet var innsynet i praksis begrenset ved at man måtte møte opp ved ligningskontor eller rådhus i en avgrenset periode. Dette endret seg brått høsten 2001, da store mediehus digitaliserte skattelistene og gjorde dem fritt søkbare på internett, slik at opplysninger om skattepliktig inntekt, formue og utliknet skatt ble langt enklere tilgjengelige.

Ved å koble detaljerte arbeidstaker- og arbeidsgiverdata med individers skatteopplysninger for perioden 1997–2006, studerer vi hvordan lønnsgapet mellom kvinner og menn utviklet seg innad i virksomhetene før og etter at skattelistene ble digitalt tilgjengelige og søkbare. Selv om skattepliktig inntekt ikke er et nøyaktig mål på arbeidsinntekten, viser vi at den kan fungere som en proxy for arbeidsinntekt i relative sammenlikninger. Vi benytter en forskjell-i-forskjell-metode og viser at lønnsgapet mellom kvinner og menn innad i virksomheter var stabilt fram til 2001, og falt deretter med om lag 2,2 prosentpoeng, tilsvarende 8,7 prosent av lønnsgapet i utgangspunktet. Reduksjonen var størst i privat sektor, i næringer som hadde store kjønnsforskjeller i lønn, samt i kommuner med begrenset tilgang til trykte skattelister før 2001. Vi finner at reduksjonen i lønnsgapet skyldtes at kvinners lønninger økte, snarere enn at menns lønninger sank.

Videre undersøker vi mekanismene bak reduksjonen i lønnsgapet mellom kvinner og menn ved å følge de samme personene over tid. Vi bruker en tidligere periode som sammenlikningsgrunnlag (en placebo-kohort) for å justere for underliggende kjønnsforskjeller i jobbmobilitet og lønnsvekst. Etter 2001 økte kvinnelige arbeidstakeres mobilitet, og lønnsgevinstene var i hovedsak knyttet til jobb-til-jobb-skifter, mens effektene var mer beskjedne blant dem som ble værende. Dette peker mot at inntektsinformasjon særlig virket gjennom jobbsøk og jobbskifter, snarere enn lønnsjusteringer i samme virksomhet. Videre viser analysene at reduksjonen i kjønnsforskjeller ikke kan tilskrives generell lønnskompresjon alene: når det kontrolleres for arbeidstakernes plassering i virksomhetens lønnsfordeling, finner vi tilsvarende reduksjoner i lønnsgapet.

1 Introduction

The gender wage gap remains one of the most persistent features of labor markets across countries. Despite substantial gains in women's educational attainment and labor force participation, women still earn less than men on average (Blau and Kahn, 2017). A common explanation is that workers have limited information about firms' pay practices. This lack of transparency shapes expectations about future wages and influences decisions such as how much effort to exert, whether to seek promotion, or when to negotiate pay (Cullen, 2024). For women, access to credible pay information may be especially important because it can correct their lower wage expectations and thereby influence outcomes such as gender pay discrimination (Reuben, Wiswall, and Zafar, 2017; Kiessling et al., 2024; Roussille, 2024).

Pay transparency has gained increasing attention as a policy instrument for promoting fair wages and reducing gender wage gaps. In recent years, a number of countries have implemented measures to improve wage transparency, ranging from modest interventions—such as prohibiting employers from restricting salary discussions—to more comprehensive requirements mandating the disclosure of wage information.[1] These developments raise the question of whether greater pay visibility can help uncover and reduce unjustified gender wage disparities.

Although these policies have gained momentum, empirical evidence on the impact of transparency on wage-setting behavior remains limited. This article contributes to the growing literature that examines the effect of pay transparency on gender wage disparities. We focus on an underexplored source of transparency: the public disclosure of 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. Public disclosure of individuals' taxable income could, however, 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 sharp and abrupt increase in the transparency of taxable income in Norway in 2001. Historically, Norwegians' tax records have been

[1] According to Cullen and Pakzad-Hurson (2023), 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.

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 by national newspapers and made available online, 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 taxable income, net wealth and total taxes paid. We leverage this sudden shift in public disclosure as an income transparency shock to identify the effects on wages among coworkers.[2]

Our analysis draws on comprehensive Norwegian administrative data that link workers to their employers and to detailed income records. The core dataset is the matched employer–employee register from 1997 to 2006, which covers the full population of workers and firms and includes information on wages, job characteristics, and demographics. We merge these data with the universe of tax returns—which report taxable income, wealth, taxes paid, and other detailed income components. We also incorporate municipality-level information on the pre-existing availability of printed tax lists.

Although the transparency shock reflected taxable income rather than wages, we show that, for full-time workers, taxable income closely tracks annual earnings over the income distribution, despite individual differences in non-wage income and deductions. These patterns suggest that individuals can plausibly use taxable income comparisons to proxy for wage differences. Because online searches required knowing individuals' names, the transparency shock primarily revealed information about people within workers' social and employment networks, such as coworkers, making the within-firm gender gap the most relevant margin of comparison. Consistent with this interpretation, we document that the rise in wage variance in Norway reversed sharply in the year of the transparency shock (2001), driven almost entirely by a reduction in within-firm variance.

We estimate the effects of increased transparency on the differential outcomes for males and females 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 interaction pairs with time and 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

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

percentage points 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 use our rich data and the nationwide disclosure of tax returns to examine how the reduction in gender wage gaps varies across sectors and industries. We find substantial heterogeneity along both dimensions. Transparency reduced gender gaps in both the private and public sectors, with markedly stronger effects in the private sector (2.6 percentage points decline compared with 1.5 percentage points in the public sector). Across industries, we find significant reductions in gender gaps in 13 of 15 aggregated classifications, and the magnitude of the reduction is positively correlated with the baseline gap: a 10% larger initial gap corresponds to roughly a 1-percentage-point greater decline. Using a two-way fixed effects decomposition following Abowd, Kramarz, and Margolis (1999), we further estimate within-firm wage premium gaps to compare 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 1.6% relative to those in industries with lower baseline disparities.

As a validation, we compare municipalities with and without access to printed versions of the tax lists prior to 2001, as used in Bø, Slemrod, and Thoresen (2015) to identify different degrees of transparency shock intensity. Our findings indicate that municipalities lacking prior access experienced a stronger transparency shock, with female workers in these areas receiving an additional 1.2 percentage points relative wage increase compared to those in municipalities where the tax lists were already available. We also show that pre-period trends are parallel for each gender across the two groups of municipalities, providing support for the identifying assumption. Importantly, this variation in access to printed tax catalogs allows us to estimate the separate effects of tax transparency for male and female workers and demonstrates 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 next three years. To account for gender-specific career dynamics unrelated to the transparency shock, we conduct a placebo analysis using the 1997 cohort to net out prevailing labor market trends when estimating effects for the 2000 cohort. This approach isolates within-group variation over time. Relative to men, women experience a 1.04-percentage-point increase in job separations, but only part of this translates into job-to-job transitions, which rise by 0.68 percentage points—about two thirds of the total separation effect. Consistent with intensified search and improved matching, women's wages rise by 0.68% relative to men three years after the baseline. To address selection concerns, we use a broader measure of gross income to show no gender-differential change in the likelihood of receiving any income; using this measure, we find a 1.2% narrowing of the gender gap in gross income.

We further assess whether the observed wage increases stem from improved offers within firms or from workers transitioning to better matches elsewhere, and the evidence favors the latter as the dominant mechanism. When restricting the sample to individuals who remained with their baseline employer for all three years after the transparency shock, female workers exhibit a modest 0.42% relative wage growth. By contrast, the effects are significantly stronger among workers who changed firms during the same period. For these movers, female workers experience an additional 0.66% wage growth relative to males, with the results benchmarked against placebo estimates from the period without the transparency shock. This contrast underscores that job-to-job transitions, rather than within-firm adjustments such as direct wage bargaining, were the primary driver of the reduction in the gender wage gap.

An important question is whether the observed reduction in the gender wage gap following increased transparency is merely a consequence of wage compression with 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 in the baseline year 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 interaction pairs 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.

The rest of the paper is organized as follows. Section 2 reviews the related literature. Section 3 describes the institutional background of Norway's public disclosure of tax information and the 2001 transparency shock. Section 4 outlines the administrative data sources, including the matched employer–employee records, tax registers, and municipality-level information on prior access to printed tax lists. Section 5 presents descriptive evidence, including the relationship between taxable income and wages, and trends in wage dispersion and gender gaps. Section 6 outlines the empirical strategy, detailing the difference-in-differences framework using repeated cross-sectional variation and the subsequent panel-based approach. Section 7 presents the main results, documents heterogeneity across sectors and industries, validates the identification using variation in prior tax-list availability, and examines mechanisms through individual mobility and wage increases. 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 income 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. In line with existing evidence, we find that transparency promotes greater wage compression and reduces within-firm gender gaps. Our contribution extends prior work in three key dimensions. First, much of the existing research has examined narrower forms of transparency, such as the disclosure of aggregate wage statistics by gender (Böheim and Gust, 2021; Bennedsen et al., 2022; Gulyas, Seitz, and Sinha, 2023; Blundell et al., 2025). Second, studies that analyze individual-level disclosures typically restrict their scope to specific sectors, such as universities (Baker et al., 2023; Card et al., 2012) or public employment (Mas, 2017). Third, because of this limited scope, the existing literature on individual-level disclosures does not follow workers once they leave their jobs, and thus provides little evidence on realized search behavior or broader labor market dynamics. 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. We also track workers after they are exposed to the transparency shock and show that job-to-job transitions induced by the shock appear to be the primary mechanism behind the reduction in the gender gap.

Closely related to our study is Baker et al. (2023), 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% decline relative to the baseline conditional gap.[3] 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, 2023; Baker et al., 2023; Bennedsen et al., 2022).

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 tax transparency 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; Hvidberg, 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. Consistent with this mechanism, we show that the information shock reduces the gender wage gap across the entire wage distribution and operates especially through job transitions. This pattern is consistent with research linking current gender gaps to women's lower earnings expectations (Reuben, Wiswall, and Zafar, 2017; Kiessling et al., 2024) and with evidence that information can correct those expectations (Roussille, 2024).

Fourth, our results demonstrate that the public disclosure of tax and income information can meaningfully affect labor market inequality. Several countries 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 reduced tax evasion among

[3] By controlling for firm-by-time fixed effects, we remove differences across firms in each period, so the conditional gender wage gap is computed using only within-firm gender gaps.

business owners by 2.7%. 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, related work also examines Norway's tax disclosure policy. Perez-Truglia (2020) shows that putting tax returns online increased the happiness gap between richer and poorer individuals by 29% and the life-satisfaction gap by 21%. Reck, Slemrod, and Vattø (2022) use individual data on searches following the 2014 reform that removed searcher anonymity. They find that social comparison, rather than tax-compliance concerns, drives most searches; notably, even without anonymity, 15% of queries target people within the same employment network, and there is no evidence that being searched induces greater tax compliance through increased reported income.[4]

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 transparency shock led to higher job separation rates and faster wage growth among lower-earning workers relative to higher-earning workers. 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

Norway has a long tradition of public tax transparency, dating back to the 19th century. However, until the early 2000s, practical access to individual tax records was limited. Individuals who wished to view someone else's tax information had to visit a local tax office or city hall in person, and only within a three-week window following the annual tax assessment. This changed dramatically in the fall of 2001, when a national newspaper digitized the tax lists and made them searchable online. Other major newspapers quickly followed, enabling anyone with internet access to view income information for the entire population.[5]

[4] Following the 2014 reform, searches had to be made through the official government website, and taxpayers could see the names of the individuals who viewed their tax returns.

[5] In later years, access to the tax lists became increasingly regulated. From 2011, users were required to log in with a personal identification number, and in 2014 anonymous searches were abolished. This allowed individuals to see who had viewed their tax records, leading to a substantial decline in the number of searches.

The online tax lists allowed users to search for individuals—typically by surname—with optional filters such as birth year, municipality, or postal code. Upon conducting a search, users could retrieve their full name and these basic identifying details along with annual taxable income, net wealth, and total taxes paid from the previous calendar year, as recorded in the annual tax assessment. These records were made publicly available each fall, following the completion of the tax assessment process.

Making tax returns available online significantly increased their accessibility for a large share of the Norwegian population.[6] Although systematic usage data from the earliest years are limited, anecdotal evidence cited in Perez-Truglia (2020) suggests that the online tax lists quickly became a popular tool for social comparison, described as a "national snooping" phenomenon. For example, one newspaper reported nearly 30 million searches in 2007, and official figures from 2013 indicated over 17 million searches conducted by more than 920,000 unique users.[7] A 2007 Synovate survey 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, with 26% reporting that they had looked up 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

Our primary data source is the Norwegian matched employer-employee dataset covering the period from 1997 to 2006. This dataset includes the full population of Norwegian workers and the firms they are employed by, allowing us to link each individual to their peers. It contains detailed information on job characteristics—including wages, industry, and contract type (full-time or part-time)—as well as individual characteristics such as gender, education level, municipality of residence, and age.

We restrict our sample to full-time employees aged 25 to 60. To ensure that wages reflect compensation within a specific job, we further limit the sample to individuals employed at a single firm in December of each year. Wages are annualized by dividing total

[6] 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.

[7] To put these numbers into perspective, the adult (18 and over) population in Norway was approximately 4 million during the mid-2000s.

yearly earnings by the number of contract days and multiplying by 365.[8] The resulting annualized wage measure is used to determine each worker's relative position within their reference group. Our baseline approach is to define the reference group as all coworkers within the same establishment of a firm.[9] However, in one specification, we use a more granular specification that includes education level and tenure. To ensure identification of gender gap effects and consistency in firm-level comparisons in our baseline estimates, we restrict the sample to firms that (i) employed both men and women at some point during the sample period—so that gender comparisons can be made within firms—and (ii) did not change industry, sector (public/private), or municipality, in order to prevent firms from switching groups in ways that could otherwise confound our heterogeneity analyses.

To check how well the disclosed tax information serves as a proxy for wages, we merge the sample of full-time employees with administrative registers of tax return records using unique individual identifiers. The tax records provide disclosed values for the individuals' taxable income,[10] net wealth, and total taxes paid, as well as additional data on non-labor income and gross income. We also combine the administrative data with survey information on which municipalities sold printed tax return transcripts before the tax data were made available online in 2001. This survey was originally conducted and used in Bø, Slemrod, and Thoresen (2015). The data were collected through interviews with chief officers at municipal tax offices and cover 138 municipalities—31 with printed tax catalogs and 107 without. The survey data are merged at the municipality level.[11]

For the analysis where we investigate the mechanisms behind the narrowing gender gaps, we impose additional sample restrictions. Using the longitudinal structure of the data, we construct two cohorts: (i) a transparency cohort—workers employed in year 2000 (the year before tax returns were accessible online)—whom we follow for the subsequent three years; and (ii) a placebo cohort—workers employed in 1997—whom we track through 2000. In both cohorts, we restrict the sample to prime-age, full-time employees (ages 25–55 at baseline) with at least one year of tenure at their baseline employer. These filters avoid mechanical attrition, as the main sample includes workers up to age 60, and ensure comparable labor-market attachment across cohorts. In some specifications, we relax the full-time restriction and use the annual tax return records to examine how the transparency shock affected the probability of earning different types of non-wage income and the overall level of gross income.

[8]This single-firm restriction also improves the use of taxable income as a proxy for wages earned in a specific job.

[9]To simplify terminology, throughout the paper we refer to firm establishments simply as firms.

[10]*alminnelig inntekt* in Norwegian tax records.

[11]See Bø, Slemrod, and Thoresen (2015) for more details.

5 Descriptive Statistics

5.1 Tax Information as a Proxy for Wage

The shift in transparency of taxable income can influence wage comparisons among coworkers if the publicly disclosed figures serve as a proxy for actual wage earnings. Taxable income differs from wage income in two key respects: it includes other sources of income—such as capital income and transfers—and it reflects both standard and individual-specific deductions. These sources of discrepancy are illustrated in the bin-scatter plot in Figure 1. To construct the figure, we use the sample of full-time workers and rank individuals into percentiles based on their taxable income in year 2000. For each percentile bin, we then plot the average taxable income, gross income, and wage income to assess whether the ranking order is preserved across these measures.

First, consider the difference between gross income and wage income, which captures the contribution of non-labor income. For most of the income distribution—except in the tails— gross income is almost entirely composed of wage income. On average, wage income accounts for more than 85% of gross income between the 10th and 91st percentiles, indicating that non-wage income plays a limited role for the majority of workers in our sample, as also shown in Figure A1 in the Appendix.

Second, the difference between gross income and taxable income reflects the impact of deductions. These include both standard deductions and deductions that vary across individuals due to differences in eligibility (e.g., interest on debt, commuting expenses, and childcare expenses). Although taxable income is systematically lower than gross income because of these deductions, the figure shows that the ranking across the income distribution is preserved.

Taken together, these patterns suggest that individuals could plausibly have used the observed taxable income ranking as a proxy for wage level ranking. While the ranking is not perfectly aligned at the individual level—due to variation in non-wage income and deductions—it remains sufficiently consistent across the income distribution to support meaningful relative wage comparisons for the majority of workers.

5.2 Trends in Wage Variance and Gender Gaps

To understand how the transparency shock in 2001 shaped firms' wage-setting behavior, we first examine overall patterns of wage dispersion. A potential concern is that within-firm wage variation might be less relevant in the Norwegian context, given the

Figure 1: Wage Income, Gross Income and Taxable Income by Taxable Income Percentiles in 2000

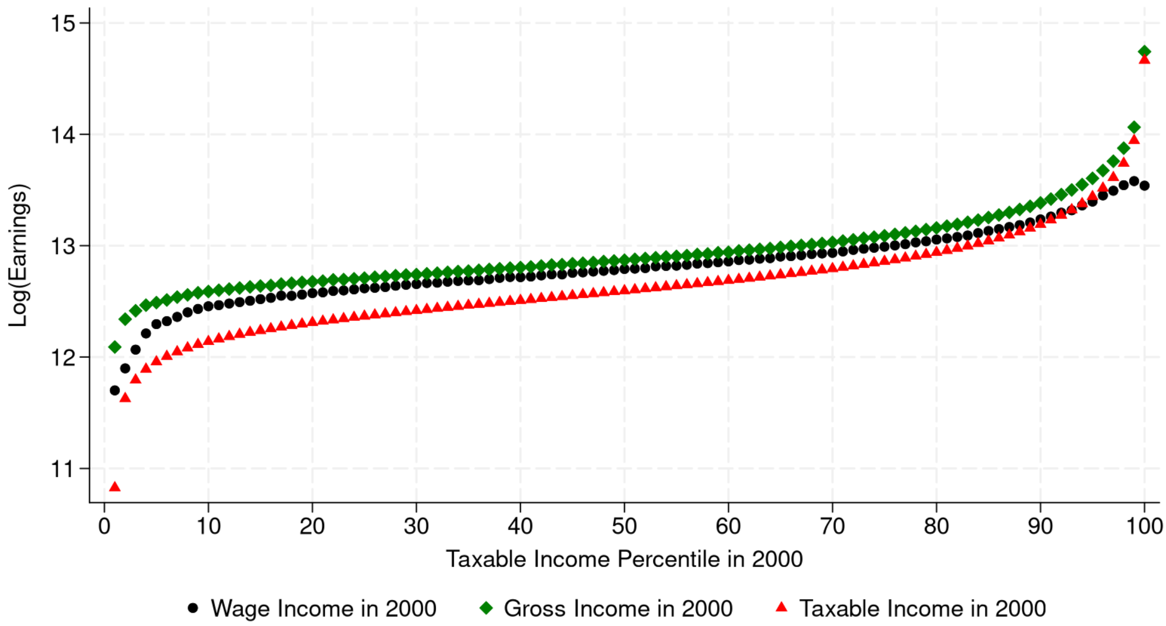

Note: Workers in our sample are ordered into percentiles based on their taxable income in 2000. For each percentile bin, we plot the average taxable income, gross income, and wage income to verify that the ranking order is preserved across these measures.

widespread coverage of union agreements, as in other Nordic countries. However, panel (a) of Figure 2 shows that within-firm variance accounted for roughly two-thirds of total wage variance from 1997 to 2006, indicating substantial discretion in wage-setting at the firm level. Moreover, the previously increasing trend in wage variance reversed abruptly in 2001, coinciding with the online searchable tax returns. This reversal was driven by a reduction in within-firm variance, while between-firm variance remained largely stable as presented in panel (b) of Figure 2.

Because the online tax records could only be searched by name—and not by occupation or employer—users were typically limited to looking up individuals they already knew, such as coworkers. This also constrained firms' ability to systematically identify and recruit workers from other employers without previous knowledge of individuals' names.

The variance pattern, together with the fact that individuals could only search for peers by name, provides strong evidence that the transparency shock operated primarily within firms. Our analysis therefore focuses on the within-firm gender gap. Gender disparities in Norway were relatively large during the period we study. Figure 3 reports the average within-firm gender wage gaps, calculated across firms employing both men

Figure 2: Wage Variance Within and Between Firms

(a) Variance (b) Variance Normalized to Year 2000

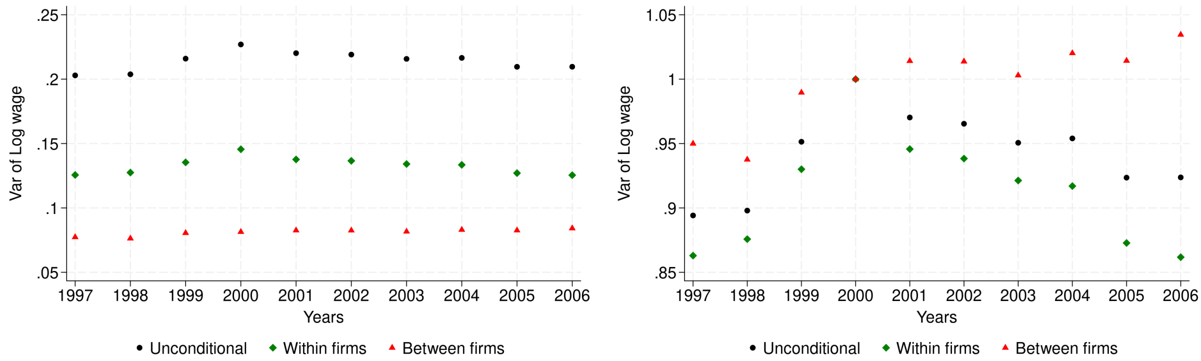

Note: In each year, the total wage variance is decomposed into within- and between-firm components as follows: $\frac{1}{N_t-1}\sum_{i=1}^{N_t}(w_{ift}-\overline{w_t})^2 = \frac{1}{N_t-1}\sum_{i=1}^{N_t}(w_{ift}-\overline{w_{ft}})^2 + \frac{1}{N_t-1}\sum_{i=1}^{N_t}(\overline{w_{ft}}-\overline{w_t})^2$, where N_t is the total number of workers in a given year, w_{ift} is the individual wage, $\overline{w_{ft}}$ is the firm's average wage, and $\overline{w_t}$ is the unconditional average wage. To preserve exact additivity—so that the within- and between-firm components sum to total variance—we do not correct the upward bias in the between-firm term. See Krueger and Summers (1988) for discussion of this bias. The naive measure can be viewed as a conservative choice for highlighting the relative importance of within-firm variation while keeping the decomposition transparent and simple.

and women and weighted by firm size. Like our main analysis, this descriptive measure emphasizes within-firm variation, but it does not account for worker mobility across firms with different baseline gaps, and it assigns weights based on total firm employment, rather than on the within-firm sizes of the male and female groups. On average, before 2001 women earned about 24.3% less than men. Following the transparency shock in 2001, the gap declined sharply to 23.5% and continued to fall, reaching 20.2% by 2006.

Descriptive statistics for the analysis sample, including demographic characteristics, mean income by source, the number of observations, counts of unique firms, and the firms' sectoral composition, are reported in Table A1 in the Appendix.

6 Empirical Strategy: Tax Transparency and Gender Gaps

To assess the causal effect of tax transparency on gender gaps, we exploit the 2001 transparency shock to estimate its differential impact on male and female workers. We use two complementary strategies. First, we measure gender-differential effects within the reference group using repeated cross-sectional data, allowing workers to change reference groups over time. This approach documents how the within-group gender gap evolves after the transparency shock. Second, we exploit the panel dimension of the data by fixing each worker's reference group as of year 2000 (prior to the transparency shock) and tracking the same individuals over the following three years to examine individual responses to the transparency shock. This approach helps identify the underlying mechanisms.

Figure 3: Average Within-Firm Gender Gaps

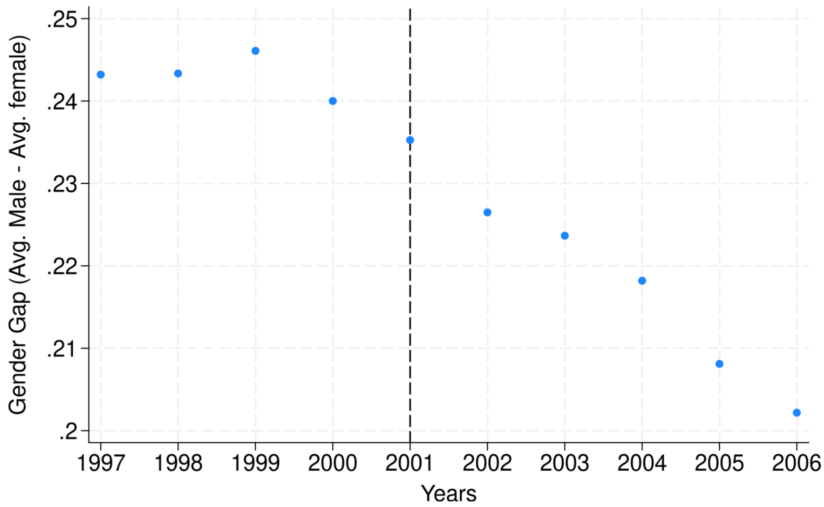

Note: For each firm employing both men and women, the within-firm gender wage gap is calculated as the difference between the average male and average female wage. The yearly series plots the weighted average of these firm-level gaps, using each firm's total number of workers as weights.

Our baseline approach is to define the reference group as all coworkers within the same firm. We focus on how the transparency shock affected within-firm gender gaps across both strategies.[12]

The first specification is

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

where $\log(w_{i,g,t})$ denotes the logarithm of annualized earnings for worker i in reference group g and year t (1997–2006), and Female_i is an indicator for female workers. $\mu_{t,g}$ are reference group-by-year fixed effects that absorb time-varying shocks at the group (e.g., firm) level, and $\mu_g \cdot \text{Female}_i$ are reference group-by-gender fixed effects that allow for time-invariant gender differences within groups; $\varepsilon_{i,g,t}$ is the error term.[13]

The event-time coefficients δ_j measure female-male differences in wage growth relative to the baseline year 2000. Pre-2001 coefficients provide a placebo assessment of parallel trends. Positive post-2001 coefficients indicate that the transparency shock raised women's wages relative to men's, narrowing the within-group gender wage gap. Because $\mu_{t,g}$ absorb group-specific shocks in each year, identification of δ_j comes from within-

[12]Note that in the second strategy we focus on how gender wage gaps evolve over time among baseline coworkers, holding the initial coworker group fixed.

[13]We omit individual fixed effects here; the panel specification introduced in Equation 3 directly addresses individual composition. Standard errors are clustered at the firm level throughout.

group changes over time—that is, comparisons of women and men exposed to the same firm environment in a given year.

We next set out the empirical strategy for the heterogeneity analyses. First, we estimate the baseline specification separately by sector and by industry. Second, we implement a triple-difference design to test whether effects are stronger in (i) industries with larger initial gaps and (ii) municipalities without prior printed tax lists—the latter serving as a validation of higher shock exposure.

The triple-difference specification is

$$\log(w_{i,g,h,t}) = \sum_{j \neq 2000} \delta_j \cdot 1\{t = j\} \cdot \text{Female}_i \cdot het_{i,h,t} + \sum_{j \neq 2000} \gamma_j \cdot 1\{t = j\} \cdot \text{Female}_i \\ + \mu_{t,g} + \mu_g \cdot \text{Female}_i + \varepsilon_{i,g,h,t} \tag{2}$$

where $het_{i,h,t}$ is an indicator for the heterogeneity dimension. In the first exercise, $het_{i,h,t} = 1$ for peer groups in industries with a within-firm wage premium gap above the median prior to the transparency shock, and 0 otherwise. In the second exercise, $het_{i,h,t} = 1$ for municipalities without printed tax catalogs prior to the transparency shock, and 0 otherwise. The coefficients δ_j are triple-difference terms: they measure the additional difference-in-difference effect on female wages relative to male wages in the het=1 subgroup compared with the het=0 subgroup at event time j (baseline j=2000). The coefficients γ_j report the corresponding difference-in-difference effects in the het=0 subgroup. The fixed effects are defined as in Equation 1: $\mu_{t,g}$ denotes reference-group-by-year (e.g., firm-by-year) fixed effects, and $\mu_g \cdot \text{Female}_i$ denotes reference-group-by-gender fixed effects. Standard errors are clustered at the firm level.[14]

Next, we implement an alternative strategy that exploits the panel dimension of the data by fixing reference groups in the year 2000 (prior to the transparency shock) and following the same individuals over the next three years to examine changes in gender wage gaps. This design tracks individual responses to the transparency shock while holding their initial reference group constant. To address mean reversion and gender-specific labor market dynamics that could bias estimates, we construct a placebo cohort by fixing reference groups in 1997 and following the same individuals through 2000. The placebo cohort captures the evolution of gender gaps and mobility patterns that would have occurred absent the transparency shock, providing a benchmark to net out pre-

[14]Because *het* is defined at the industry or municipality level, firm-by-year and firm-by-gender fixed effects subsume industry/municipality-by-year and industry/municipality-by-gender interactions; thus, including the latter would be collinear with the fixed effects.

existing gender-specific trends.

In this analysis, we restrict the sample to workers with more than one year of tenure at their firms, since the disclosed information——taxable annual income——is more accurate and informative for these individuals.

The specification is given by

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

where $y_{i,(g,t_0),c,t}$ is either the logarithm of annualized earnings or a binary mobility outcome for worker i in reference group (g, t_0), cohort $c \in \{1997, 2000\}$, and year t. Mobility outcomes include leaving the baseline full-time position (overall mobility), switching employers, experiencing non-employment, or moving to a part-time position. The coefficients δ_j are the parameters of interest and admit a triple-difference interpretation: they measure the change in the gender gap in each post-baseline year (2000-2003), net of the corresponding change observed in the placebo cohort (1997-2000). This design relies on the assumption that the placebo cohort represents the counterfactual evolution of gender gaps absent the transparency shock. Reference groups are fixed in 1997 or 2000 depending on the cohort, and we track the same individuals regardless of whether they remain at or leave the baseline firm. The reference group consists of all workers employed in the same firm in the baseline year t_0 of their cohort c. The term $\mu_{t,(g,t_0),c}$ captures reference group-by-year fixed effects.

When $y_{i,(g,t_0),c,t}$ denotes earnings, we include individual fixed effects in Equation 3, since wages vary over time for workers who remain in the labor force. By contrast, mobility outcomes exhibit limited within-person variation and are therefore not estimated with individual fixed effects; identification for those outcomes relies on cross-sectional variation.

Finally, we allow effects to vary along additional dimensions. Specifically, we compare workers who switch firms with those who stay. To disentangle responses that stem from workers' initial position in the wage distribution, we also allow effects to vary by baseline wage rank. In some specifications, we split the sample at the within-group median wage in the baseline year and estimate separate effects for individuals below and above the median.

7 Causal Results

7.1 General Effects of the Transparency Shock

We begin by presenting the difference-in-differences estimates of the average effect of the transparency shock on gender wage gaps. Table 1 reports results from models with progressively richer sets of controls. Column (1) shows estimates without additional fixed effects, while Columns (2) and (3) add reference group-by-time and reference group-by-gender fixed effects. In Column (2), the reference group consists of all coworkers within the same firm; in Column (3), it is further restricted to coworkers within the same firm who share the same education level and tenure category.[15]

Across all specifications, the estimates consistently indicate that the transparency shock increased female wages relative to male wages, thereby reducing within-group gender wage gaps. We emphasize the results from Columns (2) and (3), as their corresponding event-study graphs show parallel pre-period trends, supporting the parallel trends assumption and enabling a causal interpretation.[16] The estimated effects are statistically significant at the 1% level and imply that the transparency shock reduced the gender gap by 1.91 to 2.17 percentage points, relative to a conditional baseline gap of 22.4% to 24.9%.[17] This corresponds to a reduction of approximately 8.5% relative to its pre-shock level.

In the remainder of the analysis, we focus on specification (2), which defines the reference group as all coworkers within the same firm. Figure 4 presents the corresponding event-study coefficients, δ_j, estimated from Equation (1), illustrating the dynamic effects of the transparency shock on gender wage gaps. The results show that after the online disclosure of tax returns in 2001, female wage growth increases sharply by 0.78% relative to men. This effect persists and grows over time, reaching 3.38% by the final year of the sample (2006). The event-study patterns are similar whether the reference group is defined as all coworkers or restricted to coworkers with the same education level and tenure category.[18]

[15] Education is grouped 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 as less than one year versus more than one year.

[16] The pre-period trends for specification (1)-(3) are evaluated using the event-study graphs displayed in Figure A2 in the Appendix.

[17] "Conditional gender gaps" denotes gender gaps conditional on reference group-by-year fixed effects and are obtained by re-estimating the baseline model while replacing the reference-group-by-female fixed effects with a single female indicator; the coefficient on the female indicator summarizes the average pre-period gender wage gap within firms.

[18] See Figure A2 in the Appendix.

Table 1: Effect of Income Transparency on Female Wages Relative to Male Wages

Dep. Var: Log(Wages)	Group specification		
	(1)	(2)	(3)
After x Female	0.0147***	0.0217***	0.0191***
	(0.0021)	(0.0015)	(0.0014)
Conditional Gender Gap - Baseline	0.269	0.249	0.224
Observations	7,973,013	7,839,822	7,448,015
Reference Group	Everyone	Coworkers	Coworkers with Same Education, and Tenure
Reference Group - Time FEs	NO	YES	YES
Reference Group - Female FEs	NO	YES	YES

Note: Significance levels: * 10%, ** 5%, ***1%. Standard errors reported in parentheses are clustered at the firm level. "Conditional gender gaps" are obtained by re-estimating the baseline model, replacing the reference-group-by-female fixed effects with a single female indicator. In the original specification, the interaction between reference-group and female absorbs firm-specific baseline gender differences; in the alternative specification, the coefficient on the female indicator summarizes the aggregate conditional pre-period gender wage gap across firms.

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

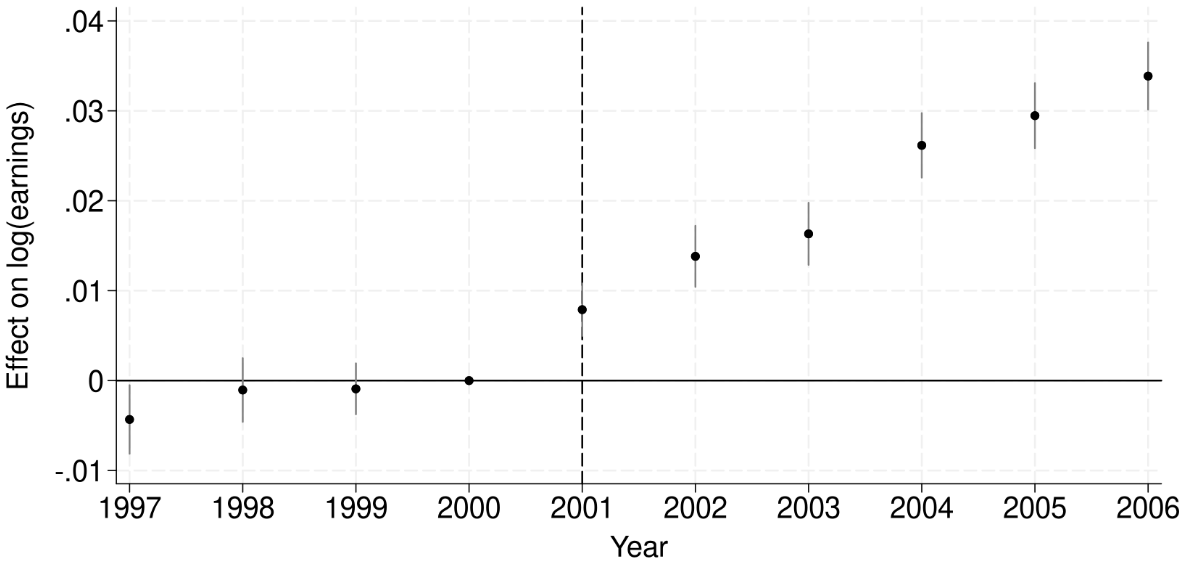

Note: Standard errors are clustered at the firm level. The regression includes year fixed effects. Confidence intervals are based on a 5% significance level. The vertical line marks the year of the transparency shock. The figure plots the interactions between the female indicator and year dummies.

7.2 Heterogeneity by Sector and Industry

Leveraging the richness of our data and the nationwide effect of the transparency shock, we next explore differential effects across sectors and industries. The results by sector are presented in Table 2. Column (1) repeats the baseline results using the full sample, while Columns (2) and (3) present results separately for private sector and public sector firms, respectively. The findings indicate that tax transparency is associated with a reduction in gender gaps across both sectors, but 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% reduction relative to the baseline conditional gender gap of 29.4%. Public sector firms saw a reduction of 1.5 percentage points, or 8.2% relative to the baseline gender wage gap of 18.3%. The smaller response in the public sector is consistent with its already highly transparent wage-setting system, which leaves less room for additional transparency to influence wages.

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

Dep. Var: Log(Wages)	Different Sector		
	All Sectors	Private	Public
	(1)	(2)	(3)
After x Female	0.0217***	0.0265***	0.0150***
	(0.0015)	(0.0021)	(0.0021)
Conditional Gender Gap - Baseline	0.249	0.294	0.183
Observations	7,839,822	4,947,051	2,728,116

Note: Significance levels: * 10%, ** 5%, ***1%. Standard errors reported in parentheses are clustered at the firm level. Results are estimated in separate regressions for each sector. All regressions include firm-by-time and firm-by-gender fixed effects."Conditional gender gaps" are obtained by re-estimating the baseline model, replacing the reference-group × female fixed effects with a single female indicator. In the original specification, the reference-group × female fixed effects absorb firm-specific baseline gender differences; in the alternative specification, the coefficient on the female indicator summarizes the aggregate conditional pre-period gender wage gap across firms. The total number of observations for the private and public sectors does not sum to the overall number of workers because some individuals are employed in unclassified sectors.

We next present results across 15 aggregated industry classifications. Figure 5 and Table 3 report the difference-in-differences estimates, stratified by the size of the baseline conditional gender gap in each industry. The point estimates are positive and statistically significant in 13 out of the 15 cases, and the magnitude of the reduction is positively correlated with the initial gap. For instance, mining and quarrying—the industry with the largest baseline gender gap of 37.5%—experienced a reduction of 4.6 percentage points, while education—the industry with the smallest gap of 14%—saw a reduction of 1.3 percentage points. A simple correlation shows that a 10-percentage-point increase in the baseline gap is associated with a 1 percentage point larger reduction after the transparency shock.

Figure 5: Effect of Income Transparency on the Gender Gap by Industry Baseline Gender Gap

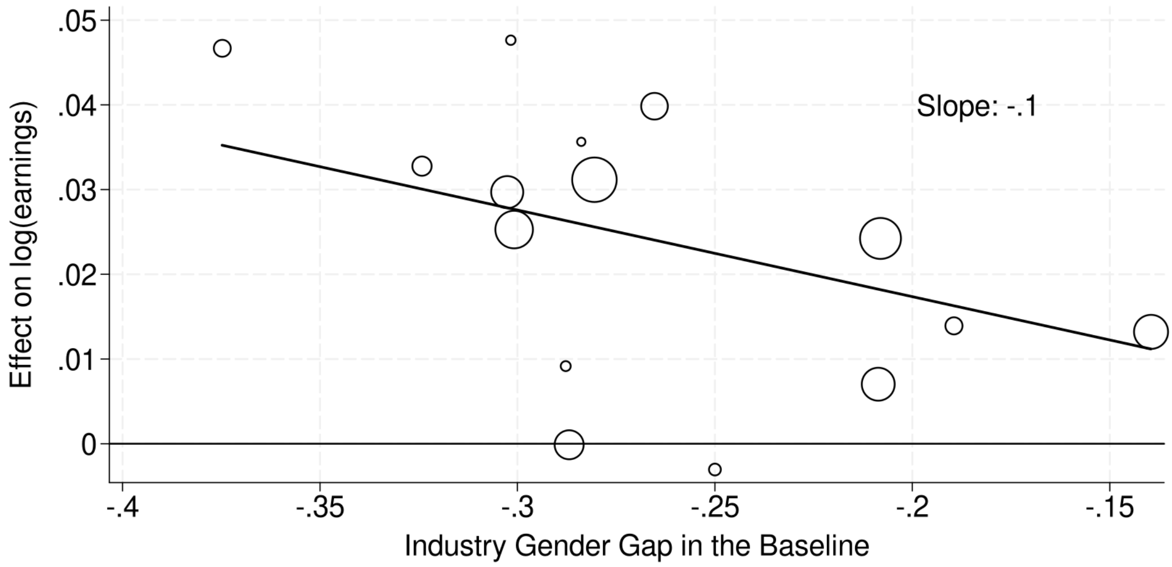

Note: The figure displays the estimated relative wage increase for females across 15 industry classifications, plotted against each industry's baseline gender wage gap. Each point represents an industry, and the scatter plot is weighted by the number of observations in that industry. The regression line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

Table 3: Effect of Income 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*	0.0467**	0.0312***	-0.0030	0.0399***	0.0253***	0.0477***	0.0092
	(0.0210)	(0.0187)	(0.0030)	(0.0082)	(0.0082)	(0.0030)	(0.0156)	(0.0157)
Conditional Gender Gap - Baseline	0.284	0.375	0.281	0.250	0.265	0.301	0.302	0.288
Observations	43,712	199,346	1,406,217	96,767	496,703	962,316	58,901	60,186
Dep. Var: Log(Wages)	Different Industries							
	Transport, Storage and Communication	Financial Intermediation	Real Estate	Public Adm	Education	Health and Social Work	Community, Social and Personal Activities	
After x Female	-0.0001	0.0328***	0.0297***	0.0071***	0.0132***	0.0243***	0.0139***	
	(0.0112)	(0.0061)	(0.0040)	(0.0023)	(0.0024)	(0.0049)	(0.0052)	
Conditional Gender Gap - Baseline	0.287	0.324	0.303	0.209	0.140	0.208	0.190	
Observations	588,276	260,316	707,673	763,070	820,273	1,181,806	193,645	

Note: Significance levels: * 10%, ** 5%, ***1%. Standard errors reported in parentheses are clustered at the firm level. Results are estimated in separate regression for each industry. All regressions include firm-by-time and firm-by-gender fixed effects. "Conditional gender gaps" are obtained by re-estimating the baseline model, replacing the reference-group × female fixed effects with a single female indicator. This coefficient on the female indicator summarizes the aggregate conditional pre-period gender wage gap across firms.

We further estimate the differential impact of the transparency shock on gender gaps across industries with high versus low average within-firm gender wage premium gaps in the pre-period. Industries are classified into these groups based on wage premium gaps estimated using a time-varying additive two-way fixed effects model for workers and firms for the year 2000, as described in Appendix B. Industries with gaps above the median are classified as high-gap, while those below are classified as low-gap. We classify at the industry level rather than the firm level to reduce the risk of spurious mean reversion effects.[19] Importantly, the results remain robust when using 1999 firm-level premium gap estimates.

Table 4 presents the aggregated results. We estimate Equation (2) both without and with firm-by-time fixed effects, shown in Columns (1) and (2), respectively. Using the wage premium classification as the additional source of variation, the estimates trace how the gender wage gap evolves across industries with different baseline disparities. The triple-difference estimates in the first row indicate that industries with larger pre-existing gender gaps experienced significantly greater relative wage gains for women, consistent with the transparency shock exerting a stronger corrective effect where disparities were initially larger. The estimated effects are positive and statistically significant in both specifications, with women's relative wages increasing by 0.69% without firm-by-time fixed effects and by 1.59% when they are included.

The transparency shock also reduced gender gaps in industries with smaller initial disparities, although the effects were more modest: female wages increased by 1.1% and 1.4% in the models without and with firm-by-time fixed effects, respectively, as shown in the second row of Table 4. The third row reports the estimated change in male wages across high- and low-gap industries, where coefficients are negative and statistically significant. This suggests that part of the narrowing gender gap is attributable to wage declines for men following the reform, provided the parallel trends assumption holds. However, panel (a) of Figure A3 in the Appendix shows that the parallel trends assumption is violated for the gender-specific difference-in-differences estimates: average wages for both men and women were already declining in high-gap industries relative to low-gap industries before the transparency shock. Despite this, the triple-difference estimates remain valid, as the pre-trends are parallel in both specifications, as illustrated in panels (b) and (c) of Figure A3 in the Appendix.

[19] Classifying units as "high-gap" based on an estimated pre-period premium can generate spurious mean reversion: if the baseline estimate is inflated by sampling noise or transitory shocks, subsequent estimates mechanically move back toward the mean even absent any effect of transparency. For example, a small firm may be labeled high-gap because of a one-off bonus or promotion cycle in 2000, an unusually male-heavy (or high-paid male) workforce mix that later normalizes through turnover, or measurement error in annualized earnings or hours in the baseline year.

Table 4: Effect of Income Transparency on the Gender Gap by Baseline Wage Premium Gap

Dep. Var. Log(Wage)	High Gap Industry	
	(1)	(2)
After x Female x Het.	0.0069**	0.0159***
	(0.0028)	(0.0029)
After x Female	0.0116***	0.0141***
	(0.0020)	(0.0023)
After x Het.	-0.0248***	
	(0.0024)	
Observations	7,955,066	7,839,822
Firm - Time FEs	NO	YES
Firm - Female FEs	YES	YES

Note: Significance levels: * 10%, ** 5%, ***1%. Standard errors reported in parentheses are clustered at the firm level. Column (1) includes education-by-year interacted with a dummy for tenure less than one year, a continuous tenure variable interacted with year, and firm-by-gender fixed effects. The model in Column (2) includes firm-by-year fixed effects and firm-by-gender fixed effects. *Het* is a dummy variable that equals 1 for peer groups in industries with a wage premium gap above the median prior to the transparency shock.

7.3 Validation using Prior Tax Disclosure Exposure

To validate our identification strategy, we exploit geographic variation in tax transparency before the online disclosure of tax records in 2001, as documented by Bø, Slemrod, and Thoresen (2015). Prior to 2001, some municipalities distributed printed tax lists, while others did not. In several areas, local groups — such as football clubs or community associations — sold copies of residents' tax records door-to-door. This variation provides a natural test of our interpretation: if the estimates truly reflect increased transparency, they should be stronger in municipalities without prior distribution of printed tax lists where the online publication represented a larger information shock.

Table 5 reports the aggregated results. We estimate the model both without (Column (1)) and with (Column (2)) firm-by-time fixed effects. Similar to the industry premium gap analysis, this specification also admits a triple-difference interpretation, capturing heterogeneous difference-in-differences effects by gender. In models without firm-by-time fixed effects, we can separately identify gender-specific impacts, which is crucial for assessing whether the narrowing of the gender wage gap is driven primarily by rising female wages or falling male wages.

The results confirm the expected pattern. The triple-difference coefficients reported in the first row of Table 5 indicate that municipalities without prior access to the tax catalog experienced a significantly larger increase in wages for female workers relative to

Table 5: Effect of Income Transparency on the Gender Gap by Tax Catalogs

Dep. Var. Log(Wage)	No Tax Catalog	
	(1)	(2)
After x Female x Het.	0.0102**	0.0120**
	(0.0045)	(0.0049)
After x Female	0.0115***	0.0091**
	(0.0037)	(0.0041)
After x Het.	-0.0015	
	(0.0037)	
Observations	1,643,234	1,615,069
Firm - Time FEs	NO	YES
Firm - Female FEs	YES	YES

Note: Significance levels: * 10%, ** 5%, ***1%. Standard errors reported in parentheses are clustered at the firm level. Column (1) includes 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. Column (2) includes firm-by-year and firm-by-gender fixed effects. *Het* is a dummy variable equal to 1 for municipalities without printed versions of the tax catalog prior to the transparency shock.

male workers, compared to municipalities where the catalog had already been available. This is consistent with the information shock being more pronounced in areas previously lacking printed versions of the tax lists. The estimated 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.

In municipalities that had prior access to printed catalogs, the transparency shock in 2001 also reduced gender gaps, although the magnitudes were smaller. The second row of the table shows that in these municipalities, female wages increased by 1.1% and 0.9% in the two specifications without and with firm-by-time fixed effects, respectively. Male wage changes are close to zero and statistically insignificant (reported in the third row), suggesting that the narrowing of gender gaps is driven by rising female wages rather than falling male wages. Panels (a)-(c) of Figure A4 in the Appendix show that the parallel trends assumption holds for the relevant difference-in-difference and triple-difference estimates, supporting a causal interpretation of the results. Importantly, the evidence also confirms that the narrowing of gender wage gaps is driven primarily by rising wages for female workers, rather than declining wages for males.[20]

[20]Estimating the effect of tax transparency on average wages using the municipality heterogeneity as the treatment in the pooled (non–gender-split) sample yields a positive but statistically insignificant coefficient.

7.4 Mechanisms

7.4.1 Worker Mobility and Individual Wage Growth

We investigate the mechanisms behind the narrowing gender gap by fixing workers' initial reference group and tracking the individuals over time, as described in Equation (3) in Section 6. This design allows us to separate gains arising from improved matching—where workers switch employers in response to new information—from gains generated by better offers within the same firm.

We begin by analyzing mobility, focusing on the likelihood of leaving the baseline employer, moving to a new firm, experiencing a non-employment spell, or switching to part-time work. Results of these linear probability models are reported in Table 6. The second row ("Female") shows pre-existing gender differences in the 1997 cohort: women were more likely than men to experience non-employment and part-time work and less likely to switch employers.[21] The first row ("Female x 1{c=2000}") reports the effect of the transparency shock on these gender differences. Column (1) shows that women's probability of any mobility-defined as leaving their full-time position in the baseline firm at least once during the three-year period-increased by 1.04 percentage points after the transparency shock.

Table 6: Effect of Income Transparency on Job Mobility

	Linear Probability Model (LPM)			
	(1) Any Mobility	(2) Job Switch	(3) Non-Employment	(4) Part-time
Female x 1{c=2000}	0.0104***	0.0068***	0.0139***	0.0121***
	(0.0025)	(0.0023)	(0.0024)	(0.0013)
Female	0.0024	-0.0172***	0.0046**	0.0469***
	(0.0019)	(0.0017)	(0.0020)	(0.0010)
Observations	1,303,813			

Note: Significance levels: * 10%, ** 5%, ***1%. Standard errors reported in parentheses are clustered at the baseline firm level. The dependent variable equals one if a worker experiences any mobility event-leaving the baseline full-time position, switching employers, experiencing a non-employment spell, or working part-time (in the same or another firm)—at least once during the three years after the baseline year. The latter three outcomes are not mutually exclusive. Yearly measures are based on employment status in December of each year. The 1997 and 2000 cohorts correspond to baseline years 1997 and 2000, respectively. The coefficient of interest is the interaction between the female indicator and the 2000 cohort dummy.

We then assess whether the higher increase in overall mobility primarily reflects job-to-job transitions, or movements into non-employment and part-time work. Column (2) shows that the probability of switching directly to a new employer increased by 0.68 percentage points—about two-thirds of the total mobility effect provided in Column

[21] Employment status is measured in December of each year.

(1). Column (3) shows a 1.39-percentage-point increase in the likelihood of experiencing a non-employment spell. Because these outcomes are measured as occurring at any year during the three-year window and are not mutually exclusive, job-to-job and non-employment probabilities do not sum to total mobility. Many mobility events involve a non-employment period before re-employment. Year-by-year transition patterns for these outcomes are shown in Figure A5 in the Appendix. Finally, Column (4) shows that the relative probability of part-time employment for women increased by 1.21 percentage points following the transparency shock.

We next examine wage effects over the three years following the transparency shock. Because women are more likely to experience non-employment after the transparency shock, we address potential selection in the next subsection. Although the online publication of tax returns modestly widened employment gaps in the short run, some workers may subsequently re-enter into better matches.

Figure 6 shows the evolution of gender wage gaps over this three-year horizon. To account for potentially changing workforce composition, we estimate models with individual fixed effects and comparison-group-by-year fixed effects, where comparison groups consist of workers employed in the same firm at baseline. This specification helps isolate wage changes among individuals who remain active in the labor market.[22]

Figure 6 reports the relative changes in the gender wage gap following the increase in tax transparency. First, both sets of estimates show that female workers experienced slower wage growth than male coworkers during the three years after the baseline period in both the 1997 and 2000 cohorts. This highlights the importance of accounting for gender-specific career dynamics that would have occurred even in the absence of the transparency shock. We address this by using the 1997 cohort as a placebo group, which nets out the underlying gender-specific trends from the estimates for the 2000 cohort. Second, the triple-difference estimates show that the transparency shock increased relative wages for women who remained in the labor market, thereby narrowing the gender wage gap. On average, women's relative wage growth following the transparency shock is 0.52% across all post-baseline years, with more pronounced gains of 0.76% in the second year and 0.66% in the third year.

To further understand the underlying dynamics of the gender wage gap reductions, 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

[22]Including individual fixed effects removes variation from individuals who permanently exit the labor market, so the difference-in-differences coefficients reflect wage changes among workers who remain active after the publication of tax returns.

Figure 6: Effect of Income Transparency on Gender Wage Gaps

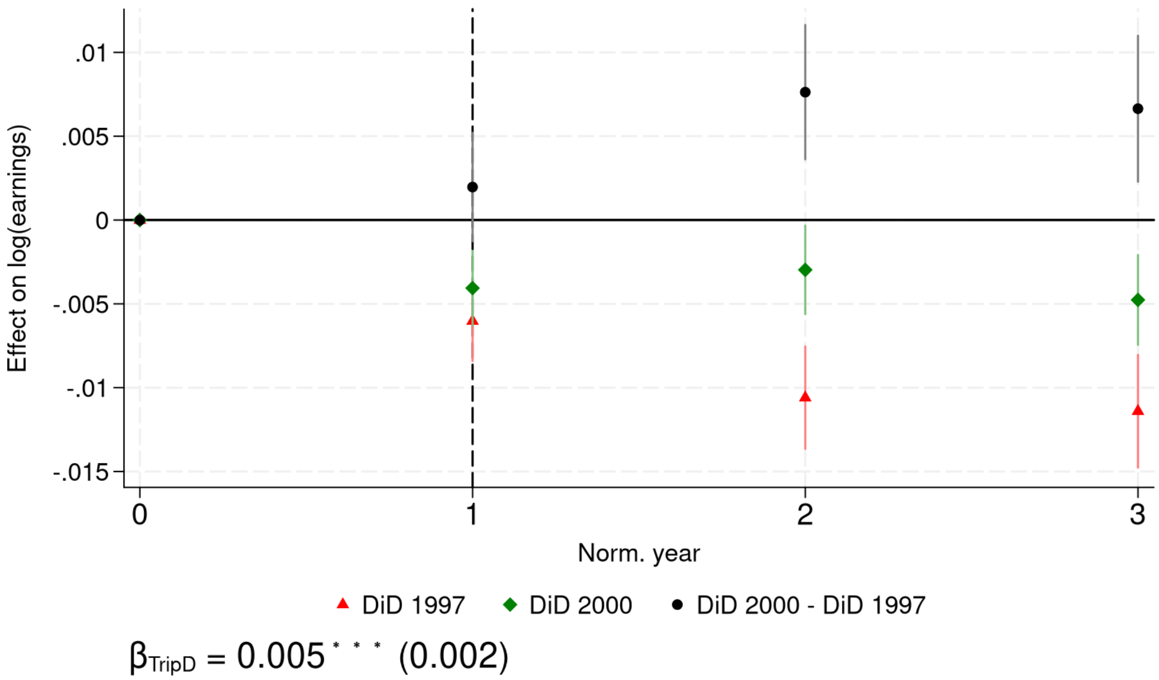

Note: Confidence intervals are shown at the 5% significance level, where standard errors are clustered at the baseline firm level. The model is estimated with individual fixed effects and firm-by-year-by-cohort fixed effects. The 1997 and 2000 cohorts correspond to samples where the baseline year (Norm. year = 0) is set to 1997 and 2000, respectively. Coefficients are reported separately for each cohort. The coefficients displayed in black circles correspond to the interaction between the female indicator, the 2000 cohort dummy, and the normalized post-baseline year dummies.

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 7. Columns (1) and (3) report 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, respectively. Columns (2) and (4) add interaction terms for switchers, isolating the effects for switchers and stayers across all post-baseline years and in the third year.

The first-row coefficients show that the average relative wage growth observed for female workers among stayers is –0.11% across all periods and 0.42% in the third year, with only the second estimate being statistically significant at the 10%-level. The second-row coefficients indicate that the female–male wage effect is larger for switchers than for stayers— 1.32% over the full post-baseline period and 0.66% in year three. The third row shows that, for men, switching is associated with –0.65% across all years (statistically significant) and 0.12% in year three (not significant). Taken together, these patterns imply that, by the third year, the post-disclosure gains from switching accrue to women, not men, and are the primary driver of the observed narrowing in the gender wage gap: women's relative wage growth is 0.42% among stayers and 1.08% among switchers.

Table 7: Effect of Income Transparency on Gender Wage Gaps for Switchers and Stayers

Dep. Var. Log (Wage)	Gender Effects			
	All Years		3 Years After	
	(1)	(2)	(3)	(4)
After x Female x $1\{c = 2000\}$	0.0052***	-0.0011	0.0068***	0.0042*
	(0.0017)	(0.0018)	(0.0022)	(0.0025)
After x Female x $1\{c = 2000\}$ x Switch		0.0132***		0.0066*
		(0.0026)		(0.0037)
After x $1\{c = 2000\}$ x Switch		-0.0065***		0.0012
		(0.0018)		(0.0026)
Observations	4,010,410	4,010,410	1,776,352	1,776,352

Note: Significance levels: * 10%, ** 5%, ***1%. Standard errors reported in parentheses are clustered at the baseline firm level. 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 during the three years after the baseline year. Workers who leave their baseline firm after the baseline year and never return to the labor market do not contribute to the identification of any coefficients, and are effectively excluded from the sample.

Table C1 in the Appendix repeats the heterogeneity analysis, comparing (i) firms with high versus low baseline wage-premium gaps and (ii) municipalities that did versus did not circulate printed tax catalogs prior to the national transparency shock. The point estimates are consistent with the main results, although the smaller panel sample reduces precision, leaving some effects statistically insignificant.

7.4.2 Effects on Different Income Sources

We next examine whether wider non-employment spells following the transparency shock had adverse consequences for women that we disregard in the previous analyses of wages. To that end, we study gender differentials in the receipt of distinct income sources, measured as receiving at least once in the three years after the baseline year: (i) any income; (ii) labor income from any work activity (employee earnings plus business/self-employment income); (iii) employee earnings; (iv) business/self-employment income; and (v) social assistance (unemployment insurance, disability insurance, and other transfers). Results are reported in Table 8.

Column (1) shows no gender-differential change in the probability of receiving any income after the transparency shock. Column (2) indicates a small statistically insignificant increase for women in the probability of earning labor income. By contrast, the composition of labor income shifts significantly: women became less likely to earn employee wages (Column 3) and more likely to earn business/self-employment income (Column 4). Finally, we find no significant gender-differential change in the probability of receiving social assistance such as unemployment insurance, disability insurance, or other transfers (Column 5).

Table 8: Effect of Income Transparency on Probability of Earning Various Income Sources

	Linear Probability Model (LPM)				
	(1)	(2)	(3)	(4)	(5)
	Gross inc.	Labor inc.	Wage inc.	Bus. Inc.	Social Assis.
Female x 1{c=2000}	0.0002	0.0005*	-0.0032**	0.0055***	-0.0004
	(0.0003)	(0.0003)	(0.0012)	(0.0011)	(0.0014)
Female	0.0009***	0.0004**	-0.0009	-0.0487***	0.0285***
	(0.0002)	(0.0002)	(0.0009)	(0.0014)	(0.0011)
Observations	1,303,813				

Note: Significance levels: * 10%, ** 5%, ***1%. Standard errors reported in parentheses are clustered at the baseline firm level. The models are estimated using cross-sectional variation (without individual fixed effects), with the dependent variable equal to one if individuals received any of the following income sources during the three years after the baseline year: (1) Gross income, (2) Labor Income (employee plus business/self-employment income), (3) Employee earnings, (4) Business/self employment income, (5) Social assistance (unemployment insurance, disability insurance, and other transfers). All income categories are here measured by annual tax information. The 1997 and 2000 cohorts correspond to baseline years 1997 and 2000, respectively. The coefficients are reported separately for each cohort. The main coefficient is the interaction between the female indicator and the 2000 cohort dummy.

As we find no gender-differential selection on the extensive margin of "any income" (gross income), we also report gross income over the three years following the transparency shock—as a broader measure that captures women's increased movement into self-employment and avoids understating effects in wage-only analyses. If this reallocation reflects improved outside options induced by the information shock, gross income should rise. Figure 7 plots the evolution of the gender gap in gross income using models with individual fixed effects and comparison-group-by-year fixed effects, where comparison groups are defined by baseline co-workers within the same firm. The triple difference estimates indicate that transparency increased women's gross income relative to men's in each post-reform year; on average, women's relative gross income rose by 1.2% over the three years.

Figure 7: Effect of Income Transparency on Gross Income Gender Gaps

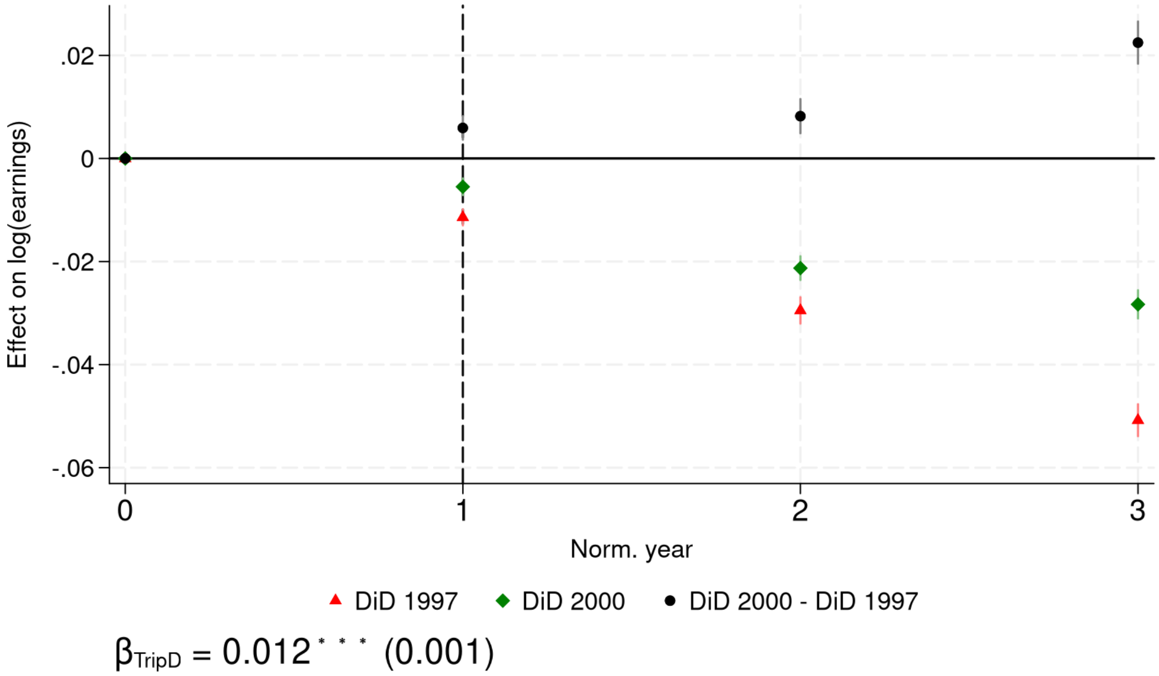

Note: Confidence intervals are constructed at the 5% significance level, where standard errors are clustered at the baseline firm level. The model includes individual fixed effects and firm-by-year-by-cohort fixed effects. Gross income is measured using the annual tax information. The 1997 and 2000 cohorts correspond to samples where the baseline year (Norm. year = 0) is set to 1997 and 2000, respectively. Coefficients displayed in black circles correspond to the interaction between the female indicator, the 2000 cohort dummy, and year dummies.

7.4.3 Are Gender Gap Reductions Driven by Wage Compression?

One possible explanation for the reduction in the gender wage gap following the transparency shock is that female workers are typically located at the lower end of the wage distribution within their comparison group. In this section, we test whether differences in their relative position during the pre-period can explain the observed change in the gender wage gaps. We begin by examining whether tax transparency caused a general 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 8: Effect of Income Transparency on Gender Wage Gaps for Low Income - Relatively to High Income Workers

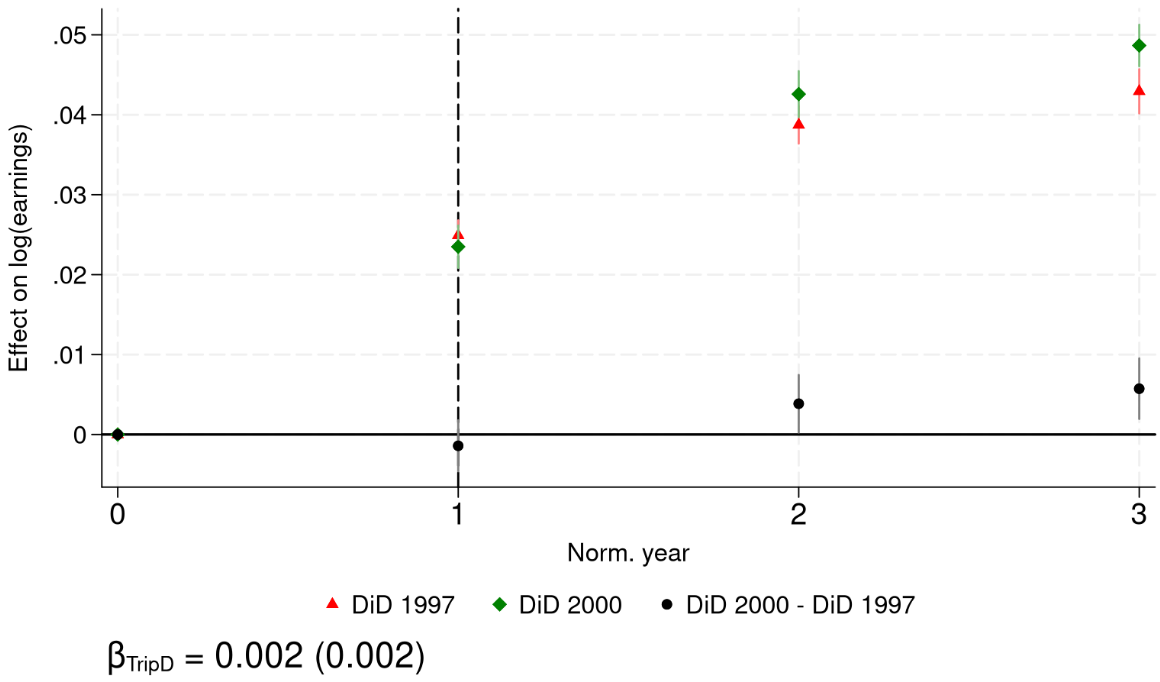

Note: Confidence intervals are constructed at the 5% significance level, where standard errors are clustered at the baseline firm level. The model is estimated using individual fixed effects and firm-by-year-by-cohort fixed effects. The coefficients are reported separately for each cohort. The coefficients displayed in black circles correspond to the interaction between an indicator for workers below the median in their baseline firm, the 2000 cohort dummy, and year dummies.

Figure 8 reports the relative wage increases of workers below the median compared to those above the median in their comparison group. First, both sets of estimates show that workers below the median experience a higher wage growth than those above the median during the three years after the baseline period, for both cohorts. This underscores the importance of controlling for mean reversion, which we address by using the 1997 cohort as a placebo to net out the effects 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 statistically significant higher relative wage growth of 0.57% in the third year after the baseline year. Related termination and job transition effects are shown in Figure A6 in the Appendix.

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 9 presents the overall impact of the tax transparency shock 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, both cohorts show that wage growth declines with a worker's relative position within the firm—those at the bottom of the distribution naturally experience higher wage growth, consistent with mean reversion. However, this pattern is significantly more pronounced for the 2000 cohort, who experienced the increased transparency. Effects are concentrated among low-wage workers. Using the 1997 cohort as a placebo control, only the bottom three baseline wage deciles show positive, statistically significant wage responses, with magnitudes that decline moving up the distribution—for example, a 3.7% relative wage increase in the first decile and 0.73% in the third—while estimates for deciles 4–10 are small and statistically indistinguishable from zero.

Figure 9: Difference-in-Differences Coefficients for the Effect of Income Transparency on Wages Across Within-Firm Wage Deciles

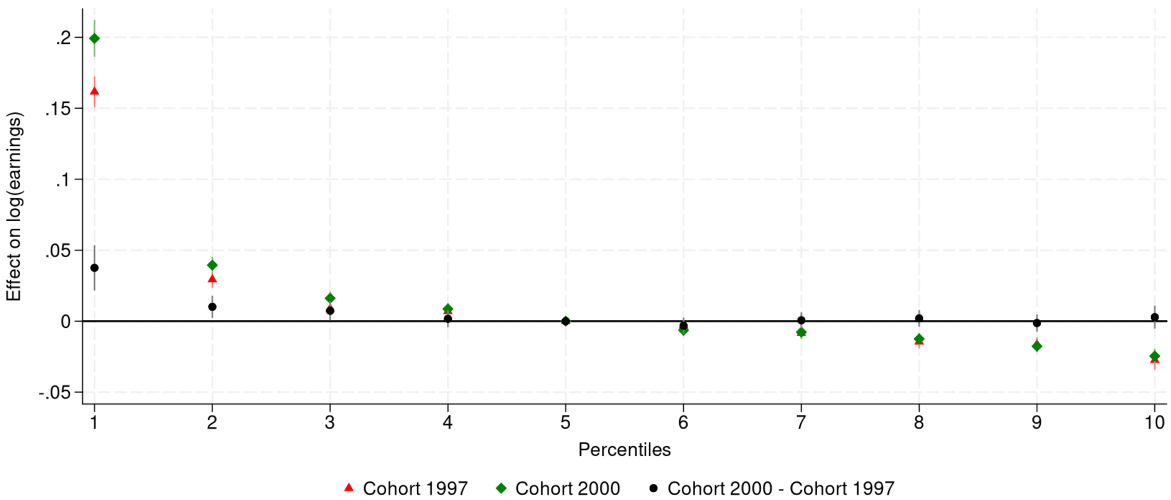

Note: Standard errors are clustered at the baseline firm level, and all regressions include year fixed effects. Confidence intervals are constructed at the 5% significance level. The model includes individual fixed effects and firm-by-year-by-cohort fixed effects. The coefficients displayed in black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

We formally test whether the decline in gender wage gaps reflects wage compression induced by the transparency reform. Table 9 reports difference-in-differences estimates that condition on workers' baseline wage rank. Column (1)—without any rank controls—shows that women's wages grow 0.86% more than men's after the reform.[23] Columns (2) and (3) sequentially add baseline wage-decile–by–year indicators and then comparison-group fixed effects; the estimated effect remains essentially unchanged, indicating that

[23]This effect is larger than in the previous section because we restrict comparison groups to those with at least 10 workers to form within-group baseline wage deciles.

the baseline wage position does not explain the narrowing of the gender gap. Columns (4) and (5) split the sample into the bottom and top halves of the baseline wage distribution; we find no statistically significant differences across halves. Overall, the reduction in the gender wage gap is broad-based across the wage distribution rather than concentrated among low- or high-wage workers, providing little support for a pure wage-compression mechanism.

Table 9: Effect of Income Transparency on Gender Wage Gaps After Controlling for Wage Compression

Dep. Var: Log(Wage)	Relative position interacted with gender				
	(1)	(2)	(3)	(4)	(5)
After x Female x 1{c=2000}	0.0086***	0.0065**	0.0092***	0.0110***	0.0068**
	(0.0024)	(0.0027)	(0.0029)	(0.0026)	(0.0032)
After x Female x 1{c=2000} $\cdot$ $\text{Below}_{i,(g,t_0),c}$				-0.0080*	0.0045
				(0.0044)	(0.0054)
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	1,383,926	1,383,926	1,147,206	1,383,926	1,147,206

Note: Significance levels: * 10%, ** 5%, ***1%. Standard errors reported in parentheses are clustered at the baseline firm level. 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 year.

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 an abrupt shift in transparency in Norway in 2001, when individual tax returns were made publicly accessible and searchable online by national newspapers. This broad increase in transparency provides a unique setting to assess how enhanced access to income 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 both private and public sectors, with the strongest effects in the private sector, and across all industries, with a stronger reduction in industries with 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 transparency shock, adjusting for typical labor market dynamics. Our findings show that increased transparency led to higher job separation rates for female workers, followed by more frequent job-to-job transitions, ultimately resulting in higher wage growth, which is also evident in broader measures of income. This wage increase is more pronounced 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, but conclude that 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, Halberstam, Yosh, Kroft, Kory, Mas, Alexandre, and Messacar, Derek (2023). “Pay Transparency and the Gender Gap”. *American Economic Journal: Applied Economics* 15.2, pp. 157–183.

Bennedsen, Morten, Simintzi, Elena, Tsoutsoura, Margarita, and Wolfenzon, Daniel (2022). “Do Firms respond to Gender Pay Gap Transparency?” *The Journal of Finance* 77.4, pp. 2051–2091.

Blau, Francine D. and Kahn, Lawrence M. (2017). “The Gender Wage Gap: Extent, Trends, and Explanations”. *Journal of Economic Literature* 55, pp. 789–865.

Blundell, Jack, Duchini, Emma, Simion, Ştefania, and Turrell, Arthur (2025). “Pay Transparency and Gender Equality”. *American Economic Journal: Economic Policy* 17.2, pp. 418–445.

Bracha, Anat, Gneezy, Uri, and Loewenstein, George (2015). “Relative Pay and Labor Supply”. *Journal of Labor Economics* 33.2, pp. 297–315.

Breza, Emily, Kaur, Supreet, and Shamdasani, Yogita (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”. *IZA Discussion Papers* 14206.

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”. *American Economic Review* 102.6, pp. 2981–3003.

Cullen, Zoë (2024). “Is Pay Transparency Good?” *Journal of Economic Perspectives* 38.1, pp. 153–180.

Cullen, Zoë B and Pakzad-Hurson, Bobak (2023). “Equilibrium Effects of Pay Transparency”. *Econometrica* 91.3, pp. 765–802.

Cullen, Zoë and Perez-Truglia, Ricardo (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”. *The American Economic Review* 109.2, pp. 620–663.

Gulyas, Andreas, Seitz, Sebastian, and Sinha, Sourav (2023). "Does Pay Transparency Affect the Gender Wage Gap? Evidence from Austria". *American Economic Journal: Economic Policy* 15.2, pp. 236–255.

Hvidberg, Kristoffer B., Kreiner, Claus T., and Stantcheva, Stefanie (2023). "Social Positions and Fairness Views on Inequality." *Review of Economic Studies* 90.6, pp. 3083–3118.

Jäger, Simon, Roth, Christopher, Roussille, Nina, and Schoefer, Benjamin (2024). "Worker Beliefs about Outside Options". *Quarterly Journal of Economics* 139.3, pp. 1505–1556.

Kiessling, Lukas, Pinger, Pia, Seegers, Philipp, and Bergerhoff, Jan (2024). "Gender Differences in Wage Expectations and Negotiation". *Labour Economics* 87, p. 102505.

Krueger, Alan B and Summers, Lawrence H (1988). "Efficiency Wages and the Inter-Industry Wage Structure". *Econometrica* 56.2, pp. 259–293.

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, pp. 1683–1721.

Perez-Truglia, Ricardo (2020). "The Effects of Income Transparency on Well-Being: Evidence from a Natural Experiment". *American Economic Review* 110.4, pp. 1019–1054.

Reck, Daniel, Slemrod, Joel, and Vattø, Trine Engh (2022). "Public Disclosure of Tax Information: Compliance Tool or Social Network?" *Journal of Public Economics* 212, p. 104708.

Rege, Mari and Solli, Ingeborg F. (2015). "Lagging Behind the Joneses: The Impact of Relative Earnings on Job Separation". *Unpublished*.

Reuben, Ernesto, Wiswall, Matthew, and Zafar, Basit (2017). "Preferences and Biases in Educational Choices and Labour Market Expectations: Shrinking the Black Box of Gender". *The Economic Journal* 127.604, pp. 2153–2186.

Roussille, Nina (2024). "The Role of the Ask Gap in Gender Pay Inequality". *The Quarterly Journal of Economics* 139.3, pp. 1557–1610.

Appendix

A Supplementary Tables and Figures

Table A1: Summary statistics of the main estimation sample

	Average of the Variables of The Years										
Year	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006	Total
Observations	706,686	745,671	776,727	786,967	810,171	836,187	829,263	808,602	831,603	841,152	799,476.9
Female	0.375	0.378	0.384	0.396	0.396	0.411	0.421	0.410	0.418	0.434	0.403
Education Level											
Secondary or less	0.652	0.646	0.623	0.617	0.608	0.594	0.586	0.580	0.568	0.563	0.602
College	0.270	0.276	0.299	0.304	0.310	0.322	0.330	0.328	0.341	0.343	0.314
Graduate Level	0.077	0.078	0.078	0.079	0.082	0.083	0.084	0.092	0.091	0.093	0.084
Wage Income	390,598	409,377	418,917	420,039	432,551	445,622	449,352	466,752	477,895	491,017	441,586
Labor income	393,967	412,713	422,073	423,545	435,856	448,787	451,660	469,808	483,376	493,725	444,926
Taxable Income	333,162	345,729	352,421	363,227	351,258	369,386	382,726	409,568	432,605	393,624	374,474
Gross Income	425,969	445,575	457,713	469,751	469,463	493,374	502,382	519,927	551,680	519,201	487,130
Matched employer-employee data											
Annualized wage income	381,100	398,599	407,271	407,534	420,100	432,870	437,924	453,899	465,563	477,962	429,610
Number of Firms	60,108	63,627	65,763	66,932	67,947	68,195	67,368	65,568	68,717	68,248	66,371.8
Number of Employees in the Firms	221.7	225.4	229.0	221.4	212.9	209.0	217.6	233.5	219.0	219.5	220.8
Sector											
Public	0.350	0.336	0.326	0.341	0.335	0.347	0.346	0.337	0.356	0.367	0.344
Private	0.627	0.642	0.652	0.637	0.642	0.630	0.631	0.639	0.620	0.610	0.633

The table reports descriptive statistics for the main estimation sample. For each year, we report the number of observations, means of background variables (gender and education), and key income measures from tax records: wage income; labor income (wage plus business/self-employment income); taxable income (revealed after the information shock), and gross income. We also include the mean annualized wage income constructed from the matched employer–employee dataset (primary outcome) and counts of unique firms overall and by sector.

Figure A1: Difference Between log(Gross Income) and log(Wage Income) by Taxable Income Percentiles in 2000:

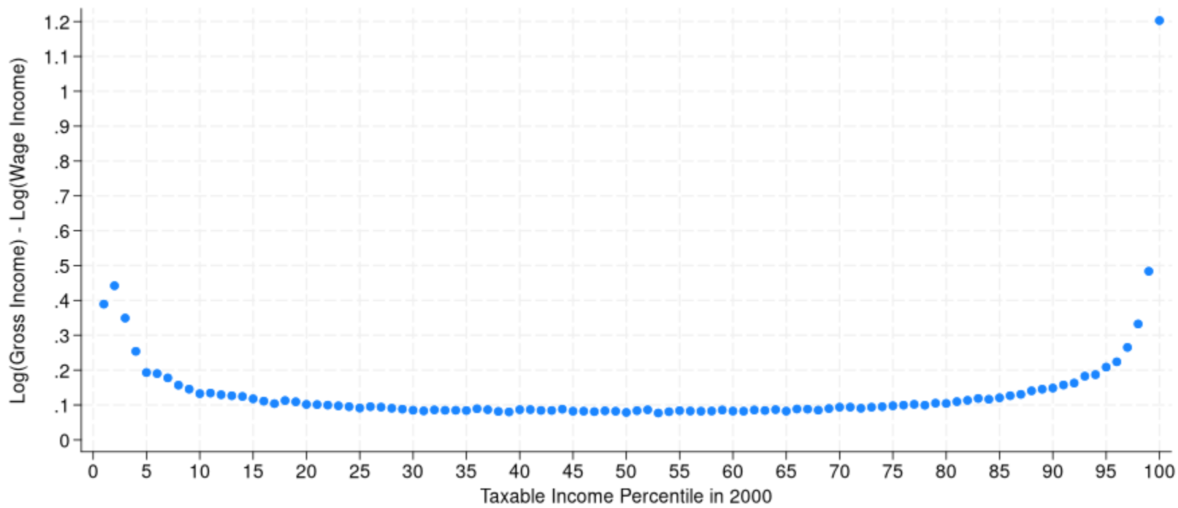

Note: Workers in our sample are ordered into percentiles based on their taxable income in year 2000. For each percentile bin, we plot the average of the individual log difference between gross income and wage income.

Figure A2: Difference-in-Differences Estimates of the Effect of Tax Transparency on Female Wages Relative to Male Wages with Alternative Controls

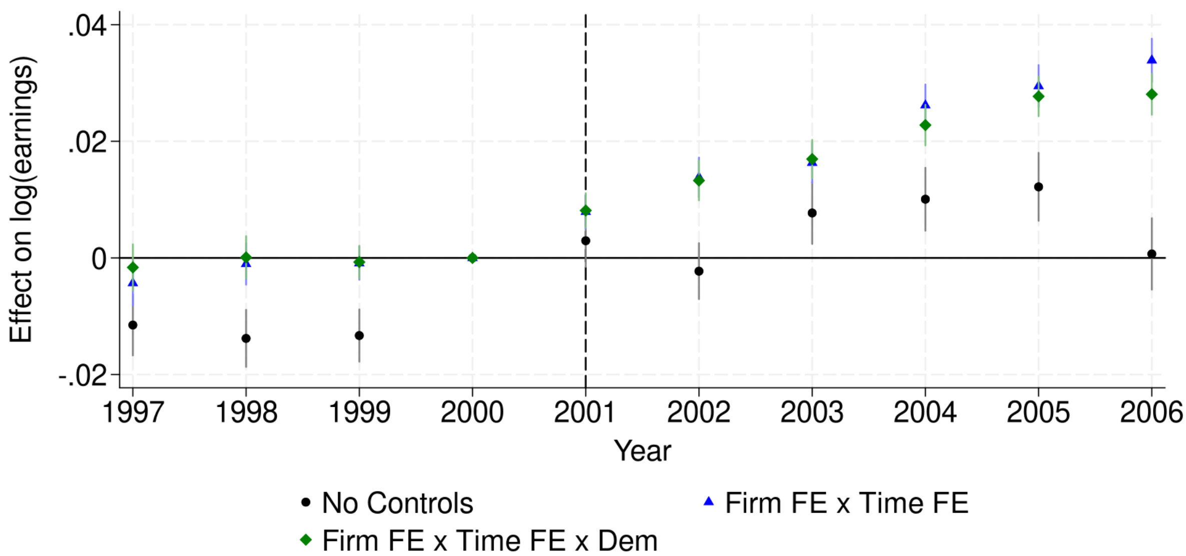

Note: Confidence intervals are generated using a 5% significance level, where standard errors are clustered at the firm level. All regressions include year fixed effects. The vertical line indicates the year tax returns became publicly available online. The figure plots interactions between the female indicator and year dummies under alternative sets of fixed effects (see legend).

Figure A3: Effect of Income Transparency on the Gender Gap by Baseline Wage Premium Gap. Estimates for High Gender Gap Industries Relatively to Low Gender Gap Industries.

(a) Gender-specific DiD Estimates
(Excluding Firm-by-year FE)

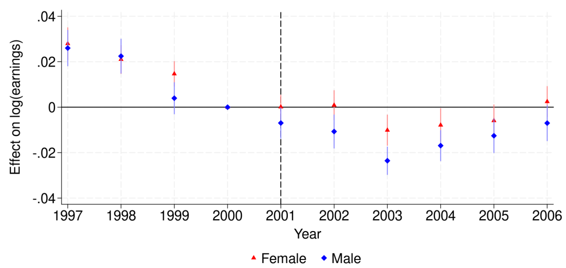

(b) Triple-Difference Estimates
(Excluding Firm-by-year FE)

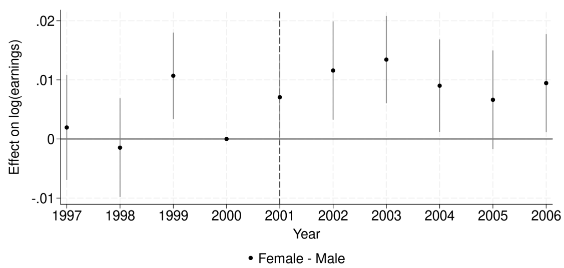

(c) Triple-Difference Estimates
(Including Firm-by-year FE)

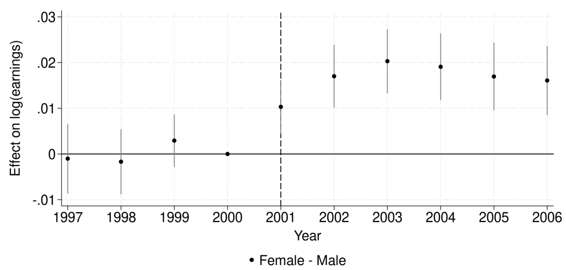

Note: Confidence intervals are constructed at the 5% significance level, where standard errors are clustered at the firm level. All regressions include year fixed effects. The vertical line indicates the year when tax returns were first published online. Models in panels (a) and (b) include education-by-year 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. Panel (a) reports coefficients equivalent to DiD estimates for males and females, using industries with a high wage premium gap as the treated group. Panels (b) and (c) report triple-difference coefficients directly. Due to collinearity, specifications with firm-by-year and firm-by-gender fixed effects cannot separately identify DiD estimates for males and females.

Figure A4: Effect of Income Transparency on the Gender Wage Gap by Prior Access to Tax Catalogs. Estimates Compare Individuals in Municipalities Without Prior Access to Those with Access.

(a) Gender-specific DiD Estimates (Excluding Firm-by-year FE)

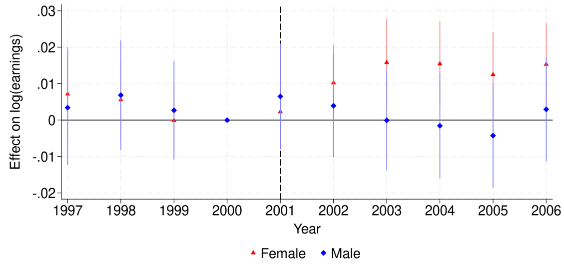

(b) Triple-Difference Estimates (Excluding Firm-by-year FE)

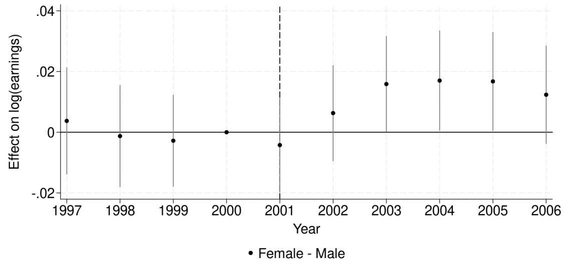

(c) Triple-Difference Estimates (Including Firm-by-year FE)

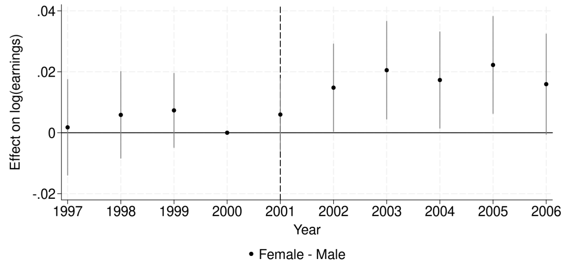

Note: Confidence intervals are constructed at the 5% significance level, where standard errors are clustered at the firm level. All regressions include year fixed effects. The vertical line indicates the year tax returns were first published online. Models in panels (a) and (b) include education-by-year 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. Panel (a) reports coefficients equivalent to DiD estimates for males and females, using municipalities without prior access to tax catalogs as the treated group. Panels (b) and (c) report triple-difference coefficients directly. Due to collinearity, specifications with firm-by-year and firm-by-gender fixed effects cannot separately identify DiD estimates for males and females.

Figure A5: Effect of Income Transparency on Job Mobility for Female Workers Relatively to Male Workers

(a) Any Mobility

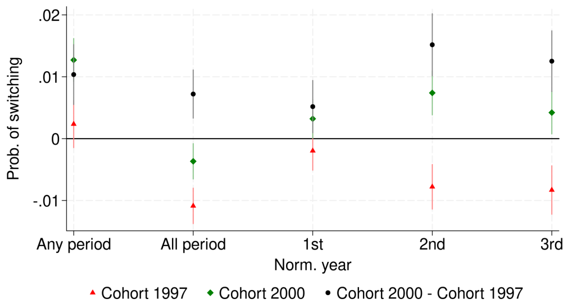

(b) Job Switch

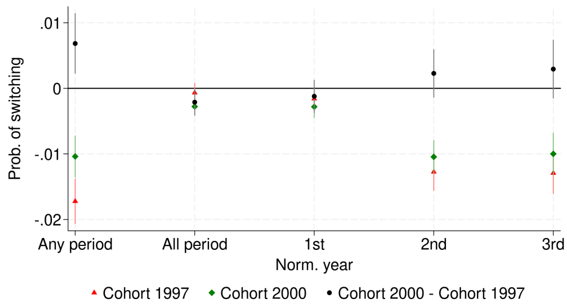

(c) Non-Employment

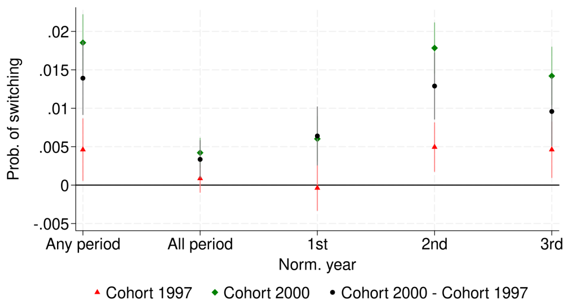

Note: Confidence intervals are constructed at the 5% significance level, where standard errors are clustered at the firm level. All regressions include year fixed effects. Models use cross-sectional variation (no individual fixed effects), with the dependent variable equal to one if there was a termination, job switch or an employment gap during the period indicated on the x-axis. The yearly measures are constructed using information on workers' employment status in December of each year (see data section). It includes firm-by-cohort fixed effects. "Any period" is a dummy variable equal to 1 if the worker is not employed at the baseline firm in at least one of the three years 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. 1st, 2nd and 3rd represents one, two and three years after the baseline year. The 1997 and 2000 cohorts correspond to samples where the baseline year is set to 1997 and 2000, respectively. The coefficients displayed in black points correspond to the interaction between the female indicator and the dummy for the 2000 cohort.

Figure A6: Effect of Income Transparency on Job Mobility for Low Income Relatively to High Income Workers

(a) Any Mobility

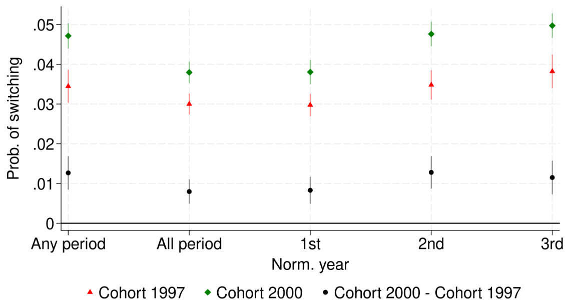

(b) Job Switch

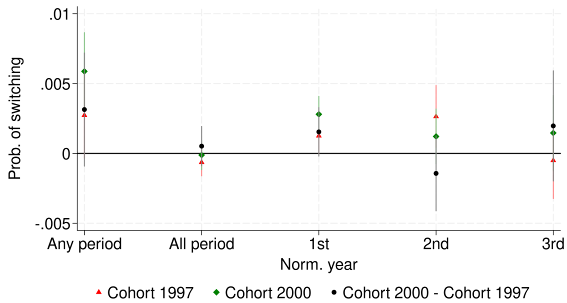

(c) Non-Employment

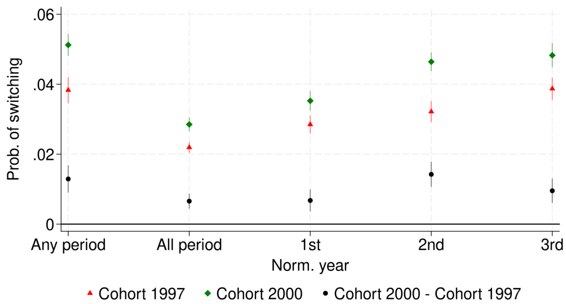

Note: Confidence intervals are constructed at the 5% significance level, where standard errors are clustered at the firm level. All regressions include year fixed effects. The models are estimated using cross-sectional variation (no individual fixed effects), with the dependent variable equal to one if there was a termination, job switch or employment gap during the period indicated on the x-axis. The yearly measures are constructed using information on workers' employment status in December of each year (see data section). It includes firm-by-cohort fixed effects. "Any Period" is a dummy variable equal to 1 if the worker is not employed at the baseline firm in at least one of the three years 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. 1st, 2nd and 3rd represents one, two and three years after the baseline year. The 1997 and 2000 cohorts correspond to samples where the baseline year is set to 1997 and 2000, respectively. The coefficients displayed in black points correspond to the interaction between an indicator for workers below the median and the dummy for the 2000 cohort.

B The 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.

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

$log(w_{it})$ represents the logarithm of the earnings variable for worker i in year t, α_i captures individual fixed effects, $\Psi^{G(i)}_{J(i,t),t}$ 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 estimated 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.

C Heterogeneous Effects of the Individual Effects

To further assess whether the observed effects are driven by increased transparency, we examine whether the estimates of the gender wage gap when we track individuals over time, as presented in Section 7.4, differ across firms located in municipalities with and without prior distribution of printed tax catalogs, as well as across industries with high and low baseline wage premium gaps. These heterogeneity analyses, reported in Table C1, focus on wage effects three years after the transparency shock. Although the coefficients are not statistically significant—reflecting the reduced sample size in these subgroup analyses—the signs of the point estimates are consistent with the main heterogeneity analyses presented in Section 7.2 and 7.3. Specifically, the first-row coefficients indicate that the wage effects for women relative to men are 0.4% and 0.66% higher in municipalities without the distribution of tax catalogs and in industries with a high baseline wage premium gap, respectively. While not statistically significant, these estimates are close in magnitude to the main effect of 0.52% reported in Section 7.4.1, lending further support to the interpretation that transparency contributed to the reduction in gender wage gaps. The second-row coefficients show that the average relative wage growth for female workers is 0.44% in municipalities with a tax catalog (Column (1)) and 0.37% in industries with a low baseline wage premium gap (Column (2)).

Table C1: Heterogeneous Effects of Income Transparency on Gender Wage Gaps

Dep. Var: Log(Wage)	Gender Effects	
	No Tax Catalog	High Gap Industry
	(1)	(2)
After x Female x 1{c=2000} $\cdot$ het_i	0.0040	0.0066
	(0.0098)	(0.0044)
After x Female x 1{c=2000}	0.0044	0.0037
	(0.0086)	(0.0025)
Observations	381,364	1,776,352

Note: Significance levels: * 10%, ** 5%, ***1%. The models include 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 municipality with prior printed versions 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)).