Pay Transparency and Mental Health*

Morten Bennedsen† Daniela Scur‡ Elena Simintzi§
Margarita Tsoutsoura¶ Daniel Wolfenzon‖

This draft: April 2024

Working paper: comments welcome

– Latest version here –

Abstract

This paper explores the impact of pay transparency on worker well-being. We exploit a pay transparency legislation in Denmark that requires firms with more than 35 employees to disclose pay information by gender. Using detailed administrative data on employee and mental health prescription drugs, we document the effect of the transparency policy. Our results suggest that the legislation led to a short-run decline in the anti-depressant use for women in affected firms. We consider two competing mechanisms behind this result: the potentially negative impact of horizontal (peer) comparisons and the potentially positive impact of organizational changes leading to the reduction of pay disparity. Taken together, our results are not consistent with a horizontal wage comparison mechanism and the evidence points to the second mechanism being the dominant force.

Keywords: Mental Health, Gender Pay Gap, Transparency

*We thank Ricardo Perez-Truglia, Zoe Cullen, John Van Reenen, Kathryn Shaw, Chris Rider for helpful discussions and suggestions. We thank seminar participants at the Empirical Management Conference 2022, Cornell, EquifFirm Workshop 2023 Bergen, Stockholm School of Economics, U Copenhagen, NOVA and Georgetown for excellent comments and discussions. Carolina Braga and Maithili Modi provided excellent research assistance.

†University of Copenhagen and INSEAD, ECGI, CEPR and DFI. E-mail: mobe@econ.ku.dk.
‡Corresponding author: Cornell University, Dyson School. E-mail: dscur@cornell.edu.
§University of North Carolina at Chapel Hill. E-mail: elena_simintzi@kenan-flagler.unc.edu.
¶Olin Business School, Washington University in St. Louis. E-mail: tsoutsoura@wustl.edu.
‖Columbia Business School. E-mail: dw2382@gsb.columbia.edu.

1 Introduction

Gender pay gaps are prevalent in the economy. Governments often attempt to nudge firms to reduce these gaps by introducing pay transparency mandates. Despite the prevalence of those mandates, our understanding of their effect on workers has been limited to studying their effect on wages, where the impact has been generally to help decrease in gender pay gaps (Bennedsen et al.; 2022; Passaro et al.; 2023). Pay transparency may also have important implications on employees' well-being, though the direction of the effect is not clear ex-ante. Transparency can facilitate income comparisons, which may have a negative effect on job satisfaction (e.g. Card et al.; 2012) as differences in peers' salaries are often considered arbitrary and unfair. But as policies often lead to changes in pay, these and other positive firm responses might generate a correspondingly positive effect on employee well-being.

In this paper, we explore the impact of gender pay gap transparency legislation on worker well-being using detailed administrative data from Denmark. We exploit a 2006 law that required firms with more than 35 employees to report salary statistics by gender and occupation, highlighting the within-firm gender pay gap. A key empirical challenge in studying well-being is having a measure that can compare employees' well-being both accurately and consistently across firms and over time (c.f Andreoni et al.; 2024). We overcome this challenge by using mental health drug prescription records for all employees in Denmark; this includes the population of prescriptions, as Denmark has a universal healthcare system. We use information on the type and dosage of drugs used to treat mental health conditions – henceforth, anti-depressant (AD) drugs – and assume that an increase in dosage over time is a proxy for worsening mental health conditions and vice-versa.[1] To examine the effect of transparency in anti-depressant use, we use a difference-in-differences (DiD) strategy in a narrow window around the 35 employee cutoff, defining firms with 35 to 50 employees as "treated" and firms with 20 to 35 employees as "control".

We find a significant impact of the pay transparency law on workers' anti-depressant use, especially for female workers. Anti-depressant use by female workers in treated firms was significantly lower relative to those in the control firms. The effect is driven by female workers with below-median wages, those working in firms with high wage gaps in the pre-period, and firms that had at least one female manager. The effect is, however, short-lived. There is an immediate reduction in year 1 that persists in year 2, but use reverts to the mean by year 3.

[1] We further verify that our measure is indicative of employees' well-being by matching drug use data to a contemporaneous workplace survey data at the worker level, though this survey is not available for the same time period as the policy change.

Quarterly data shows a gradual decrease followed by a gradual increase after about 5 to 6 quarters. This is consistent with the economics literature of the effect of life events on well-being, which tends to be short-lived (Oswald and Powdthavee; 2008), as well as the behavioral literature on the phenomenon of the "hedonic treadmill" (Brickman and Campbell; 1971). Antidepressant usage among male workers, in contrast, increases slightly. The granularity of our data allows us to control for time-invariant characteristics by including individual by firm fixed effects, so we can recover a causal effect of the law keeping the person-firm match constant, as well as time-varying individual characteristics such as worker age and experience. Although we consider a narrow window of firm sizes around the 35 employee cutoff, we additionally show our results are robust when controlling for a measure of firm size.

We attribute this short term reduction in anti-depressant use by female workers to the successful reduction in within-firm wage inequality, as well as a set of accompanying organizational changes that led to a more equitable work environment. We do not find evidence of horizontal peer comparisons negatively impacting morale (as in Card et al. (2012) and Cullen and Perez-Truglia (2022)); likely because firms seem to have addressed the inequality relatively quickly, and the comparisons were not direct, personal comparisons but rather more general, group-based comparisons. Our findings are also not a result of a direct effect of the law on individual wages; the gap grew smaller due to slower male wage growth relative to no impact on female wages (Bennedsen et al.; 2022). As female wages did not rise in absolute terms, it was the reduction in *relative* wage gaps that seems to have mattered for employee well-being. Further, the transparency law also seems to have triggered broader positive organizational changes that led to a more equitable work environment. Beyond the reduction to the gender pay gap, treated firms also increased their rates of hiring and promotion of women, including to managerial positions.

Our effect is economically important. To fix ideas, we examine the effect of other life events on individuals' anti-depressant use. The magnitude of the effect of transparency is, for example, about 20% that of divorce and very small relative to losing a spouse. Importantly, the short-lived effect on well-being is not unique to pay transparency legislation. We show that even more extreme life events, such as losing a spouse, have a temporary effect on anti-depressant use as well.

Our paper contributes to the literature on the impact of firm organization and pay practices on worker outcomes. One strand of this literature considers pay inequality specifically, generally finding satisfaction drops following the availability of direct horizontal comparisons with colleagues (Breza et al.; 2018; Card et al.; 2012; Cullen and Perez-Truglia; 2022).

Another strand focuses on organizational changes, and finds that organizational change is generally stressful for employees (Dahl; 2011; Dahl and Pierce; 2019) and employee well-being is key for aggregate healthcare costs (Goh et al.; 2016), but positive organizational changes (such as "wellness programs") can have positive impacts on workers health (Gubler et al.; 2018). We contribute to this literature by linking the peer comparison and organizational changes channels, showing that a group-based pay transparency policy yielded positive organizational changes and mental health outcomes for women, albeit short-lived, and minimal negative impact on men.

A related part of the pay transparency literature examines how it affects wage setting. Cullen and Pakzad-Hurson (2021) explores how transparency affects wage setting in organizations, and argue that higher transparency results in more equal but lower wages and higher profits for employers. Mas (2017) shows that top earners in municipal jobs experience a drop in wages following the public disclosure of wages, which he argues is due primarily to public aversion to visibly exorbitant salaries. Bennedsen et al. (2022) and Baker et al. (2019) show pay transparency reduced the gender pay gap, with this effect driven primarily by lower growth of male wages. We contribute to this literature by showing that pay transparency can impact employee outcomes beyond wages, and specifically, their mental health.

2 Conceptual framework and context

2.1 Danish Context: The Pay Transparency Law

The law requiring gender pay gap disclosure was passed on June 9, 2006 (Act no. 562), mandating disclosure of gender-based disaggregated statistics with the purpose of informing employees of within firm gender wage gaps. The law requires firms with more than 35 employees to report annually gender-segregated wage statistics for occupations (defined by six-digit DISCO code) with at least 20 employees of each gender.[2] This information, however, was not required to be disclosed to the wider public but only to the employee representatives who then convey the information the other employees.

In the empirical analysis we focus on the criterion that firms with 35 or more employee are required to report wage statistics and do not consider the second criterion that firms report only for occupations with more than 20 employees of each gender. We do this because most firms followed this criterion only. First, we investigated how the law was communicated to firms and find that the description of the law by the EU and the International Labor

[2]The requirement does not extend to companies in the fields of farming, gardening, forestry, and fishery.

Organization (ILO) only mentions the first criterion.[3] Second and consistent with the way the law was communicated, many firms that met the first but not the second criterion reported wage statistics. According to a report by the Confederation of Danish Employers (DA), the largest employer association covering about a third of Danish employees, 35% of member firms that report gender-disaggregated wage statistics do not satisfy the second criterion but, all of them have more than 35 employees. Moreover, Bennedsen et al. (2022) provides a examples of implementation challenges that led firms to not follow the DISCO requirement as well as additional details about the law and its implementation.

The passage of the law was not anticipated. While the government had promised to enact a transparency law during the campaign for the snap election in February 2005, there was little discussion about fulfillment of this promise after winning the election.[4] This neglect was not surprising given that the same administration had tabled a similar proposal years earlier. However, eight months after re-election, in October 2005, the government announced that it would introduce a bill to amend the Equal Pay Act. The Ministry of Economics introduced the bill to Parliament in December 2005. The proposal was adopted on June 2006, and the new provisions came into force in January 2007.

2.2 How Pay Transparency Impacts Mental Health: Conceptual Framework

In this section we briefly outline a conceptual framework to serve as a guide to our empirical analysis. We propose that pay transparency at the firm level can impact worker wellbeing via two key channels: a response to peer comparisons, and a response to organizational changes within the firm (if any). Much of the literature on pay transparency has focused on the former, and we propose that while the very short-term impact is likely due to peer comparisons, the medium to longer term impact is linked to the latter.

The literature on peer comparisons has primarily shown the impact of direct comparisons; that is, workers were able to directly see their peers' salaries (e.g. Breza et al.; 2018; Card et al.; 2012). Such directives, however, are only enforceable in public sector organizations

[3]The European Commission's Directorate General for Internal Policies issued a report entitled Gender Equality in Denmark that describes the law as follows: "Since 2007, companies with 35 employees or more should carry out gender disaggregated pay statistics and elaborate status reports on the efforts to promote equal pay in the workplace." (EC; 2015). ILO's description of the law reads as follows: "Employers employing 35 or more workers are required to prepare annually gender-disaggregated statistics or, alternatively, an equal pay report and action plan.".

[4]The snap election was announced on January 18, one year earlier than scheduled, and took place on February 8.

where the government is the employer. For private firms, unless firms voluntarily take up such policies, directives such as the one in Denmark are more feasible to enforce, where firms are compelled to publish aggregate gender statistics to enable group-wise wage comparisons rather than reveal individuals personal salary information. Indeed, this is the most common type of policy in the OECD countries (OECD; 2023). Under this type of policy, the information workers are exposed to allows for an indirect comparison and triggers responses based on preferences for equity, or inequality aversion. Similar to direct comparisons, we expect that the direction of the impact mirrors the direction of the inequality: if one's own wage is lower (or part of the lower-wage group) relative to peers, that leads to a decrease in wellbeing. In this paper we focus on the indirect, group-based type of peer comparison and expect that firms with higher wage gaps would have a negative impact in this component wellbeing.

There is less evidence on how the second component of worker wellbeing, related to the internal organization of their firms, is affected by pay transparency. While organizational change is generally stressful for employees (Dahl; 2011; Dahl and Pierce; 2019), positive organizational changes can have positive impacts on workers health (Gubler et al.; 2018). In the case of pay transparency, equilibrium wages tend to fall along with inequality following the implementation of such policies (Baker et al.; 2019; Bennedsen et al.; 2022; Cullen and Pakzad-Hurson; 2021)). We propose that this organizational component has two parts: positive and negative changes. Specifically, if there are positive organizational changes (expected or actual), these would have a positive impact on worker well-being (and vice-versa). "Positive" is, of course, relative. While increased fairness can be seen as a generally positive change and unambiguously better for a female worker, it could be either positive or negative for a male worker whose salary decreased to achieve this fairness. An individual may experience an impact from both positive and negative organizational changes, as well as from peer comparisons.

Relative impact of components. While our data does not allow us to decompose the effect, it does allow for inference based on the relative sign of the components. That is, we expect that, for female workers, the impact of peer comparison is at least non-positive (though likely negative) and that the impact of female-worker friendly organizational changes is positive. If the impact of peer comparison is larger than that of organizational changes, then we would see an aggregate negative impact on wellbeing. Conversely, if the impact of organizational changes is larger, we would see an aggregate positive impact on wellbeing.

For male workers, we also expect the peer comparison effect to be, similar to the women, at least non-positive. The impact of female-worker friendly organizational changes is less

clear ex-ante. If the male worker has pro-social beliefs and ranks fairness highly he might experience a positive wellbeing impact from seeing greater fairness in his organization. If his own standing in the firm declines, however, he might experience a negative impact. The balance of these, alongside the peer comparison impact leads to the aggregate impact on wellbeing.

Timing. In psychology, the concept of the "hedonic treadmill" suggests that people have a baseline level of wellbeing (a "set point") and that while this level of happiness may rise or fall in response to life events, it eventually reverts to its "set point mean" (Brickman and Campbell; 1971). As this is the case with even major life events, such as the loss of a job, the birth of a child or the loss of a loved one, we similarly expect the impact here to be short-lived.

In our setting, the policy mandates that firms disclose the gender pay gap annually, usually towards the end of the year. The firms we focus on are also relatively small, such that workers would see organizational changes being implemented. That is, while hiring a few female workers or promoting a female worker to manager may not be visible to everyone in a 100-person firm, it is likely to be so in a 20 to 50-person firm.

3 Data and context

3.1 Data Sources and Sample Description

We combine three primary data sources: data on individual workers' prescription use from the Danish Health Authority (SSI), employee work history from the Integrated Database for Labor Market Research (IDA) from Statistics Denmark, and firm data from the Danish Business Register (Experian).

3.1.1 Prescription Data: Danish Health Authority

Denmark's Health Authority compiles comprehensive dataset (SSI) on the distribution of pharmaceuticals throughout the country. This dataset started in 1994 and includes all records of drugs dispensed via prescription at a pharmacy and hospitals to individuals. Each record includes the individual's personal identification number (CPR number), drug name, date of purchase or date of delivery, the drug's Anatomical Therapeutic Chemical code (ATC) as defined by the World Health Organization (WHO), the strength of the active ingredient(s) included in the drug (for example, "500mg"), and number of units delivered.

Our main proxy of workers' mental health captures the intensive margin: the total sum of the volume of anti-depressant medication an individual worker is prescribed within a year. We also use the measure at the extensive margin, namely whether the volume of anti-depressant use was larger than zero anytime during the year. We identify antidepressants by their four-digit ATC code. Specifically we use the following categories of mental-health related drugs: antidepressants (N06A), anxiolytics (N05B), sedatives (N05C), and antipsychotics (N05A). We refer to these drugs used to treat mental health conditions, broadly, as anti-depressants (AD). Since the various antidepressant medications have different active ingredients and thus their quantity cannot be directly compared, we convert the amount of active ingredient into a standardized measure. For each drug, the WHO provides the Defined Daily Dosage (DDD), which is the average maintenance dose per day for a drug used for its main indication in adults. Our detailed information on prescriptions, allows us to define, for each individual, their annual antidepressant dosage in number of DDDs.[5]

3.1.2 Employee data: Integrated Database for Labor Market Research

We use the matched employer-employee dataset from the Integrated Database for Labor Market Research (IDA) from Statistics Denmark. This dataset includes the employer identification number (CVR), the employee identification number (CPR), compensation, demographic information (age, gender, education) and occupation (job type and hierarchy). This dataset covers all employees and firms in Denmark and hence allows us to follow employees over time and across firms.

3.1.3 Firm data: Danish Business Register

Firm-level data comes from the Danish Business Register. This dataset covers all firms incorporated in Denmark and includes the information firms are required to file with the Ministry of Economics and Business Affairs, including the value of total assets, number of employees, and revenues. Even though most firms in this dataset are privately held, external accountants audit firm financial information in compliance with Danish corporate law.

[5]Figure A1 shows the number of DDDs of mental health drugs that employees in our sample take by categories. Antidepressants are the most commonly prescribed drug, with sedatives (often prescribed for sleep disorders and some types of anxiety) and anxiolytics (often prescribed for anxiety) as the next two most common prescription types.

3.2 Sample construction

Our starting point is the universe of the limited liability companies in Denmark and their employees. We focus the analysis around the 2006 reform, covering the period 2003 to 2008. The empirical analysis involves comparing firms affected by the law (treated firms) to firms that are not affected (control firms). To keep the treated group as comparable as possible to the control group, our sample includes firms within a narrow band around 35 employees, specifically from 10 to 50 employees. We exclude firms in industries not affected by the law (those in farming, gardening, forestry and fishery) and firms with missing financial information.

We focus on full-time workers and exclude CEOs. Most importantly, our preferred sample includes employees with no pre-existing mental health conditions as it avoids including workers with long-term chronic conditions, who are unlikely to be affected by the transparency. Specifically, our sample includes only workers who had not used anti-depressants from the time our data starts, 1995, to 2002, before our window of analysis. Our final analysis dataset includes almost 187,000 person-year observations across over 54,000 unique individuals in over 3,200 unique firms.

3.3 Summary statistics

Table A1 reports summary statistics for treated and control firm in our sample. Treated firms are those that employ between 35 and 50 employees before the 2006 law and control firms are those that employ 20 to 34 employees in the same period. Panel A presents information on firm characteristics, while Panel B present employee-level information.

By design, treated companies are larger in terms of sales and number of employees. However, the gender composition and education level of their respective employees is very similar. Albeit the difference in age is statistically significant, the magnitude is quite small at only 0.4 years. The anti-depressant use in the pre-period is, however, not statistically different between employees in treated versus control firms.

4 The Effects of Transparency Legislation

4.1 Pay transparency and mental health

We first consider the effect of the transparency law on antidepressant usage of employees using a differences-in-differences (DiD) framework with data at the annual level.

We estimate the following OLS regression for individual i in firm j and year t, separately for male and female employees:

$$ln(AD\_volume)_{ijt} = \alpha_{ij} + \alpha_t + \beta_1 I(Treated_{ij} \times Post_t) + \varepsilon_{ijt}$$

where $Post_t$ is an indicator that takes a value of one for 2006, 2007, and 2008 and a value of zero for 2003, 2004 and 2005, $Treated$ is an indicator that takes a value of one for firms that employ 35 to 50 employees prior to the introduction of the law and zero for firms that employ 20 to 34 employees. α_t are year fixed effects that absorb aggregate macroeconomic shocks. We also include individual $\times$ firm fixed effects, α_{ij}, to control for time-invariant individual characteristics, time-invariant firm characteristics, and the match between firms and employees.[6] With this specification, we compare the *same employee at the same firm* before and after the regulation, avoiding composition effects. We include controls for firm and employee time-varying characteristics (sales, age, and experience) as a robustness check. Standard errors are clustered at the individual and firm level.

We also test whether the difference between female and male employees is statistically significant using a triple difference (DiDiD) specification:

$$\begin{aligned} ln(AD\_volume)_{ijt} &= \alpha_{ij} + \alpha_t + \beta_1 I(Treated_{ij} \times Post_t) \\ &+ \beta_2 I(Post_t \times Female_i) \\ &+ \delta\, I(Treated_{ij} \times Post_t \times Female_i) + \varepsilon_{ijt}, \end{aligned} \tag{1}$$

where the notation is the same as in Equation 1.[7] δ is the triple differences coefficient estimate that tests whether the effect of the transparency law is different between female and male employees. In this specification, we can further include firm $\times$ year firm fixed effects, α_{jt}, which controls for any firm level shocks. Standard errors are clustered at the individual and the firm level. Table 1 reports our main results.

4.1.1 Intensive margin: changes in anti-depressant volume of use

Column (1) compares the change in anti-depressant use in treated relative to control firms for all female employees and Column (2) for male employees. We observe the transparency law

[6] Note that we omit $Treated_j$ and $Post_t$ as these coefficients are absorbed by the fixed effects.

[7] Note that we omit $Female_i$, $Treated_j$, $Post_t$, and $Treated_{ij} \times Female_i$ as these coefficients are absorbed by the fixed effects.

had a differential impact in men and women: we find that female employees' anti-depressant use in treated firms 3.7 percentage points lower relative to control firms. Male employees' use in treated firms, however, was slightly higher, by 1.6 percentage points. Column (3) reports the result of estimating the triple difference model presented in Equation 1. The triple-difference coefficient $Treated \times Post \times Female$ compares the effect of the law on female vs male employees. Female antidepressant usage growth is 5.2 percentage points lower than male usage growth, and the effect is statistically significant. In Column (4), we additionally control for interacted firm $\times$ year-fixed effects that absorb any time-varying changes at the firm level.[8] The triple-difference coefficient in this more stringent specification is stronger both in terms of magnitude and significance.

To understand the dynamics of the effect, we report the estimation of a linear panel model with dynamic effects (Figure 1) at the quarterly level, with Q1 of 2005 as the base year-quarter. We show the dynamic plots of the event study coefficients and their associated confidence intervals with the log of anti-depressant use as the outcome variable for female workers (green solid triangles) and male workers (red hollow circles). In Panel 1a there are no differential trends in antidepressant use for both female and male employees and, consistent with our main results in Table 1, women in treated firms experience a significant drop in the rate of their anti-depressant usage after the introduction of the transparency law. Specifically, we observe a gradual decline in the growth rate of antidepressant use in each of the quarters of 2006 and 2007. However, the effect is relatively short-lived and it dissipates by 2008, which explains the relatively smaller coefficients in the differences in differences average.

4.1.2 Extensive margin: changes in incidence of use

Columns (5) through (8) of Table 1 repeat the same set of specifications but for the extensive margin: whether a worker takes any anti-depressant in a particular year (quarter). Again there is a significant difference between male and female workers (Columns 7 and 8), but it is driven primarily by a slightly higher incidence of use among men relative to no change for female workers. This suggests much of the impact in the aggregate volume slowdown in consumption is driven by relative reductions in use by individual workers, rather than just fewer female workers using anti-depressants in general. There is also a significant impact on the incidence of use in Panel (B), though even more short-lived.

Looking at the dynamic effect, Figure 1b shows no pre-trend in 2004 and 2005, though a small bump in 2003. All quarters in 2006 and 2007 have a similar lower share of users in

[8] Firm $\times$ year fixed effects subsume the coefficient on $TreatedxPost$.

the female worker sample and slightly higher share of users in the male sample, though the female workers revert to the mean in all quarters of 2008. The slightly lower trend (albeit not significant) in 2003 along with the mean reversion in 2008 is likely the reason we do not see any significant difference in the incidence of use on average in the diff-in-diff annual-based analysis.

4.1.3 Timing as a mediator: mean reversion

There are two potential explanations for this mean-reversion three years after the law was introduced. First, this is consistent with a hedonic adaptation (or, "hedonic treadmill") documented in the literature. For example, learning about relative income as in Cullen and Perez-Truglia (2022), positive unexpected effects on employee compensation ('gift-giving') as in Gneezy and List (2006) or firm acquisitions as in Silva et al. (2021) have a temporary effect on individuals' satisfaction. Even important life-changing events, such as divorce or a spouse's death have large effects at the time of the event but they tend to subside by the third year (Figure 5). An alternative explanation is that anticipatory effects of expected improvements could have not materialized, thereby resulting in a reversion to the prior status quo.

4.2 Placebo tests, alternative specifications and benchmarking

In this section we consider whether the relationships we see could be linked to trends around more general drug use, other contemporaneous events that affect firms more generally, and whether a counterfactual of smaller firms is appropriate. We show that workers in treated and control firms have similar use of other, non-mental health related types of drugs, and alternative specifications reinforce our identification strategy.

4.2.1 Alternative drugs

An alternative concern could be that health coverage changes for firms above 35 employees affecting their employees' general health, including their anti-depressant use. However, unlike the U.S. where health coverage varies by firm, healthcare in Denmark is government-funded and does not depend on employment. To mitigate this concern, we estimate our main model using as dependent variables alternative drugs (cholesterol and diabetes). Figure 3 shows the event study plots for female and male employees. These plots indicate that there is no effect on other medications for treated employees either before or after the transparency law.

4.2.2 Alternative cut-offs

A possible concern is that a contemporaneous event is causing our results. To address this concern, we perform several placebo tests where the treatment threshold takes each of the values between 15 and 100 employees, excluding the range of 20-50 employees. We plot the average placebo coefficients in Figure 2, in gray hollow markers, together with our "real" effect coefficients in solid green and red markers (respectively). The average placebo coefficients are close to zero both before and after the law in all cases, suggesting that the observed effect is specific to the 35-employee threshold.

4.2.3 Alternative specification: Regression-discontinuity design

An additional identification concern with our empirical strategy might be that treated and control groups are not good counterfactuals for each other, given that the control group includes smaller firms by definition. To mitigate this concern, we implement a regression discontinuity approach (RDD), given that the law applies only for firms above the 35 employees threshold. For each worker, we compute the difference between their average antidepressant usage in 2006-2007 and 2004-2005 (in logs). For this analysis, we use a narrow window around the event (two years before and two years after), effectively excluding the mean-reversion year (2008) as shown in Figure 1. Figure 4 shows the results of fitting a third-degree polynomial on each side of the discontinuity.[9] In Panel (A) we observe a clear discontinuous drop at the 35-employee threshold for female employees (significant at the 1% level), while in Panel (B) we find no significant effect for male employees. These results are consistent with those in Table 1.

4.2.4 Benchmarking the effect: comparison with Other Life Events

To put the magnitude of the effect into context, we compare the change in AD use to the magnitude of other life events that are potentially impactful for the mental health of individuals.

We first consider the change in antidepressant use around the death of a spouse, for women and men separately. For these tests we use the whole Danish population excluding those with pre-existing mental health issues. We plot the event study coefficients around the year of the death of the spouse (t=0). The green horizontal line in the plots of Figure 5

[9] In Appendix Figure A2 we show a McCrary test indicating the distribution of number of employees is continuous around the threshold (McCrary; 2008).

shows the absolute value of our baseline magnitude. Panel A shows a sharp increase of anti-depressant use for women following a partner death, twenty times larger than to our baseline estimate (and, of course, in the opposite direction). Notably the effect dissipates within a couple of years and reverts to the mean by year 3. Panel B shows there is an impact in male anti-depressant use following a partner death, though their response is half relative to that of the women. Similarly, the economic magnitude gradually converges to zero by year four.

In Panels C and D we repeat the same exercise for divorces as the main event. As expected the magnitude of the effect is lower than death for both women (Panel C) and men (Panel D), and between 5 to 6 times higher relative to the magnitude of our baseline. We observe that for divorces the effect starts in t=-1 because in Denmark there is a mandatory two year period of living separately before a legal divorce can be issued. We also see a similar pattern here where the use of anti-depressants reverts to the mean two years after the event.

5 Mechanisms: why do female workers take less anti-depressants in the short term?

In this section we present evidence that is consistent with "positive" organizational changes driving lower antidepressant use female workers. We define "low wage workers" and "high wage workers" as *workers* below and above median in their within-firm wage distribution. We define "low pay gap" and "high pay gap" as *firms* that fall below or above the median pay gap in 2005, the year before the reform.

5.1 Wages

Prior literature consistently documents that wage transparency in the private sector leads to a deterioration of employee morale and loss of productivity (Breza et al.; 2018; Cullen and Perez-Truglia; 2022), but if pay transparency also leads to changes in the organizational structure of the firm that can be viewed as positive these could have a counterbalancing effect. As detailed in our conceptual framework, the balance of these factors would predict a negative effect on mental health only employees who perceive changes geared towards equality as negatively affecting them, and potentially a positive effect on employees who perceive equality as positive and/or who are personally positively affected by such changes. As our data allows us to look beyond the very short term impact of immediate peer comparisons, we instead capture the new equilibrium resulting from learning about inequality combined

with actions to address it.

Indeed, the Danish transparency legislation was aimed at informing employees of gender wage gaps and seems to have led firms to make organizational changes to address gender inequality. As Bennedsen et al. (2022) showed, the pay gap reduction was driven by slower growth in male wages. They also show that the pay gap reduction was largest at the bottom of the income distribution, with the ratio of male to female wages decreasing in the 10th and 50th percentiles of the wage distribution but not on the 90th. In Table 3 we expand this result and show the change in wages for female workers in Panel A and for male workers in Panel B, separating workers by high and low wage, and those who work in high and low pay gap firms.[10] We confirm that, indeed, female wages did not grow differentially in treated versus control firms in any of the sub-samples of interest. Male wages did, however, grow slower in all these sub-samples. We also explicitly rule out this channel in our data. We estimate our main tests including employee wages as a control and neither the magnitude nor significance of the main coefficient estimates on antidepressant use change substantially.

5.2 Organizational changes

While wage changes are not likely behind the decrease in female workers' anti-depressant use, organizational changes promoting a more equitable workplace seem to be the primary drivers in the improvements in employee mental health, however short-lived. We document a greater impact for female workers in firms where there was more room for change, that is, lower-wage workers and firms with larger pre-period pay gaps, as well as where changes likely promoted a female-friendly work environment with female workers represented in managerial positions (Tate and Yang; 2015).

Bennedsen et al. (2022) documented a significant difference in male and female promotions, with more female workers being promoted in treated firms relative to control firms. We expand on this by considering differences in promotion rates of employees in high pay gap firms and low pay gap firms. Table 4 shows that high pay gap firms had a higher number of promotions in general, but they were distributed to both male and female workers (Columns 1 and 2) at quite similar rates relative to control firms. The triple difference in Column (3) is, thus, not significantly different from zero. Low pay gap firms, however, do promote more female workers relative to male workers in treated vs control firms. At the firm-level, we instead consider whether the share of female managers has changed and whether there is at least one female manager in the firm. While there is no difference in the share of female

[10] We define high and low as above or below median for each variable set.

managers, high pay gap firms are more likely to have at least one female manager after the reform.

Further as in Bennedsen et al. (2022), we find that treated firms hired more female workers post-reform and did not lose them. Considering these labor flows across high pay gap and low pay gap firms, both have similar patterns for female workers: significantly higher share of joiners and no difference for leavers. Men, however, make up a significantly lower share of joiners in high pay gap firms – suggesting high pay gap firms slowed their hiring of men – but also made up a smaller share of leavers.

5.3 Heterogeneous impact on anti-depressant use

We estimate our benchmark model for the sample of male and female workers across the key firm heterogeneity factors above. Figure 6 presents the results of the event study with quarterly data. Panels (A) and (B) show the changes in anti-depressant use for high and low wage workers, and the results indicate a larger antidepressant reduction in the sample of women at the bottom of the wage distribution rather than those at the top. Panel (C) and (D), in turn, consider the impact on female workers in "high pay gap" firms and "low pay gap" firms and shows a larger reduction in anti-depressant usage for women in the higher pay gap firms. We find no significant change for men in either group of firms or workers. This evidence is consistent with the effect being driven by female employees that likely expected to have greater gains from reduction in pay gaps.

In Panels (E) and (F) we examine the role of having at least one female manager in the *pre-period*. That is, as in Tate and Yang (2015), if female managers promote a more female-friendly workplace, female workers could be more confident in expecting changes in their firm. The results show that the anti-depressant use of female workers decreased relatively more in firms that had female managers in the pre-period, though this result is less clear at the quarterly level (it is significant in more aggregated annual and average iterations). These results suggest that female representation at the top of the organization can help push for positive change once an opportunity arises.

Taken together, these results suggest that actual or anticipatory organizational changes that are expected to benefit female employees are consistent with the positive mental health outcomes for female employees in treated firms.

6 Discussion

As mental health issues become more widespread and a growing concern in many countries, it is important to understand the potential contribution of workplace factors to people's well-being. We explored the impact of pay transparency legislation on worker well-being using detailed employee and medical prescription administrative data from Denmark using two empirical strategies. As documented elsewhere, the policy of pay transparency was effective in changing the structure of pay in Denmark. Our results suggest that it also led to a short-run decline in the relative growth rate of anti-depressant use for women in affected firms. As our measure of well-being is anti-depressant use, our results are likely an upper bound. Many employees suffering with stress or job-related anxiety may not resort to using anti-depressants as treatment in the short run or for "milder" cases. We consider two competing mechanisms behind this result: the potentially negative impact of horizontal (peer) comparisons and the potentially positive impact of organizational changes leading to the reduction of pay disparity. Taken together, our results are not consistent with a horizontal wage comparison mechanism and the evidence points to the second mechanism being the dominant force.

References

Andreoni, J., Bernheim, B. D. and Jia, T. (2024). Do people report happiness accurately?, *Working Paper 32208*, National Bureau of Economic Research.
URL: *http://www.nber.org/papers/w32208*

Baker, M., Halberstam, Y., Kroft, K., Mas, A. and Messacar, D. (2019). Pay transparency and the gender gap, *Working paper No. 25834*, National Bureau of Economic Research.

Bennedsen, M., Simintzi, E., Tsoutsoura, M. and Wolfenzon, D. (2022). Do firms respond to gender pay gap transparency?, *The Journal of Finance* **77**(4): 2051–2091.
URL: *https://onlinelibrary.wiley.com/doi/abs/10.1111/jofi.13136*

Breza, E., Kaur, S. and Shamdasani, Y. (2018). The morale effects of pay inequality, *Quarterly Journal of Economics* **133**(2): 611–663.

Brickman, P. and Campbell, D. T. (1971). Hedonic relativism and planning the good society, *in* M. H. Appley (ed.), *Adaptation-level theory*, Academic Press, New York, pp. 287–305.

Card, D., Mas, A., Moretti, E. and Saez, E. (2012). Inequality at work: The effect of peer salaries on job satisfaction, *American Economic Review* **102**(6): 2981–3003.

Cullen, Z. B. and Pakzad-Hurson, B. (2021). Equilibrium effects of pay transparency, *Working Paper 28903*, National Bureau of Economic Research.

Cullen, Z. and Perez-Truglia, R. (2022). How much does your boss make? the effects of salary comparisons, *Journal of Political Economy* **130**(3): 766–822.
URL: *https://doi.org/10.1086/717891*

Dahl, M. (2011). Organizational change and employee stress, *Management Science* **57**: 240–256.

Dahl, M. and Pierce, L. (2019). Pay-for-performance and employee mental health: Large sample evidence using employee prescription drug usage (forthcoming in academy of management discoveries).

EC (2015). The policy on gender equality in denmark - update, *Report*, European Commission.

Gneezy, U. and List, J. (2006). Putting behavioral economics to work: Testing for gift exchange in labor markets using field experiments, *Econometrica* **74**(5): 1365–1384.

Goh, J., Pfeffer, J. and Zenios, S. A. (2016). The Relationship Between Workplace Stressors and Mortality and Health Costs in the United States, *Management Science* **62**(2): 608–628.

Gubler, T., Larkin, I. and Pierce, L. (2018). Doing Well by Making Well: The Impact of Corporate Wellness Programs on Employee Productivity, *Management Science* **64**(11): 4967–4987.
URL: *https://ideas.repec.org/a/inm/ormnsc/v64y2018i11p4967-4987.html*

Mas, A. (2017). Does transparency lead to pay compression?, *Journal of Political Economy* **125**(5): 1683–1721.

McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test, *Journal of Econometrics* **142**(2): 698–714.

OECD (2023). Gender pay gap reporting and equal pay audits, *Policy brief*, OECD.

Oswald, A. and Powdthavee, N. (2008). Does happiness adapt? a longitudinal study of disability with implications for economists and judges, *Journal of Public Economics* **92**(5-6): 1061–1077.
URL: *https://EconPapers.repec.org/RePEc:eee:pubeco:v:92:y:2008:i:5-6:p:1061-1077*

Passaro, D. G., Kojima, F. and Pakzad-Hurson, B. (2023). Equal Pay for Similar Work, *Papers 2306.17111*, arXiv.org.
URL: *https://ideas.repec.org/p/arx/papers/2306.17111.html*

Silva, R., , Baghai, R. and Bos, M. (2021). How Do Acquisitions Affect the Mental Health of Employees?, *CEPR Discussion Papers 16657*, C.E.P.R. Discussion Papers.
URL: *https://ideas.repec.org/p/cpr/ceprdp/16657.html*

Tate, G. and Yang, L. (2015). Female leadership and gender equity: Evidence from plant closure, *Journal of Financial Economics* **117**(1): 77–97.

Tables and Figures

Figure 1: Effect of transparency policy on anti-depressant use

(a) Continuous ln(antidepressant use)

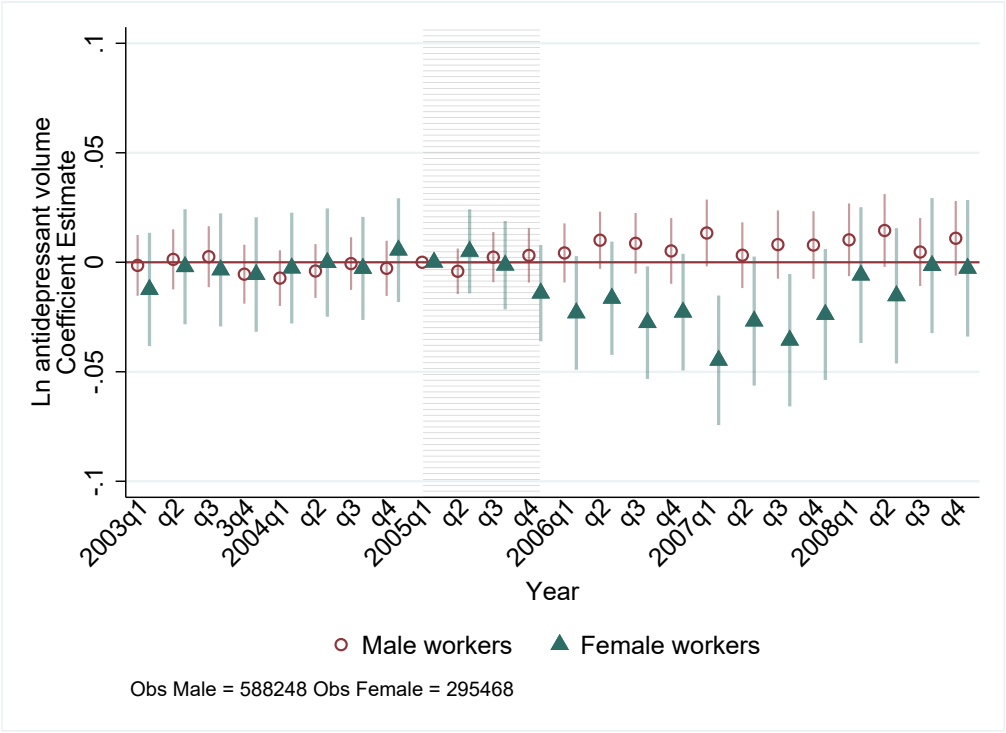

(b) Binary outcome (user = 1)

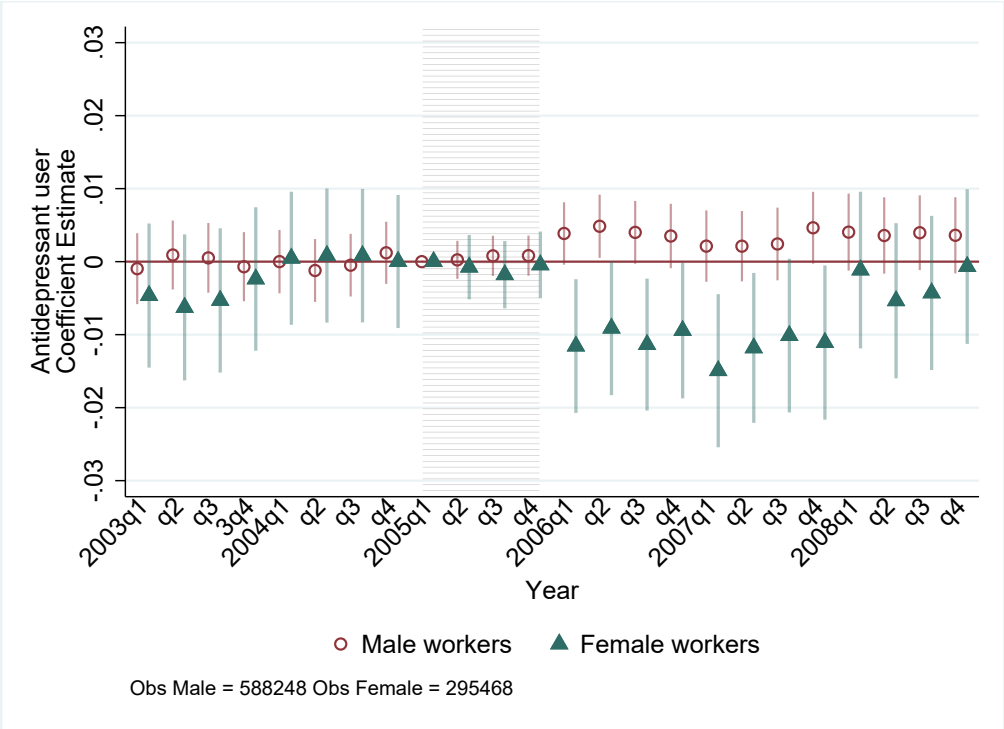

Note: Data from Danish statistics, quarterly use of antidepressants by workers between 2003 and 2008. This figure plots the coefficients of two event studies: Panel (A) uses the log of antidepressant use as the outcome variable, and Panel (B) uses a binary user = 1 / non-user = 0 dummy as the outcome variable. Hollow circles indicate male workers. Filled triangles indicate female workers.

Figure 2: Placebo test: different cut-off

(a) Different cut-off, female

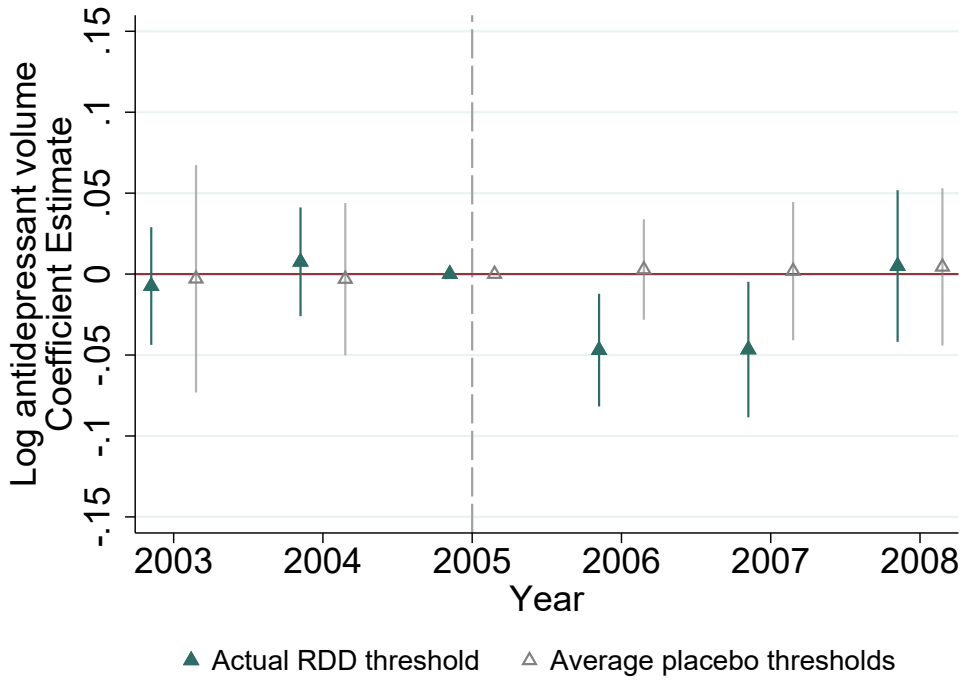

(b) Different cut-off, male

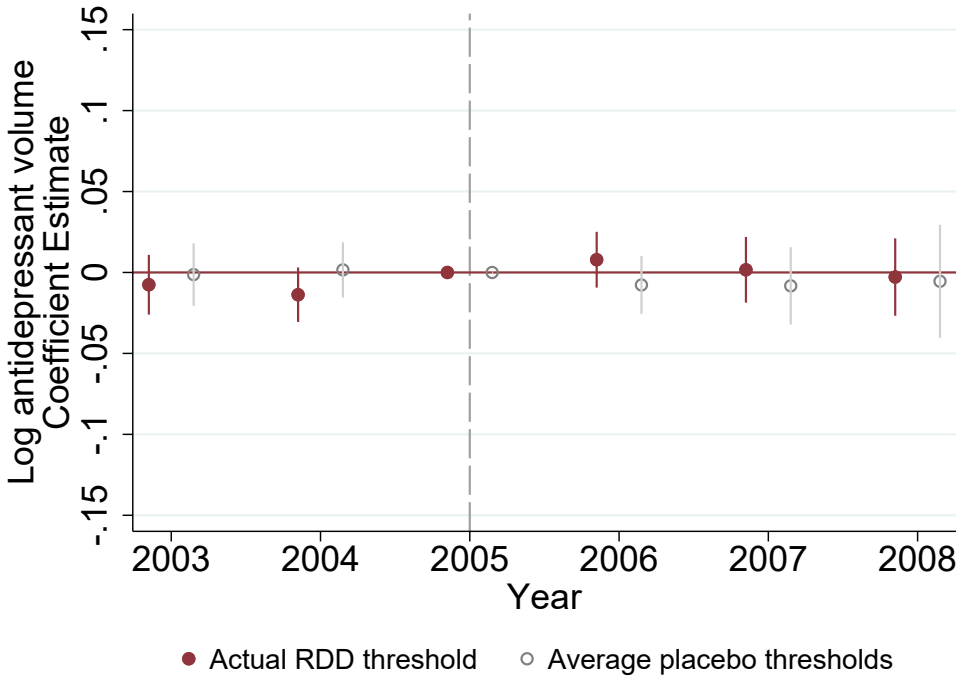

Note: Data from Danish statistics, annual use of antidepressants by workers between 2003 and 2008. For Panels A and B, we perform several placebo tests where the treatment threshold takes each of the values between 15 and 100 employees, excluding the range of 20-50 employees. We plot the **average placebo coefficients** together with our baseline coefficient estimates shown in Panels.

Figure 3: Placebo test: different drugs

(a) Cholesterol drugs

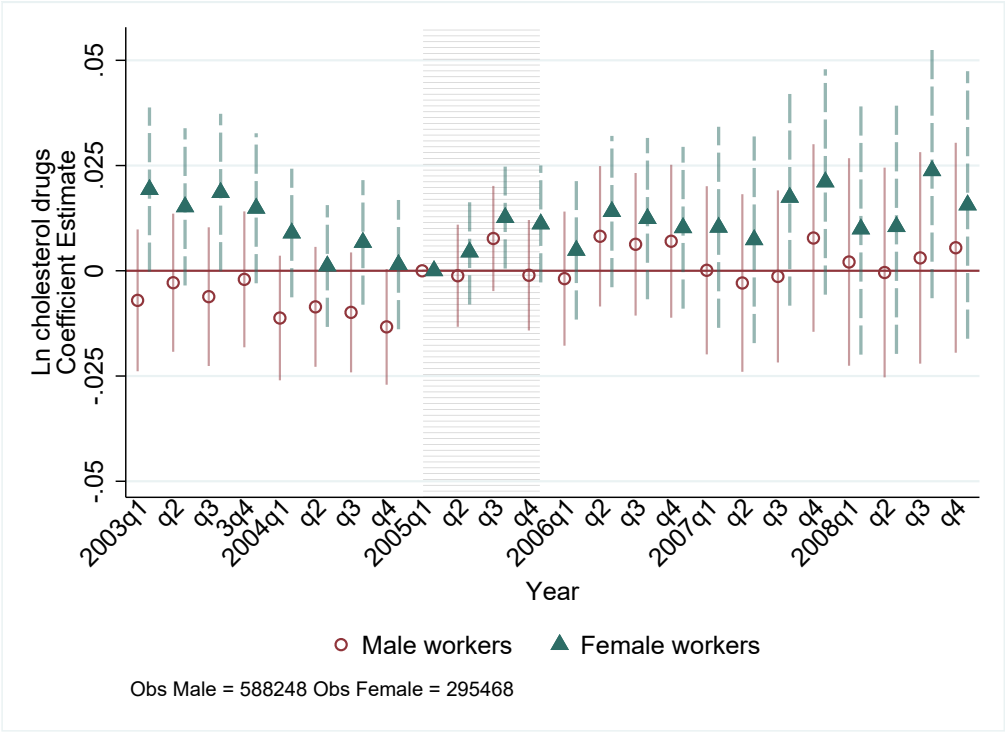

(b) Diabetes drugs

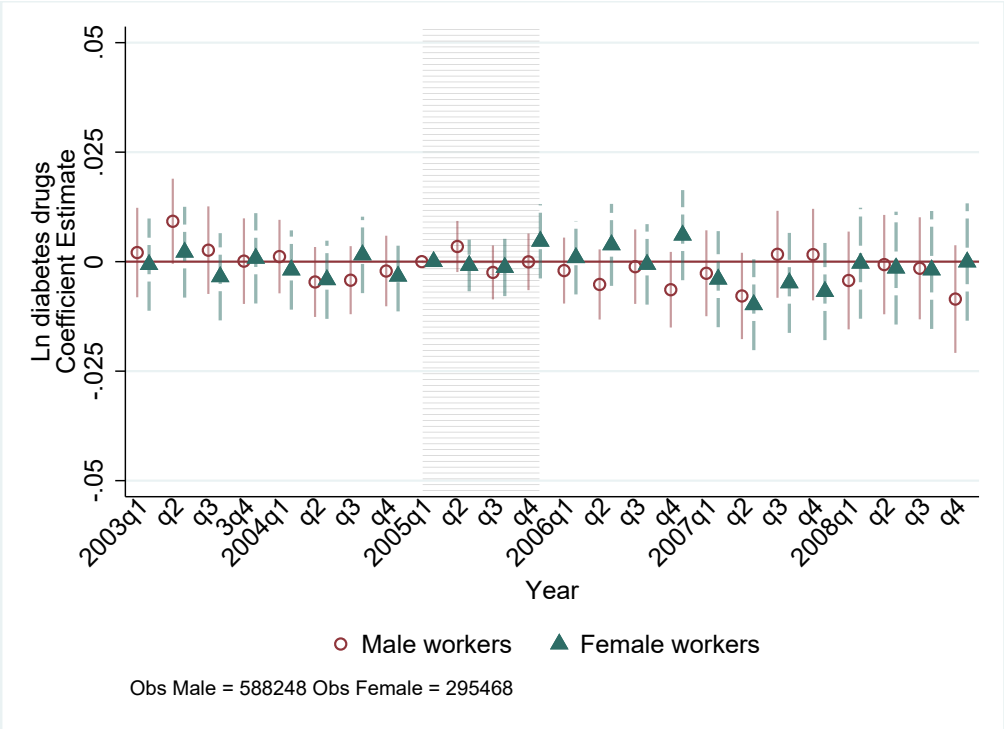

Note: Data from Danish statistics, quarterly use of cholesterol and diabetes drugs by workers between 2003 and 2008. This figure plots the event study of drug use for non-mental health drugs for the same sample as our preferred specifications. Panel (A) shows the results for cholesterol drugs and Panel (B) shows the results for diabetes drugs.

Figure 4: Alternative identification: Regression Discontinuity

Panel (A) Female workers

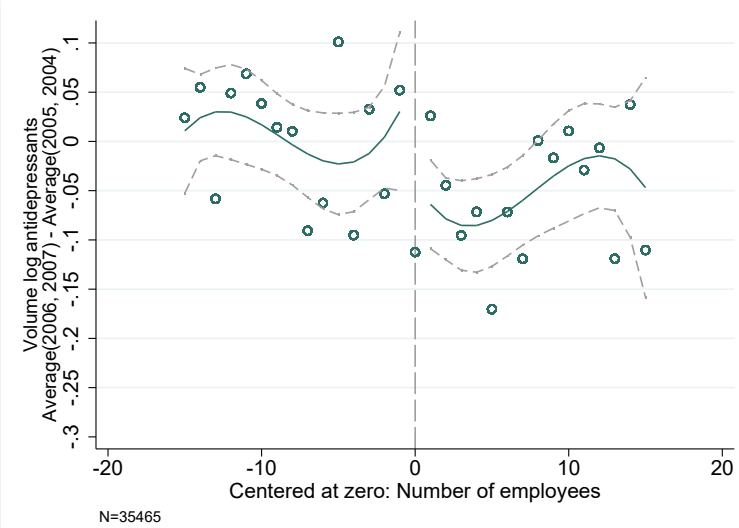

Panel (B) Male workers

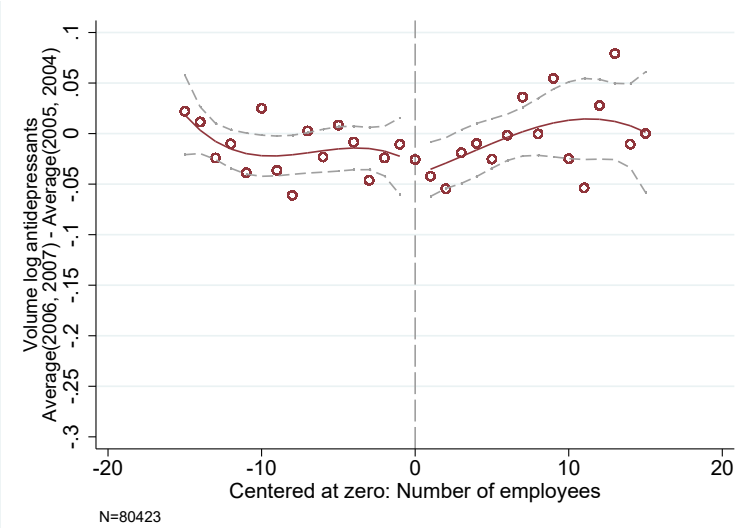

Note: Data from Danish statistics, annual use of antidepressants by workers between 2003 and 2008. Panels A and B plot results for log antidepressant volume for female and male workers. The vertical axis in both panels shows the difference between employees' log usage of antidepressants in avg(2006, 2007) and avg(2004,2005). The horizontal axes show the number of employees in the firm minus 35. The employment number is based on year $t-1$. The center line corresponds to a third-order polynomial fit; the dotted lines represent the 95% confidence interval. Scatter points are average changes in individual log usage between avg(2006, 2007) and avg(2004,2005) for firms with a given employee size. The p-value of the difference in Panel (A), female workers, is 0.003 and the p-value of the difference in Panel (B), male workers, 0.57.

Figure 5: Magnitudes: comparison with other life events

Panel (A) Partner death, female workers

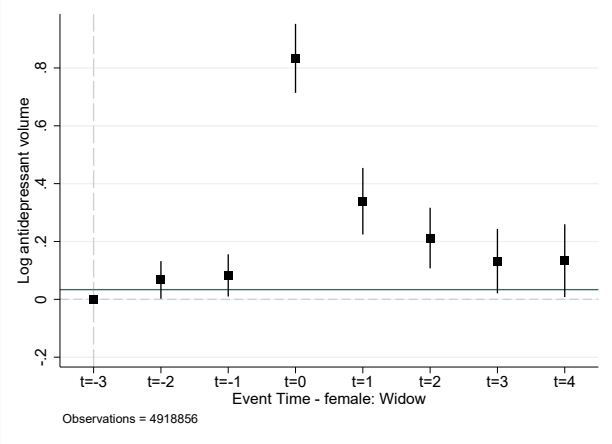

Panel (B) Partner death, male workers

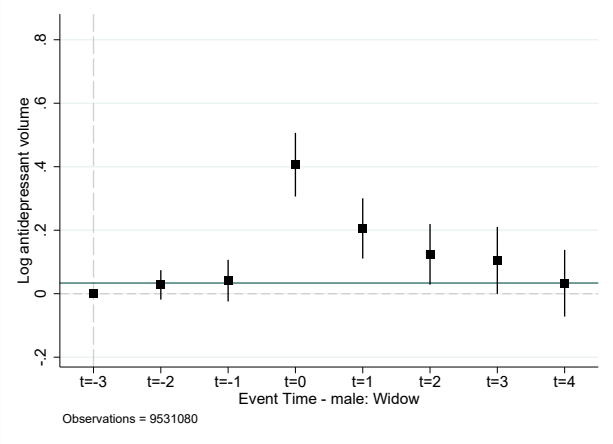

Panel (C) Divorce, female workers

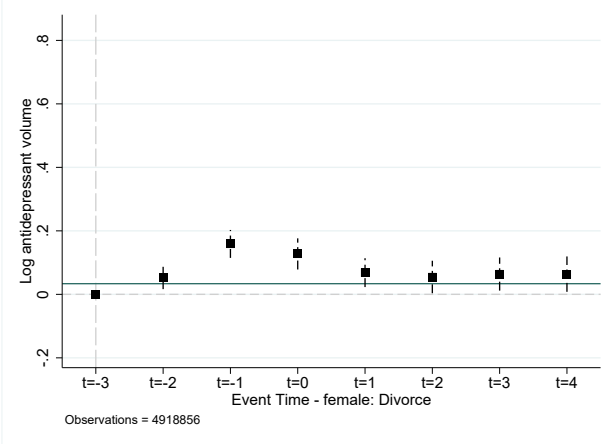

Panel (D) Divorce, male workers

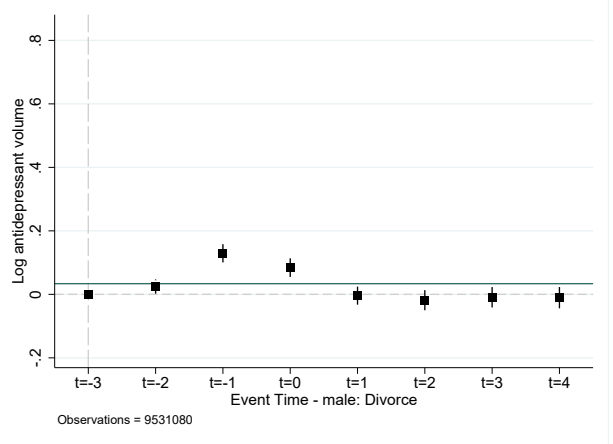

Note: Data from Danish statistics, annual use of antidepressants. The solid green line denotes the benchmark result for the post-pay gap law period. Panels A and B show the results for change in anti-depressant use following a partner's death, for women and men respectively. Panels C and D show the results for change in anti-depressant use following a divorce, for women and men, respectively.

Figure 6: Mechanisms: anti-depressant use volume, by category

Panel (A) High wage workers

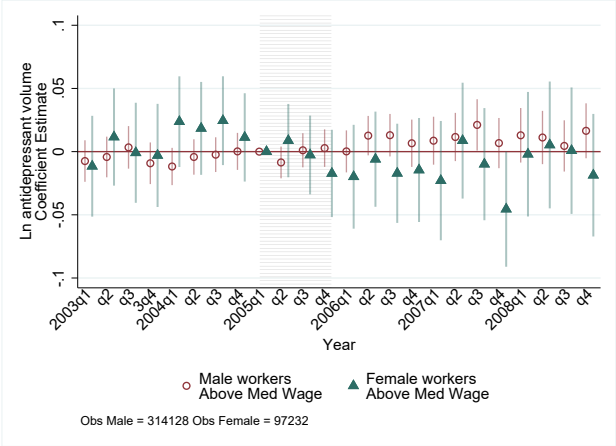

Panel (B) Low wage workers

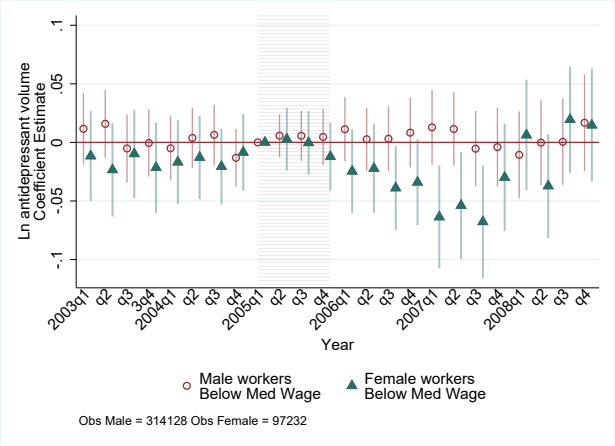

Panel (C) High firm pay gap

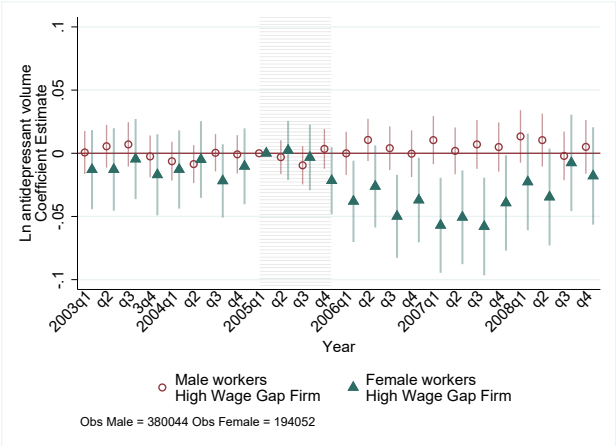

Panel (D) Low firm pay gap

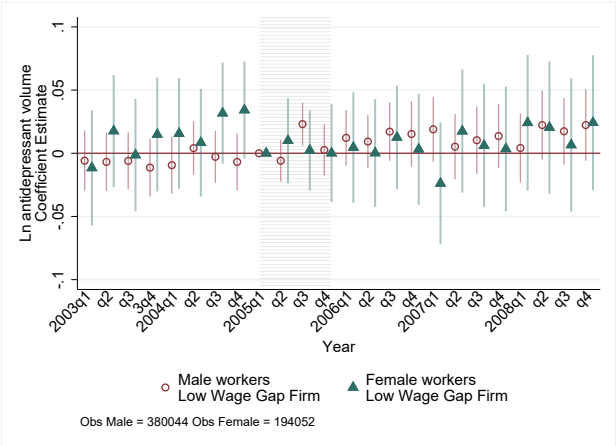

Panel (E) No female manager in pre-period

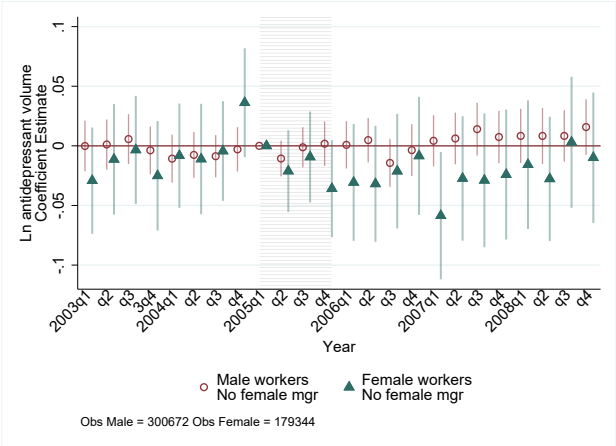

Panel (F) At least one female manager in pre-period

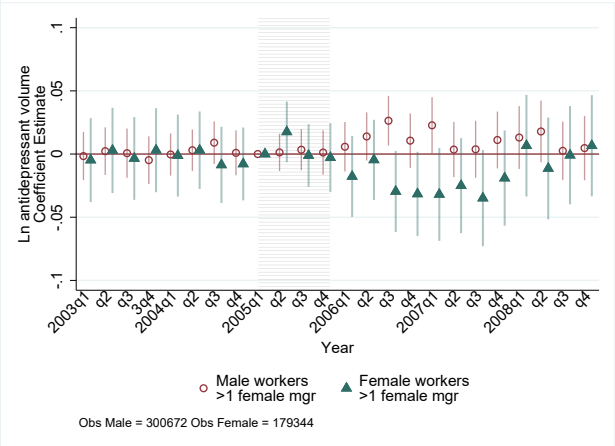

Note: All graphs have log antidepressant use as the outcome variable. Panel A and B split the sample of workers into above and below median within-firm wage ("high wage" and "low wage", respectively). Panels C and D split the sample by above and below median firm-level pay gap in the pre-period ("high" and "low" pay gap, respectively). Panels E and F split firms as having zero or at least one female manager in the pre-period.

Table 1: Effects of pay transparency on mental health

Dependent variable:	**ln(antidepressant volume)**				**Antidepressant user = 1**			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treated × Post	0.016**	-0.037**	0.016**		0.004*	-0.005	0.004*	
	(0.007)	(0.015)	(0.007)		(0.002)	(0.004)	(0.002)	
Female × Post			0.059***	0.062***			0.012***	0.013***
			(0.012)	(0.014)			(0.003)	(0.004)
Treated × Post × Female			-0.052***	-0.070***			-0.008*	-0.013**
			(0.016)	(0.019)			(0.005)	(0.005)
R^2	0.718	0.699	0.710	0.737	0.613	0.614	0.615	0.652
Year FE	✓	✓	✓	✓	✓	✓	✓	✓
Firm × Year FE				✓				✓
Person × Firm FE	✓	✓	✓	✓	✓	✓	✓	✓
# Employee-Years	125577	61369	186946	185805	125577	61369	186946	185805
# Employees	36123	18324	54447	54447	36123	18324	54447	54447
# Firms	3283	3283	3283	3283	3283	3283	3283	3283
Dep Var Mean	0.11	0.21	0.144	0.144	0.036	0.065	0.045	0.045
Sample	**Male**	**Female**	**All**	**All**	**Male**	**Female**	**All**	**All**

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. The dependent variable is the log volume of antidepressant use (in DDD) by each year. Treated is a dummy variable that takes the value of one for individuals working in firms with 35 to 50 employees before the introduction of the law and zero for employees in firms with 20 to 34 employees. Post takes the value of zero for 2003-2005, and value of one for 2006-2008. Standard errors are two-way clustered at the individual and firm levels and reported in parentheses.

Table 2: Placebo: effects of pay transparency on other drug use

	Cholesterol			**Diabetes**		
	(1)	(2)	(3)	(4)	(5)	(6)
	ln(volume)	ln(volume)	ln(volume)	ln(volume)	ln(volume)	ln(volume)
Treated × Post	-0.000	0.013	0.013	-0.002	-0.006	-0.005
	(0.015)	(0.012)	(0.012)	(0.006)	(0.005)	(0.005)
Female × Post			-0.024*			-0.013**
			(0.013)			(0.006)
Treated × Post × Female			-0.014			0.003
			(0.019)			(0.008)
R^2	0.781	0.808	0.802	0.906	0.915	0.913
Year FE	✓	✓	✓	✓	✓	✓
Person × Firm FE	✓	✓	✓	✓	✓	✓
# Employee-Years	186946	186946	186946	186946	186946	186946
# Employees	54447	54447	54447	54447	54447	54447
# Firms	3283	3283	3283	3283	3283	3283
Dep Var Mean	0.18	0.11	0.16	0.07	0.05	0.08
Sample	**Male**	**Female**	**All**	**Male**	**Female**	**All**

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Columns (1) to (3) outcome variable is the log of cholesterol drug volume (in DDD) by each worker, columns (4) to (6) outcome variable is the log of diabetes drug volume (in DDD) by each worker. Treated is a dummy variable that takes the value of one for individuals working in firms with 35 to 50 employees before the introduction of the law and zero for employees in firms with 20 to 34 employees. Post takes the value of zero for 2003-2005, and value of one for 2006-2008. Standard errors are two-way clustered at the individual and firm levels and reported in parentheses.

Table 3: Mechanisms at the employee level: change in earnings

Panel A: Female workers	**All**	**High wage**	**Low wage**	**High pay gap**	**Low pay gap**
Dependent variable: ln(wage)	(1)	(2)	(3)	(4)	(5)
Treated × Post	-0.002	0.006	-0.003	-0.008	0.010
	(0.005)	(0.005)	(0.007)	(0.006)	(0.006)
R^2	0.836	0.879	0.768	0.829	0.846
Year FE	✓	✓	✓	✓	✓
Person × Firm FE	✓	✓	✓	✓	✓
# Employee-Years	61319	23003	27532	40127	21192
# Employees	18301	5960	7526	12104	6188
# Firms	1153	3933	6435	8696	12614
Dep Var Mean	12.761	12.922	12.652	12.785	12.716
Level	Employee	Employee	Employee	Employee	Employee
Panel B: Male workers	**All**	**High wage**	**Low wage**	**High pay gap**	**Low pay gap**
Dependent variable: ln(wage)	(1)	(2)	(3)	(4)	(5)
Treated × Post	-0.014***	-0.008**	-0.018*	-0.015***	-0.012*
	(0.004)	(0.003)	(0.010)	(0.005)	(0.007)
R^2	0.861	0.883	0.770	0.855	0.851
Year FE	✓	✓	✓	✓	✓
Person × Firm FE	✓	✓	✓	✓	✓
# Employee-Years	125428	74838	29556	81817	43611
# Employees	36061	18754	8095	23450	12610
# Firms	9643	18727	19129	3899	24424
Dep Var Mean	12.986	13.104	12.747	13.049	12.869
Level	Employee	Employee	Employee	Employee	Employee

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Outcome variable in all columns is the log of employee wage. Treated is a dummy variable that takes the value of one for individuals working in firms with 35 to 50 employees before the introduction of the law and zero for employees in firms with 20 to 34 employees. Post takes the value of zero for 2003-2005, and value of one for 2006-2008. Standard errors are two-way clustered at the individual and firm levels and reported in parentheses. High wage workers are defined as those with above median wages in their within-firm distribution. Low wage workers are defined as those with below median wages in their within-firm distributions. High pay gap firms are those with above-median pay gap in the pre-period. Low pay gap firms are those with below-median pay gap in the pre-period.

Table 4: Mechanisms at employee level: change in promotions

Dependent variable:	**High pay gap firms**			**Low pay gap firms**		
Employee was promoted = 1	(1)	(2)	(3)	(4)	(5)	(6)
Treated × Post	0.014*	0.017*	0.014*	0.001	0.019**	0.001
	(0.008)	(0.010)	(0.008)	(0.009)	(0.009)	(0.009)
Female × Post			-0.001			-0.011
			(0.007)			(0.007)
Treated × Post × Female			0.003			0.019*
			(0.011)			(0.011)
R^2	0.322	0.332	0.325	0.329	0.333	0.330
Year FE	✓	✓	✓	✓	✓	✓
Person × Firm FE	✓	✓	✓	✓	✓	✓
# Employee-Year	68002	32321	100323	37000	17951	54951
# Employees	17713	8740	26364	9840	4819	14661
Dep Var Mean	0.059	0.056	0.058	0.055	0.048	0.053
Sample	**Male**	**Female**	**All**	**Male**	**Female**	**All**

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Outcome variable in all columns is an indicator for whether the employee was promoted. Treated is a dummy variable that takes the value of one for individuals working in firms with 35 to 50 employees before the introduction of the law and zero for employees in firms with 20 to 34 employees. Post takes the value of zero for 2003-2005, and value of one for 2006-2008. Standard errors are two-way clustered at the individual and firm levels and reported in parentheses. High pay gap firms are those with above-median pay gap in the pre-period. Low pay gap firms are those with below-median pay gap in the pre-period.

Table 5: Mechanisms at employee level: female managers

	High pay gap		**Low pay gap**	
	(1)	(2)	(3)	(4)
	% fem mgr	>1 fem mgr	% fem mgr	>1 fem mgr
Treated × Post	0.009	0.050**	-0.011	-0.000
	(0.008)	(0.019)	(0.010)	(0.021)
N	8013	8013	6128	6128
R^2	0.739	0.738	0.768	0.755
Year FE	✓	✓	✓	✓
Firm FE	✓	✓	✓	✓
# Firms	1655	1655	1322	1322

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Outcome variable in Columns (1) and (3) are the percentage of managers in a firm who are female. Columns (2) and (4) are an indicator that takes a value of one if there is at least one female manager in place. Treated is a dummy variable that takes the value of one for individuals working in firms with 35 to 50 employees before the introduction of the law and zero for employees in firms with 20 to 34 employees. Post takes the value of zero for 2003-2005, and value of one for 2006-2008. Standard errors are clustered at the firm level and reported in parentheses. High pay gap firms are those with above-median pay gap in the pre-period. Low pay gap firms are those with below-median pay gap in the pre-period.

Table 6: Mechanisms at firm level: change in labor flows

Panel A: Female workers	**All firms**		**High pay gap**		**Low pay gap**	
	(1) % Joiners	(2) % Leavers	(3) % Joiners	(4) % Leavers	(5) % Joiners	(6) % Leavers
Treated × Post	0.010***	-0.000	0.009**	0.003	0.010**	-0.003
	(0.003)	(0.003)	(0.004)	(0.004)	(0.005)	(0.004)
N	17656	17656	9457	9457	8199	8199
R^2	0.399	0.381	0.416	0.406	0.376	0.355
Year FE	✓	✓	✓	✓	✓	✓
Firm FE	✓	✓	✓	✓	✓	✓
# Firms	3552	3552	1882	1882	1670	1670
Panel B: Male workers	**All firms**		**High pay gap**		**Low pay gap**	
	(1) % Joiners	(2) % Leavers	(3) % Joiners	(4) % Leavers	(5) % Joiners	(6) % Leavers
Treated × Post	-0.006	-0.007*	-0.011*	-0.010*	-0.000	-0.004
	(0.004)	(0.004)	(0.006)	(0.006)	(0.007)	(0.006)
N	17656	17656	9457	9457	8199	8199
R^2	0.374	0.385	0.368	0.388	0.380	0.382
Year FE	✓	✓	✓	✓	✓	✓
Firm FE	✓	✓	✓	✓	✓	✓
# Firms	3552	3552	1882	1882	1670	1670

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Outcome variables in Columns (1), (3) and (5) are the share of new employees (female in Panel A and male in Panel B) divided by total employment in a given year. Treated is a dummy variable that takes the value of one for individuals working in firms with 35 to 50 employees before the introduction of the law and zero for employees in firms with 20 to 34 employees. Post takes the value of zero for 2003-2005, and value of one for 2006-2008. Standard errors are clustered at the firm level and reported in parentheses. High pay gap firms are those with above-median pay gap in the pre-period. Low pay gap firms are those with below-median pay gap in the pre-period.