Gender Disparities in the Welfare Effect of the Minimum Wage*

Decio Coviello, Erika Deserranno, Nicola Persico

March 16, 2025

Abstract

We study how women and men working in the same minimum-wage supported job respond to, and benefit from, a minimum wage increase. Using administrative data from a major US retailer, we find that the welfare of women increases less with the minimum wage hike than that of men, even though both attain comparable pay raises. We show that this occurs because women exert more effort in response to the minimum wage increase, driven by their greater need for job retention due to less favorable outside options. This evidence points to a generalizable mechanism whereby disparities *outside* the firm account for welfare disparities in the impact of an important gender-neutral policy (i.e., the minimum wage) *inside* the firm.

Keywords: gender inequality, welfare gap, minimum wage, outside option.
JEL Classification: M52, J30, J24, J60.

*Decio Coviello (HEC Montreal). Erika Deserranno (Northwestern University and Bocconi University). Nicola Persico (Northwestern University). We thank Paul Oyer, Maddalena Ronchi, Alan Manning, Till von Wachter for insightful comments on the paper. We also thank participants at the NBER Law and Economics Group, CEPR Labour Workshop, CSEF-IGIER Symposium on Economics and Institutions, CEPR Joint Workshop on Incentives, Management and Organization and Entrepreneurship, Strategy and Firm Dynamics, IGIER B.L.O.W. Workshop, Yale Kuznets Conference and many seminar participants. We are grateful to Elizabeth Jaramillo, Paulina Koenig, Federico Lenzi, Virginia Manzo, and Francesco Marchese for their (more than) excellent research assistance. The authors declare they have no relevant or material financial interests that relate to the research described in this paper. This work has been screened to ensure no confidential information is revealed. Data and institutional background are provided in such a way as to not disclose information that may allow the firm to be identified. This paper is undertaken, in part, thanks to funding from the Canada Research Chairs program. E-mails: decio.coviello@hec.ca; erika.deserranno@kellogg.northwestern.edu; nicola@nicolapersico.com.

1 Introduction

This paper addresses an important fairness question: when a job's working conditions improve, do women and men benefit equally? We ask this question in the context of the minimum wage, using data from salespeople at a major US retailer.

We show that women and men who work in the same position and under the same pay scheme react similarly to a minimum wage increase if, and only if, their outside option is similar. However, whenever women's outside options differ – which is often the case in our setting, and typically means they are lower than those of men – women respond differently than men: they exert more effort because, we argue, they are more concerned about retaining their job due to worse employment conditions outside the firm. Utilizing a new formula to assess the welfare effect of the minimum wage, we find that women derive less welfare benefit than men in the same position.

These findings demonstrate empirically that, even if the *pay scheme inside the firm* is scrupulously gender-neutral, the *overall incentive scheme* is not gender neutral whenever the workers' outside options differ by gender. As a consequence, the welfare effects of gender-neutral policies may differ by gender. This empirical finding highlights a little-noted but generalizable consequence of efficiency wage theory: when fear of termination is part of the overall incentive scheme for workers, disparities outside the firm beget disparities inside the firm.

Our evidence comes from salespeople who work at a large US retailer employing more than 10% of department store employees nationwide, and operating more than 2,000 stores across all fifty states. Our workers' pay is, in part, based on individual productivity (sales per hour) which is recorded administratively. When a worker's average hourly pay falls below the minimum wage, the employer is required to pay a "top-up" to make up the difference.

Our data cover 70 minimum wage increases at the state and local levels. Using a border-discontinuity research design, we study the differential gender effects of the minimum wage on pay and welfare by comparing gender differences in stores where the minimum wage has increased ("treated" stores) with those in stores where it has not ("control" stores) across the *same* county border. Because our stores are composed of two departments, and women disproportionately work in the lower-paying one, our *ceteris-paribus* specification includes department×store fixed effects, and thus effectively compares women and men in the same working conditions within the firm (although the workers' outside options are not held fixed).

We find that, *ceteris paribus*, women receive a pay raise comparable to that of men when the minimum wage increases. However, women respond by putting forth extra additional productivity (larger increase in sales per hour), and are rewarded with extra job stability (larger increase in retention). We attribute this stronger *productivity* response to women exerting *more effort* when the minimum wage increases. We argue that women exert more effort because, while a minimum wage increase makes the current job more worth keeping for all workers, the incentive effect is especially strong for women, as they have worse outside options than men; this drives them to exert extra effort.

Attributing the observed productivity response to worker effort is consistent with the following empirical findings. *First*, we find that all workers become (weakly) more productive after a minimum wage increase. This makes sense because, for workers who benefit from the minimum wage in their current job, a minimum wage increase makes their job more valuable,[1] and so workers exert (weakly) more effort to avoid being terminated. *Second*, we find that the magnitude of the productivity response to the minimum wage is the same for women and men with similar outside options (i.e., similar employment opportunities outside the firm). This makes sense if workers who have a similar fear of termination respond by exerting an equal amount of additional effort. *Third*, we find that this productivity response decreases (for both genders) with the value of a worker's outside option. This makes sense because the incentives to keep one's job are strongest for those who most fear termination. *Fourth*, we find that there is no productivity response (for either gender) during times when workers

[1]The minimum wage happens to have a much smaller effect on the workers' outside option, that is, on their expected employment opportunities outside the firm.

are unmonitored, in which case their probability of termination is less dependent on individual productivity. This makes sense because, if workers cannot affect the risk of termination by working harder, there is no reason to work harder after a minimum wage increase.

The second above-mentioned empirical finding indicates that the observed gender disparities in our workers' response to the minimum wage arise entirely from different outside option levels. That is, male and female workers respond similarly when they face similar outside employment opportunities; it is only because women *generally* have fewer such opportunities that, *on average*, female workers are more responsive than male workers to the minimum wage (because, on average, women are more fearful of being terminated). Several alternative explanations for why women have a stronger productivity response than men are ruled out; these include gender-specific traits (e.g., differences in marginal cost of effort, innate job aptitude, risk aversion, propensity to reciprocate, or childcare constraints), post-minimum wage firm-level adjustments that might disproportionately affect women (e.g., reduced hours or increased monitoring), and gender-specific demand shocks or product price changes. Each of these explanations would predict that men and women with similar outside options respond differently to the minimum wage – yet, the data do not support this.

The four above-mentioned empirical findings point to an efficiency wage model à la Rebitzer & Taylor (1995) where workers exert effort in order to retain their job, and the incentives to exert effort are a function of the difference between inside option (wage scheme, including current minimum wage level) and outside option (opportunities in the outside labor market). A natural interpretation of these empirical findings is that, in our setting, male and female workers differ mainly in their outside option and not much in other dimensions, including effort cost – though we note that gender differences in the cost function do not affect the welfare analysis below.

In an efficiency wage model such as described above, worker welfare involves not only the present (current employment), but also the future (probability of staying with the firm). Therefore, the minimum wage affects worker welfare through several countervailing forces. Empirically, we see that women benefit less than men from a minimum wage increase because they work extra hard after a minimum wage increase (effort cost) and, also, because their pay is topped up by the minimum wage less often (because they work harder). On the other hand, women benefit more than men because they are retained more. To boil down these countervailing effects to a single number, we turn to theory.

We derive a novel (to our knowledge) theory-based formula for the impact of the minimum wage on the welfare of minimum-wage supported workers. The formula says that a worker's welfare gain from a minimum wage increase is the product of two terms. The first term is a capitalization factor that captures a worker's expected tenure in her current job, and is higher for women. The second term, the flow benefit of an increased minimum wage, tends to be larger for men because *ceteris paribus* – i.e., comparing women and men in the same department – the men's pay is more frequently topped up. Conveniently, both terms can be calculated without knowledge of the worker's (unobservable) effort cost: the direct effect of effort on welfare "cancels out" in the formula due to an envelope condition. In our calibrations, the second effect dominates, leading us to conclude that the welfare benefits of the minimum wage are larger for male than female workers, *ceteris paribus.*

The welfare estimates flip in a *non-ceteris-paribus* analysis where we remove the department×store fixed effects, allowing the estimates to capture the consequences of women being employed in the lower-paying department. In this analysis, women benefit more than men from the minimum wage because, mechanically, their pay is topped up more often by the minimum wage. This fact implies that the minimum wage is a force for gender equalization – simply because female workers are dissimilarly situated than male workers. However, among similarly situated workers, a higher minimum wage disproportionately benefits men, largely due to their more favorable employment options outside the firm.

This paper makes two novel contributions. First, our case study illustrates empirically a little-noted but generalizable consequence of efficiency wage theory: that disparities *outside* the firm (gender differences in outside options) beget disparities *inside* the firm – in our case, disparate welfare impact of a gender-neutral policy.

Therefore, a "systemic" gender disparity outside the firm may require affirmative correction even within a scrupulously gender-neutral firm. Second, ours is the first paper, to our knowledge, that quantifies the impact of a gender-neutral policy (here, the minimum wage) on the gender gap *in welfare*, as opposed to pay. This is important because, when costly effort is endogenous and workers care about retention, pay is not the same as welfare. This observation suggests that policies should be evaluated by considering their effects on the gender welfare gap, rather than focusing solely on the gender pay gap, as these may not align.

Our paper contributes to several literatures. Starting with the literature on the disparate impact of the minimum wage by gender, Caliendo & Wittbrodt (2022), Blau et al. (2023), and Paul-Delvaux (2023) study the differential gender effect of the minimum wage on wages. In line with our *non-ceteris-paribus* results, they find that a higher minimum wage reduces the gender pay gap because women tend to be overrepresented in lower-paying positions.[2] Whereas these papers focus on wages, we also document the disparate effects of the minimum wage on a rich set of outcomes including productivity, retention and, most notably, welfare, in addition to wages. Furthermore, the existing estimates in the literature are not *ceteris paribus*, i.e., they do not compare women and men in the same role. However, when evaluating the "fairness" of a policy, we show that it is important to also make comparisons among workers in similar positions, as these may differ (or even reverse) from comparisons between women and men in different positions.

Unrelated to gender, a relatively small literature focuses on the productivity effects of the minimum wage. This literature finds that the minimum wage increases worker productivity (Ku, 2022; Coviello et al., 2022; Ruffini, 2024), as do we. Unlike our paper, these studies do not focus on the role of the outside option. In Flinn (2006), the minimum wage impacts the search activity of the unemployed but it is assumed to have no effect on the productivity of employed workers.

[2]Unrelated to gender, some studies investigate the effects of the minimum wage on other groups of workers, including teenagers (e.g., Giuliano 2013, using personnel data), higher-paid co-workers (Dube et al., 2019), and minorities (e.g., Derenoncourt & Montialoux 2021, using macro data).

Unrelated to the minimum wage, a number of papers have studied the role of the outside option on worker productivity and incentives to exert effort. Lazear et al. (2016) show that workers employed in a large US firm are less productive in times of low unemployment, when their outside option is better, and attribute this effect to lower individual effort. However, their analysis does not differentiate by gender. Separately, improvements in workers' outside options given by an extension of unemployment benefits have been shown to increase worker absenteeism and reduce worker productivity (Ahammer et al., 2023; Lusher et al., 2022). These findings align with ours; we add to this literature by studying the effect of the minimum wage and by shifting the focus to worker welfare.

Finally, we contribute to the literature on gender disparities arising from ostensibly gender-neutral pay policies. Bolotnyy & Emanuel (2022) find that a gender-neutral pay scheme among bus drivers results in a gender wage gap because women do less overtime. Biasi & Sarsons (2022) show that wage flexibility favors male teachers over female teachers due to stronger bargaining abilities, and Antecol et al. (2018) find that gender-neutral tenure-clock stopping policies increase gender gaps in tenure in high-skilled professions. We complement these papers by studying a different policy (the minimum wage) and by studying welfare.

2 Data and Identification Strategy

2.1 Institutional setting and data

Our data cover more than 40,000 consultative sales associates working in more than 2,000 stores at a nationwide US retailer from February 2012 to June 2015. Our border discontinuity research design (see Section 2.2) restricts attention to a subsample of more than 200 stores that share an administrative border. This "border store sample" covers over 10,000 consultative sales associates, about 7,000 of which are administratively classified as men.[3]

[3]Henceforth, all the information we report refers to this "border store sample."

Consultative sales associates assist walk-in customers by answering their questions and demonstrating product features. These tasks, collectively referred to as "exerting effort", involve warmly greeting the customers, patiently explaining and persuading, up-selling higher-margin products, and cross-selling items such as warranties, loans, and credit cards. Each store has, on average, 16 consultative sales associates, a store manager, and typically one supervisor per department. In what follows, we describe the summary statistics of male and female consultative sales associates. These statistics are reported in Table A.1.

Age, tenure, and termination The workforce is relatively young: average age is 36 (median: 27), with similar age distributions for men and women. Women have longer average tenure (58 vs. 44 months; median: 27 vs. 22 months) and lower termination rates (4.1% vs. 4.8% per month). Table A.2 shows that termination correlates with low productivity (low sales per hour) for both genders, particularly when workers are supervised (more on supervision in Section 4.3). We generally combine voluntary and involuntary terminations due to the subjectivity of the distinction, but also present findings for involuntary terminations. Career advancement is limited, as few consultative sales associates are ever promoted to managerial roles.

Pay and department allocation Our consultative sales associates are paid by the hour based on a nationwide compensation scheme. Compensation includes a fixed and a variable portion; the latter is based on customer purchases which each associate claims as her own sales. A sale associate's *regular pay* includes a fixed component (*base hourly pay*) and a variable component (*commissions* based on customer purchases which each associate claims as her own sales). On average, regular pay is $12 per hour, with $6 stemming from the fixed component and another $6 from commissions. If the weekly average of a worker's regular pay per hour falls below the minimum wage, the employer is required to make up the difference as prescribed by the Fair Labor Standards Act.[4] We create a variable called *minimum wage top-up* which equals the average hourly amount paid by the employer to comply with the minimum

[4]Under this law, commissioned workers can occasionally be deemed "exempt" and thus not receive a top-up. Based on administrative records, however, all of the workers in our sample are non-exempt.

wage. Approximately 42% of our workers receive some top-up in at least one week of a month and, among these workers, the average top-up amount is $0.50 per hour. However, only 3.2% of our workers receive a top-up in every week of the month (and so are paid exactly minimum wage in that month). Later in our analysis, we will refer to *total pay* as regular pay plus any top-up.

Within a store, employees work in different units that sell different product types. Following an internal company classification, we group units into two "departments," denoted A and B for confidentiality. Each department in a store has its own supervisor. Employees in department A earn significantly more than their counterparts in department B: refer to Figure A.1, panel A.[5]

The gender composition differs across departments, with men making up 75% of workers in department A and only 9% in department B. Across our firm, men earn more than women and are substantially less likely to be situated at the lower end of the pay distribution.[6] However, these gender pay disparities are entirely explained by the disproportionate allocation of women to the lower-paying department B. In fact, within a department, the gender pay gap disappears: if anything, women appear to earn slightly *more* than men despite facing a similar compensation scheme as men,[7] and are substantially *less likely* to be at the lower end of the pay distribution.[8] This suggests that, mechanically, women are less likely to benefit from the minimum wage relative to the men in their *same* department.

Sales per hour/productivity Both female and male sales associates work an average of 28 hours per week.[9] We compute "sales per hour" as the value of sales divided by the number of hours worked. We refer to sales per hour interchangeably

[5]Median pay is $12.6 per hour in department A and $9.9 in department B. The share of workers who receive a top-up in a month is 34% in department A vs. 69% in department B.

[6]See Figure A.1, panel B. The gender pay gap is 4.5% across our firm: median pay is $11.3 for men and $10.8 for women. 36% of men receive a "top-up" during the month, compared to 53% of women.

[7]Table A.3 shows that, within each department, women and men have a similar base hourly rate and commission rate. Within department A (resp., B), the median pay per hour is $12.2 for women vs. $11.3 for men (resp., $9.6 for women vs. $9.2 for men).

[8]See Figure A.1, panels C and D.

[9]75% of male and female employees work 24 hours per week, while the rest work 35-40 hours.

as "productivity." Because women disproportionately work in department B, they sell less than men on average (Figure A.2, panel A); however, this is no longer the case when holding the department fixed (panel B). Notably, within a department, women's sales per hour are less likely to be at the lower end of the distribution compared to men's.[10]

Our workers resemble US hourly workers In Appendix B, we show that our male and female workers resemble US hourly workers in terms of their exposure to the minimum wage, earnings, and termination rates. Hourly paid workers make up 58% of US workers and likely account for an even greater proportion of minimum wage recipients.[11]

2.2 Identification strategy

Between February 2012 and June 2015, the stores in our sample experienced 70 minimum wage increases: 49 at the state level and 21 at the county or city level. The prevailing minimum wage in a locality is determined as the highest rate across the state, county, or city levels.[12]

Our empirical specification implements a border discontinuity design in the spirit of Card & Krueger (2000), and closely follows Dube et al. (2010) and Allegretto et al. (2011). Specifically, workers on the side of the border where the minimum wage increased (treatment group) are compared to workers on the other side, where the minimum wage did not increase (control group). This research design has the advantage of ensuring that, apart from the minimum wage change, treated and control groups are similar in terms of local economic conditions and demand shocks. The

[10] The units of the sales per hour measure are shrouded for confidentiality reasons, and are rescaled by a factor between 1/50 and 1/150 relative to dollar value.

[11] The Current Population Survey (CPS) does not report the fraction of minimum wage recipients who are hourly workers, however, in the CPS data, hourly workers' weekly earning are about half of non-hourly workers' earnings, suggesting that hourly workers are more affected by the minimum wage (US Bureau of Labor Statistics, 2015a).

[12] In our sample, the average minimum wage is \$7.84 per hour, with a s.d. of \$0.50, and the average increase is \$0.54. Appendix C presents a full list and map of the minimum wage changes.

main disadvantage of this approach is the risk of cross-border worker movements from control to treated stores (Neumark et al., 2014), but we will argue that this risk is minimal in our setting (see page 15).

Following Card & Krueger (2000), Dube et al. (2010, 2016) and Allegretto et al. (2017), we restrict our sample to stores (and their respective workers) located in adjacent counties that share a border and whose centroids are less than 75 km apart. This subset comprises over 200 stores and more than 10,000 salespeople, half of which experience variations in the minimum wage during our study period.[13]

***Ceteris-paribus* impact of the minimum wage by gender** In Section 3, we will assess the causal effect of the minimum wage for women vs. men *under the same working conditions*. We estimate the following specification:

$$Y_{idjpt} = \alpha + \beta M_{jt} + \gamma M_{jt} \times Woman_i + \eta X_{idjpt} + \delta_i + \zeta_{dj} + \phi_{pt} + \varepsilon_{idjpt}. \qquad (1)$$

Y_{idjpt} is the outcome variable of interest (pay, retention, productivity, and, later, welfare) for worker i in department d of store j of county-pair p in month t. $Woman_i$ is an indicator for whether worker i is a woman. M_{jt} is the prevailing minimum wage in store j's jurisdiction in month t, expressed in dollars. The coefficients β and $\beta + \gamma$ capture the effect of increasing the minimum wage by \$1 on men and women, respectively.[14] The coefficient γ captures the differential effect of the minimum wage by gender, which is the focus of this paper. To ensure that this differential effect does *not* capture different working conditions across gender, equation (1) includes department×store fixed effects ζ_{dj}, thus effectively comparing women and men in the same department within the same store. We also include worker fixed effects δ_i to

[13] See Appendix C.2 for details on data construction. We use the 75 km threshold which is standard in the literature. In Section 3.2, we will show that our main results are similar if we use different thresholds, e.g., stores in bordering counties with centroids less than 37.5 km or 18.75 km apart. This is reassuring because by narrowing down the definition of bordering counties, we increase the comparability between treated and control stores, although it reduces the sample size.

[14] A \$1 increase corresponds to an increase of two standard deviations in the minimum wage, or a 13% rise relative to the average minimum wage level.

account for time-invariant worker characteristics such as ability.[15]

We implement the border discontinuity design by including county-pair × month fixed effects in equation (1), thus effectively restricting the comparison to "treated" and "control" stores/workers on either side of the *same* border. We estimate this equation by "stacking" the data, meaning that stores/workers located in a county sharing a border with n other counties appear n times in the final sample. The standard errors are two-way clustered at the state and border-segment level.[16]

In our main specification, X includes $M_{jt} \times Department_d$ to control, for example, for the fact that a higher minimum wage may increase demand in one department more than another. The results are robust across several alternative specifications (see Section 3.2), including adding department×store×month fixed effects (ζ_{djt}) to control for time-varying department characteristics.

***Non-ceteris-paribus* impact of the minimum wage by gender** Replacing the department×store and worker fixed effects in equation (1) with store fixed effects yields *non-ceteris-paribus* estimates of the differential effect of the minimum wage by gender (γ) which incorporate the fact that women are disproportionately represented in the low-paying department relative to men. Estimates are discussed in Section 6.

3 *Ceteris Paribus*, Men and Women Respond Differently to the Minimum Wage

3.1 Main results

Table 1 documents the *ceteris paribus* impact of the minimum wage by gender on pay, individual productivity (sales per hour), and retention. It also displays the impact

[15] We can identify worker and department×store fixed effects because we observe nearly 20% of our workers switching department or store. In Section E.2, we will show that the minimum wage does not affect the likelihood that female and male workers switch department or store.

[16] We cluster standard errors this way because the presence of a single county in multiple pairs along a border segment induces a mechanical correlation across county-pairs, and potentially along an entire border segment (Dube et al., 2010). Refer to Appendix C.2 for more details.

on welfare, which will be discussed later in Section 5.

Table 1: Impact of the Minimum Wage on Productivity, Pay, Retention, and Welfare by Gender (Ceteris-Paribus Analysis)

	(1)	(2)	(3)	(4)	(5)	(6)
	Productivity		**Pay**		**Retention**	**Welfare**
Dep.Var.	Sales per hour	Total pay per hour = col. (3)+(4)	Regular pay per hour (fixed+variable)	MinW top-up per hour	Retained	Discounted synthetic pay per hour
MinW	0.059	0.556***	0.215	0.341***	-0.004	20.307***
	(0.040)	(0.127)	(0.162)	(0.061)	(0.005)	(3.913)
MinW × Woman	0.055***	0.082	0.338**	-0.256*	0.020***	-9.007**
	(0.016)	(0.163)	(0.124)	(0.126)	(0.003)	(3.792)
Observations	217,746	215,558	215,558	215,558	217,746	197,333
Mean Dep.Var.	2.085	12.271	12.046	0.225	0.954	195.876
p-value MinW+MinW×Woman	0.024	0.027	0.037	0.330	0.020	0.028
Effect MinW for Men (%)	2.5%	4.5%	1.8%	194.9%	-0.4%	10.6%
Effect MinW for Women (%)	6.8%	5.3%	4.7%	26.6%	1.7%	5.5%

Notes: Each observation represents a worker-month. The table reports the estimates of β and γ from specification (1), with the p-value for the null hypothesis H0=β+γ=0 presented at the bottom of the table. All regressions include store×department fixed effects, worker fixed effects, pair×month fixed effects and control for MinW×department. Standard errors are two-way clustered at the state and border-segment level. *** p<0.01, ** p<0.05, * p<0.1. "Sales per hour" are the sales per hour rescaled by a factor between 1/50 and 1/150 relative to its $ value. "Total pay per hour" is the monthly total pay (in $ per hour). "Regular pay per hour" is the total amount earned from the base hourly rate and variable pay (commission rate × sales per hour), without the top-up. "MinW top-up per hour" is the monthly total minimum wage adjustment paid by the company to the worker (in $ per hour). The sample size is smaller for the pay variables because we trim the top 1% of the observations due to presence of outliers. "Retained" is a dummy variable that equals one if the worker is retained that month (i.e., not terminated). "Discounted synthetic pay per hour" is the synthetic pay per hour—i.e., the hourly pay the company would have paid the worker had they made the same sales as in the month before the minimum wage increase, calculated as the maximum of total pay per hour in *t-1* and the minimum wage in *t*—multiplied by the discount factor [(1+r)/(1+r-π)], where r is the monthly discount rate and π is the average monthly retention rate by gender (lagged). "MinW" is the predominant minimum wage in deviation from its sample mean (in $). "Effect MinW for Men (%)" [resp., "Effect MinW for Women (%)"] is the percent effect of a $1 increase in MinW relative to the mean of the outcome variable for men [resp., women].

A $1 increase in the minimum wage increases *total* pay similarly for both genders: +$0.638 per hour (+5.3%) for women and +$0.556 per hour (+4.5%) for men, with the gender difference not statistically significant: see Table 1, column 2. But this similarity masks a very different behavioral response: women's productivity increases by 6.8% (significant at the 5% level), while men's rises by only 2.5% (not significant), with the gender difference significant at the 1% level: see column 1. This is why men receive their pay increase mostly through an increase in top-up, whereas women attain their pay increase mostly through the *regular* component of pay (inclusive of variable pay): see columns 3 and 4.[17,18]

But, if women become more productive after a minimum wage increase compared to men, how are they rewarded for their extra productivity? Table 1, column 5 shows that they are rewarded in the form of greater retention: female retention goes up by 1.6 percentage points (1.7%) with the minimum wage, with no corresponding effect on the retention of male workers (coefficient of -0.004, not statistically significant). Similar results are obtained when terminations are limited exclusively to "involuntary terminations" – see Table A.5, column 5.

Table A.6 shows that the minimum wage does not differentially impact other organizational variables by gender, including hours worked, assignment to better shifts (e.g., busier shopping hours), transfers across departments, and the likelihood of being supervised. See Appendix E.1 for more details on each of these adjustments. Additionally, Table A.7 shows that the minimum wage does not alter the termination rule (the function mapping higher productivity to lower termination rates) differentially by gender. Thus, the increase in women's retention following a minimum wage hike is consistent with them working extra hard relative to men, as we argue in Section 4.

3.2 Robustness checks and threats to identification

Robustness checks Table A.8, column 1 shows that the results are robust to controlling more flexibly for time by adding department$\times$store$\times$month fixed effects in specification (1). Columns 2-5 show that the results are robust to controlling for potential correlates of gender (worker tenure, age, childrearing age, home-work distance)

[17] Whereas the increase in regular pay is more than twice as large for women than for men (\$0.553 vs. \$0.215), the increase in top-up is four times larger for men than for women (\$0.341 vs. \$0.085). The minimum wage increases the share of men and women receiving some top-up over the month by 18.9 and 12.6 percentage points, respectively. While these effects are both large and significant, the effect is significantly larger for men. Similar results are obtained for the number of weeks per month in which the worker receives a top-up. See Table A.5, columns 1 and 2.

[18] These effects do not reflect an adjustment in the compensation scheme by gender. Indeed, in Table A.5, columns 3-4, we find no evidence that the firm adjusted the compensation *scheme* (base or commission rate) to the minimum wage for either gender. This makes sense because the compensation scheme is set nationally and does not respond to local minimum wage adjustments.

and their interaction with the minimum wage, indicating that the heterogeneous effects by gender are not driven by disparities in these potential correlates of gender. Table A.9 shows that the results hold in both department A and B, suggesting they are not driven by workers in a single department. Therefore, we do not believe that department-specific behavioral mechanisms (e.g., female workers in department B trying to maintain top-up frequencies) explain their stronger response. Finally, Table A.10 confirms robustness to using a Poisson pseudo-likelihood or log-log regression.

Pre-trends We rule out the possibility that observed gender differences arise from pre-existing trends in outcome variables within a department before minimum wage changes. To test this, we examine gender-specific pre-trends in the 12 months before the minimum wage change using an autoregressive distributed lag model: see Appendix C.3 for details.[19] Table A.11 shows no gender differential pre-trends.

Worker selection Our estimates could be confounded if, after a minimum wage increase, women's productivity increases more than the men's because of selection into and out of the worker pool by ability, in a way that differs by gender. For example, stores may have retained more-capable women and shed less-capable ones after a minimum wage increase. Although the inclusion of worker fixed effects in our specification should mitigate these concerns because we effectively compare the "same" worker at two minimum wage levels, we dig deeper and replicate our findings in the "non-selected" subsample of workers who were present on the first and the last day of our sample period.[20] When we do this, the sample size drops but the results on productivity are similar to the main sample: see Table A.12, column 1. The results in Section 4 will further rule out this selection story.

Cross-border movements Border-discontinuity research designs are vulnerable to the critique that workers might move from control to treated counties (Neumark

[19] This model, commonly used in the minimum wage literature (Dube et al., 2010), has the advantage of taking into account the sequential occurrence of changes in the minimum wage level.

[20] A caveat: not all the workers in this subsample are employed *continuously* throughout our sample period. Restricting to continuously employed workers leaves us with few observations for our analysis.

et al., 2014). This becomes an issue for our identification strategy if women are less inclined to cross borders compared to men, and if men who cross borders are of particularly low/high ability, leading to a change in the ability composition of the female vs. male workforce in both treated and control counties following the minimum wage increase. In our specification, this confounder is mitigated by the presence of worker fixed effects, and the fact that very few of our workers transfer to a different store on the opposite side of the same county.[21] Moreover, the results are similar if we restrict our analysis to bordering counties with centroids less than 37.5 km, or 18.75 km apart, instead of using the 75km threshold (Table A.12, columns 3-4). If cross-border movements were an issue, we should observe changes in the results as we narrow down the definition of bordering counties.

Spillovers across genders It could be that, after a minimum wage increase, the higher sales by women depress men's sales – a business-stealing spillover effect. However, we find that the estimated impact of the minimum wage on men's productivity remains unchanged when controlling for the department's proportion of female employees interacted with the minimum wage (Table A.8, column 6) or when restricting the analysis to male workers in all-male departments (Table A.12, column 2).

4 Outside Options Explain the Differential Response to the Minimum Wage by Gender

The previous section has shown that women's productivity responds more strongly than men's to a minimum wage increase, and that this has implications for their pay composition and their retention. This section explores why women's productivity response differs from men's. We find that gender differences in outside options fully explain this disparity: men and women with similar outside options respond similarly. Moreover, the effect arises only among supervised, as opposed to unsupervised, workers. The evidence, we argue below, supports an efficiency wage model.

[21] Only 1.2% of male and female workers transfer to a store on the opposite side of the same county.

4.1 Proxy for the workers' outside option

We seek an empirical proxy for our workers' outside option that varies by gender and home zip code. This measure is intended to capture, by gender, the local opportunities and hourly wages available to salespeople residing in a zip code, as well as unemployment duration. In the spirit of Schubert et al. (2024), we define the outside option index (OOI) of a gender-g worker living in zip code z in year y as:

$$OOI_{gzy} = \nu_{gy} \sum_{o} \theta_{og} \cdot \frac{s_{ogzy}}{s_{ogy}} \cdot w_{ogzy},$$

where θ_{og}, the nationwide probability that a gender-g salesperson transitions to new occupation o, is adapted to local×year conditions by the "occupation availability" factor $\frac{s_{ogzy}}{s_{ogy}}$, which is the relative share in year y of gender-g salespeople in occupation o with home zip-code z compared with the nationwide average. The product of these two terms proxies for the availability of different jobs a salesperson can transition to *from their home zip code*.[22] The term w_{ogzy} is the hourly wage that a gender-g worker living in zip code z earns if they are in occupation o in year y. The scaling factor ν_{gy}, which is equal to (employment - unemployment spell)/employment spell, captures unemployment duration by gender at the national×year level; it is decreasing in unemployment duration and equals 1 if unemployment lasts zero weeks.[23] The OOI is measured in dollars per hour, and is computed based on data from the American Community Survey (ACS) and Current Population Survey (CPS). Appendix D.1 details its construction.

Figure 1 presents the distribution of OOI by gender. Panel A shows the distribution of the median OOI at the store-year level for women (red line) and men (blue line). Women tend to have substantially worse outside options than men. This is true, also, when comparing women and men working in the same store (panel B).

[22] Simply put, we "localize" nationwide transition frequencies for salespeople using the prevalence of each occupation among a zip code's residents. This idea is borrowed from Schubert et al. (2024), although our measure is more granular geographically and is gender-specific. Paul-Delvaux (2023) borrows the same idea.

[23] We will show that the results are robust to omitting this scaling factor from the OOI measure.

Figure 1: Distribution of Outside Option by Gender

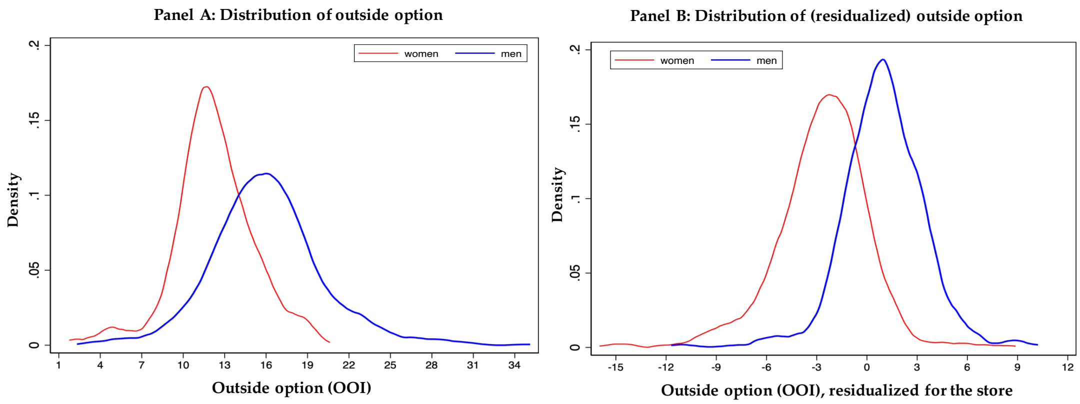

Notes: Panel A shows the distribution of the average outside option index (OOI) in a store-year, separately for women and men. Panel B shows the same distribution for the residualized OOI. The residualized OOI is calculated as the residuals from a regression of the OOI on store fixed effects.

Extensions and cross-validation We can extend the OOI to account for our workers' commuting preferences. In Appendix D.1, we build an extended OOI that takes into account how far a worker lives from their store assuming, in the spirit of Caldwell & Danieli 2023, that workers who live farther from their workplace are more willing to commute and thus have greater access to outside options. Incorporating commuting preferences into the OOI will not materially affect the results.[24] In Appendix D.1, we also show that our results do not change if we remove the scaling factor (ν_{gy}) from the OOI measure, or if the OOI is built using alternative measures of job-to-job transitions (θ_{og}).

To cross-validate the OOI as a good proxy of our worker's outside option, we construct a separate and different proxy for a subset of our workers based on anonymized financial transaction data. We leverage a dataset from a large financial aggregation and analytics firm that happens to cover a third of our workers and contains information on their next-job earnings. Regressing next-job earnings, whenever this measure is available, on the OOI yields an estimated coefficient of 0.605 (significant at the 1% level).[25] This high correlation level is a reassuring cross-validation of the OOI. Refer

[24]Although commuting preferences are believed to differ by gender (Le Barbanchon et al. 2021), in our sample, average home-work distances are similar for men and women (8.9 km vs. 8.4 km).
[25]\$1 per hour increase in the OOI is associated with a \$0.605 increase in next-job hourly wage.

to Appendix D.2 for more details on the data and the correlation.

4.2 Main results

We test whether the productivity response to the minimum wage is a function of the outside option by estimating:

$$\begin{aligned} Y_{idjpt} &= \alpha + \beta M_{jt} + \gamma M_{jt} \times Woman_i + \lambda M_{jt} \times OOI_{gzy-1} + \psi M_{jt} \times OOI_{gzy-1} \times Woman_i \\ &\quad + \theta OOI_{gzy-1} + \mu OOI_{gzy-1} \times Woman_i + \eta X_{idjpt} + \delta_i + \zeta_{dj} + \phi_{pt} + \varepsilon_{idjpt}, \end{aligned} \tag{2}$$

where Y_{idjpt} is the productivity of worker i in department d of store j of county-pair p in month t, and the other variables are defined as in specification (1). OOI_{gzy-1} is the outside option for a gender-g worker living in zip code z during the calendar year $y-1$ prior to month t. We lag the OOI to capture backward-looking expectations and also, conveniently, to ensure that it is pre-determined and exogenous to subsequent minimum wage changes. The coefficient λ (resp., $\lambda + \psi$) measures how the men's (resp., women's) response to the minimum wage varies with their outside option. The coefficient ψ captures the gender differential in the response as a function of the outside option (triple-interaction term).

The identification strategy exploits two sources of variation in the OOI: variation across workers and variation over time within a worker. Decomposing the standard deviation of the OOI reveals that the cross-worker variation is 2.3 times larger than the within-worker variation over time. This indicates that most of the identifying variation in our setting comes from workers who are employed in the same store but reside in different zip codes. These workers experience the same minimum wage increase and store-specific conditions but face differing outside options based on their home zip codes. This source of variation in the workers' outside options is not vulnerable to the concern that our estimates of λ and ψ are confounded by cross-store correlation between the OOI and unobserved store-specific factors such as product prices, customer demand, or gender composition of the clientele.

Columns 1-2 of Table 2 present the results. The coefficient for "MinW×OOI" captures the slope of the workers' productivity response profile as a function of the outside option. The estimated coefficient is negative, indicating that workers' positive productivity response to the minimum wage attenuates as their outside option improves. This finding is a novel empirical contribution to the efficiency wage literature and speaks to the key mechanism of our paper, as described in Section 4.3 below.

Table 2: Impact of the Minimum Wage on Productivity by Gender and Outside Option (Linear)

	(1)	(2)	(3)	(4)	(5)	(6)
Dep.Var.	Sales per hour					
Sample	*Full sample*		*Supervised*		*Unsupervised*	
MinW	0.224***	0.215***	0.300***	0.308***	-0.053	-0.054
	(0.040)	(0.044)	(0.066)	(0.063)	(0.086)	(0.094)
MinW × Woman	-0.008	0.019	0.001	-0.034	-0.023	-0.014
	(0.013)	(0.067)	(0.023)	(0.066)	(0.040)	(0.119)
MinW × OOI	-0.008***	-0.007***	-0.010***	-0.010***	-0.001	-0.001
	(0.001)	(0.001)	(0.002)	(0.002)	(0.004)	(0.004)
MinW × Woman × OOI		-0.002		0.003		-0.001
		(0.004)		(0.005)		(0.007)
Observations	212,443	212,443	162,730	162,730	40,859	40,859
Mean Dep.Var.	2.087	2.087	2.141	2.141	1.900	1.900

Notes: Each observation represents a worker-month. The table reports the estimates of β, γ, λ, and ψ from specification (2) in that order. All regressions include store×department fixed effects, worker fixed effects, pair×month fixed effects and control for MinW×department (estimates omitted in the table for clarity). Standard errors are two-way clustered at the state and border-segment level. *** p<0.01, ** p<0.05, * p<0.1. Columns 3-4 (resp., columns 5-6) restrict the sample to workers with (resp., without) a direct supervisor assigned during that month. Missing OOI data makes the sample in this table slightly smaller than in Table 1, column 1, where results still hold in this sample. The supervisor variable is also missing for some workers, which explains why the number of observations in 2-3 and 5-6 does not sum to that in 1-2. OOI refers to the outside option index, which varies by the worker's home zip code and gender and is lagged by one year. See Appendix D.1 for details on how the OOI variable is constructed.

The coefficient for "MinW×Woman" in column 1 is small and not statistically significant, in stark contrast to Table 1, column 1, where, without controlling for "MinW×OOI," the coefficient was large and significant. That the coefficient for "MinW×Woman" disappears once we control for "MinW×OOI" indicates that gender differences in productivity responses to the minimum wage are entirely driven by differences in outside options.

Column 2 confirms this using the full specification (2): the coefficients for "MinW×Woman" and "MinW×Woman×OOI" are both statistically indistinguishable from zero, indicating that women's productivity response profiles with respect to outside options – both in terms of intercept and slope – are similar to men's. More broadly, the fact that all coefficients involving "×Woman" are small and statistically insignificant reinforces that men and women with similar outside options exhibit similar productivity responses to the minimum wage. This is the key result of this section.

Allowing for a more flexible, non-parametric relationship between responsiveness to the minimum wage and outside option – rather than assuming a linear relationship as in specification (2) – reinforces the main result: the productivity response to the minimum wage is fundamentally similar for men and women (Figure 2, panel A).

Figure 2: Impact of the Minimum Wage on Productivity by Gender and Outside Option (Non-linear)

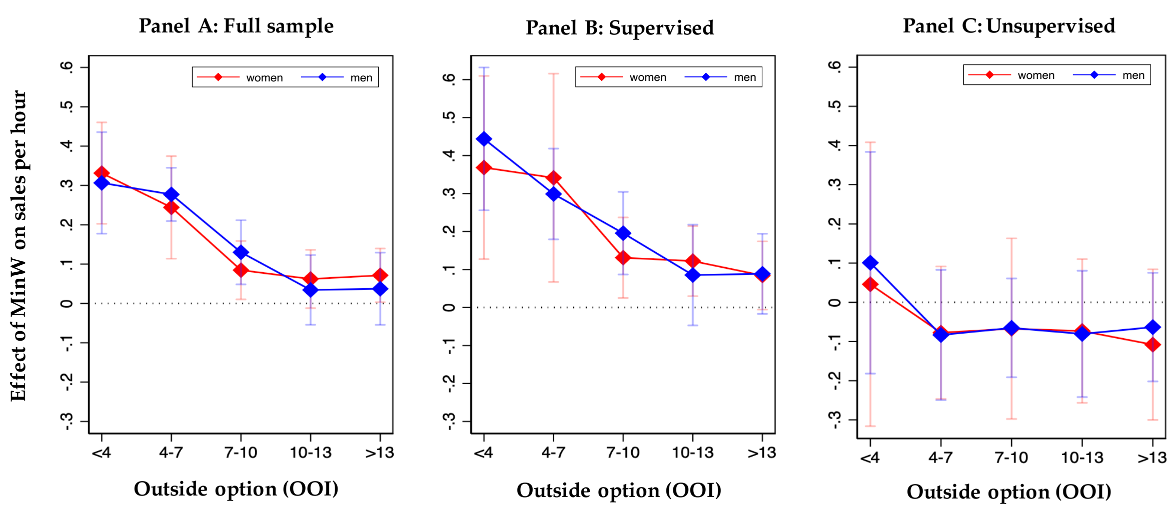

Notes: The red (blue) line represents the estimated effects of the minimum wage on sales per hour for women (men) as a function of the OOI. The estimates capture the "triple interaction" coefficients of the minimum wage interacted with gender (man for the blue line and woman for the red line) and with five indicators for different OOI bins (those on the x-axis). Panel A uses the full sample of workers. Panel B is restricted to workers in months with a direct supervisor assigned to them. Panel C is restricted to workers in months with no direct supervisor assigned to them. The OOI varies with the worker's home zip code and gender and is lagged by one year. See Appendix D.1 for details on how the OOI variable is constructed.

Finally, we show that the main results are robust to tweaking the definition of the OOI measure. These tweaks include: omitting the adjustment for unemployment duration by excluding the scaling factor ν_{gy} from the OOI; extending the OOI to incorporate commuting preferences; and varying how job transitions (θ_{og}) are measured. See Table A.13, and the discussion in Appendix D.1.

4.3 Interpretation

Evidence supporting efficiency wages The estimates from Table 2, column 2, show that workers respond positively to the minimum wage, but this response is less pronounced for those with a better outside option. Figure 2, panel A, provides visual confirmation: both lines are positive and decreasing. Why?

The answer, we claim, lies in an "efficiency wage" logic rooted in the empirical observation that a minimum wage increase improves our workers' "inside option" (i.e., their current job's value) more than their outside option. Indeed, whereas the minimum wage has no significant effect on the OOI for either gender (see Table A.14, columns 3-4), the minimum wage has a much larger effect (by one order of magnitude) on pay in our workers' *current* job: see Table 1, column 2.[26] The small impact on the outside option is expected, as the minimum wage does not affect earnings during unemployment spells, and the probability of securing another minimum wage job after reemployment is relatively low for our workers.[27]

As the gap between inside and outside option widens, workers are incentivized to work harder to reduce the probability of termination, consistent with efficiency wage theory à la Shapiro & Stiglitz (1984) or Rebitzer & Taylor (1995). This explains why the lines in Figure 2, panel A, are positive. Moreover, this incentive effect is stronger for workers with worse outside options, who value job retention the most, and weaker for those with better outside options. This explains why the lines in Figure 2, panel A, are decreasing.

[26] A \$1 increase is associated with a 0.5%-1% change in OOI, indistinguishable from zero despite tight standard errors. By contrast, the effect on current pay is +5%.

[27] See Appendix E.3 for a more formal discussion.

Further evidence in favor of efficiency wage theory comes from a placebo test. Because supervisors occasionally turn over, there are periods during which some teams (departments within a store) are unsupervised.[28] During these periods, we can confirm that the empirical relationship between low productivity and termination becomes much weaker under no supervision.[29] When effort does not affect retention much, efficiency wage theory predicts that worker effort should not respond to a minimum wage increase, regardless of their outside options (because working harder does not increase the probability of retaining one's job). This prediction is borne out in the data: Table 2 column 6, and Figure 2 panel C show that, during periods of no supervision, workers of both genders do not become more productive following a minimum wage increase, regardless of their outside option. This placebo test shows that the outside option matters when, and only when, it is predicted to matter by efficiency wage theory.

Finally, Table 2 (columns 3-6) and Figure 2 (panels B and C) demonstrate that workers of both genders respond similarly to the outside option *even within monitoring regime* (supervised vs. unsupervised).This finding further supports the efficiency wage mechanism.

Evidence against alternative mechanisms The analysis has shown that men and women with similar outside options respond similarly to the minimum wage. Therefore, the gender differences observed in Table 1 (column 1) are fully explained by gender disparities in outside options, as documented in Figure 2, panel A.

This finding rules out several alternative explanations, unrelated to outside options, for why women have a stronger productivity response than men. These include gender-specific traits (e.g., differences in marginal cost of effort, innate job aptitude, risk aversion, propensity to reciprocate, or childcare constraints), post-minimum wage firm-level adjustments that might disproportionately affect women

[28] About 30% of our workers experience some periods without supervision. The likelihood that a worker is supervised is unrelated to the workers' outside options (Table A.4, column 4) and is not impacted by the minimum wage (Table A.6, column 5).

[29] See Table A.2, columns 3-6. Also, as expected, unsupervised workers are less productive than supervised workers: see Table A.4, column 3.

(e.g., reduced hours, increased monitoring), and "gendered" effects of demand shocks or product price adjustments. Each of these mechanisms would predict that men and women with similar outside options respond differently to the minimum wage – yet, the data do not support this.[30]

Finally, the powerful moderating role played by monitoring (Table 2, columns 3-4 vs. 5-6, and Figure 2, panels B vs. C) strongly points to the family of efficiency wage models where incentives are driven by the fear of termination (Shapiro & Stiglitz 1984; Rebitzer & Taylor 1995), as opposed to "gift exchange" models (Akerlof 1982; Akerlof & Yellen 1990) where, potentially, women might increase their effort more than men out of greater gratitude following a minimum wage increase. Because the evidence points to an efficiency wage model à la Shapiro & Stiglitz (1984); Rebitzer & Taylor (1995), we use it in the next section to evaluate the welfare effect of the minimum wage.

5 Welfare Effect of the Minimum Wage by Gender

This section quantifies the effect of the minimum wage on the welfare of our female vs. male workers. Several competing forces create gender differences in the welfare effect of the minimum wage. On the one hand, women benefit less than men because they work extra hard after a minimum wage increase (effort cost) but receive a similar pay increase. On the other hand, women benefit more than men because they are retained more. We provide a model within which we derive the formula for the welfare effect of increasing the minimum wage which boils down these countervailing effects to a single number. After calibrating this formula for our male and female workers separately, we find that, *ceteris paribus*, the minimum wage increases the welfare of women less than that of men.

[30] Out of an abundance of caution, in Appendix E.2, we offer additional evidence that rules out mechanisms other than the outside option.

5.1 Model

The model that follows is in the spirit of Rebitzer & Taylor (1995)'s efficiency wage model. A worker (in our empirical setting, a salesperson whose job is to interact with a customer) chooses effort under two incentives: the probability of being terminated, and the wage. The probability of termination is decreasing in worker effort. The expected wage is based on individual performance (in our setting, sales per hour) and is increasing in effort. By law, the wage cannot fall below the minimum wage. The fine details about the model are provided in Appendix F.1.

Primitives Worker effort is denoted by e and has cost $c(e)$. Worker performance (in our case, sales per hour) is a random variable $Y(e)$ that enjoys the strict monotone likelihood ratio property (MLRP) in e. Intuitively, the MLRP means that greater effort produces stochastically higher output.[31]

Consider any continuous nondecreasing compensation scheme $\overline{w}(\cdot)$ that transforms performance into pay. For example, $\overline{w}(Y) = b + cY$, where b represents the base salary and c the commission rate. Since in our firm the compensation scheme is set uniformly at the national level, in our model we cannot assume that the compensation scheme $\overline{w}(\cdot)$ is optimally adapted to the local parameters, including the minimum wage M. We assume, instead, that when a locality increases M, $\overline{w}$ does not change.[32] Thus, in a store that is subject to a local minimum wage M, the expected wage is:

$$w(e; M) = \mathbb{E}\left(\max\left[M, \overline{w}(Y(e))\right]\right). \tag{3}$$

The function $w(e; M)$ is bounded below by M and is nondecreasing in all its arguments.[33]

[31] The MLRP implies first-order stochastic dominance.

[32] This assumption is validated empirically in Table A.5, columns 3 and 4, where we show that when a locality increases M, base pay and commission rates in the store do not change.

[33] It is obviously nondecreasing in M. It is nondecreasing in e by stochastic dominance, because the function $\max[M, w(Y)]$ is nondecreasing in Y.

The worker's effort choice problem The worker's effort choice problem is:

$$V^E(M) = \max_e w(e;M) - c(e) + \frac{1}{(1+r)}\left[\pi(e)V^E(M) + (1-\pi(e))V^U(M)\right]. \quad (4)$$

Here, $V^E(M)$ represents the lifetime welfare of a worker who is currently employed by our firm. The numbers $r > 0$ and $V^U(M)$ represent, respectively, the discount rate and the lifetime value of becoming unemployed. The function $\pi(e)$ represents the probability of continued employment, which is assumed to be strictly increasing and continuously differentiable over $[0,1]$.[34]

To simplify the worker's problem, subtract the equation $[r/(1+r)]\, V^U(M) = u^U(M)$ from (4). We get:

$$V(M) = \max_e u(e;M) + \frac{1}{(1+r)}\pi(e)V(M), \quad (5)$$

where $V(M) = V^E(M) - V^U(M)$ represents the additional lifetime welfare of a worker who is currently employed by our firm relative to being unemployed, and

$$u(e;M) = w(e;M) - c(e) - u^U(M)$$

represents the flow value of employment, net of flow opportunity cost $u^U(M)$, of a worker who is currently employed and exerts effort e.

To ensure that the maximization problem in (5) is strictly concave in e, we assume $u_{ee} < 0$ and $\pi_{ee} \leq 0$. Concavity of u in e may be imparted to u by either of its components, w and c. For example, $u_{ee} < 0$ if the wage w is identically equal to the minimum wage, provided that the cost function is strictly convex in e. These assumptions guarantee that the worker's optimal effort $e^*(M)$ is the unique solution to the first-order conditions for problem (5).

The next lemma establishes that, for fixed M, $e^*(M)$ is strictly decreasing in the worker's outside option $V^U(M)$.

[34]Expressions (3) and (4) imply risk neutrality on the worker's part. This assumption is consistent with the fact that, empirically, worker response to the minimum wage is uncorrelated with the variance of the worker's outside option (details available on request).

Proposition 1. *For any given minimum wage level, the worker's optimal effort $e^*(M)$ is strictly decreasing in the worker's outside option.*

Proof. See Appendix F.1.1. ∎

This property is intuitive: in an efficiency wage model, the worker is motivated to exert effort by the fear of being terminated. When the consequences of being terminated improve, this fear factor attenuates, and effort decreases. Conversely, workers exert more effort when their outside option worsens.

5.2 Formula for worker welfare effect of the minimum wage

Next, we compute a formula for how the welfare of a generic worker changes as a function of M. In the following sections, the formula will be taken to the data separately for men and women, meaning that we do not need to *assume* that women and men have the same pay or retention schedules, or the same cost of effort. In fact, the worker's cost of effort happens to drop out of the formula, which is helpful because this function is unobservable.

First (and most insightful) step to compute the welfare formula Rewrite problem (5) as follows:

$$V(M) = u(e^*; M) + \frac{1}{1+r}\pi(e^*) V(M), \tag{6}$$

and rearrange (6) to get:

$$V(M) = \underbrace{\frac{1+r}{1+r-\pi(e^*)}}_{\text{dynamic factor}} \cdot \underbrace{u(e^*; M)}_{\text{static factor}}. \tag{7}$$

Intuitively, the lifetime welfare $V(M)$ of an employed relative to an unemployed worker is the product of two factors. The second factor is the flow difference between the employed and unemployed state; we call this a static factor. The first factor is a capitalization term that converts flows into stocks, and depends on the probability

π that the worker remains employed; we call this a dynamic factor. Both factors depend on the effort level e^* chosen by the worker.

Differentiate with respect to M and use the envelope (or first-order) condition for problem (5) to get:

$$\frac{dV(M)}{dM} = \left[\frac{1+r}{1+r-\pi(e^*)}\right] u_M(e^*; M). \tag{8}$$

The calculation is presented in Appendix F.1.2. This expression represents the formula for the change in $V(M)$ due to a change in the minimum wage M. A significant empirical advantage is that formula (8) does not depend on $c(e)$, the worker's cost function in her current employment: conveniently, this term dropped out due to an envelope condition whose economic content is discussed next.

Intuition for formula (8) To get an intuition for expression (8), observe that this formula is simply the partial derivative of $V(M)$ with respect to M, without accounting for the changes in the worker's optimal effort $e^*(M)$. Technically, the reason why these changes drop out of the algebra is the envelope (or first-order) condition for problem (5). Intuitively, the reason why the change in effort does not affect the worker's welfare is that by definition the effort level e^* maximizes the worker's lifetime welfare $V(M)$, so any small change in effort around the baseline level e^* only has second-order effects on $V(M)$.

Even more intuitively, the worker sets e^* to optimally balance two countervailing effects: increasing e increases the probability of retention, and hence the first (dynamic) factor in (7); and it decreases the second (static) factor $u(e; M)$ because, as is apparent from inspecting equation (5), the function $\pi(\cdot)$ being strictly increasing motivates the worker to exert *excessive effort* relative to what is justified solely by static incentives. At the optimal effort choice e^*, changing the worker's effort causes these two effects to move in opposite directions in a way that *exactly offsets* each other. As a result, the only effect on welfare is the one directly caused by the variable M which is not under the worker's control. In other words, when M increases, the envelope condition implies that any welfare change that is mediated by a shift in effort (e.g., increased effort cost, increased wage due to more effort, and increased probability of retention due to a change in effort) has no welfare implications. Only the direct effect of the minimum wage on pay (and, potentially, on the outside option) matters for the welfare calculation.

Expression (8) is deceptively elegant, but its empirical implementation is delicate because it requires computing a counterfactual. Indeed, e^* represents the counterfactual effort that the worker would have exerted *in the absence of a change in the minimum wage*. We will deal with this empirical challenge in the next subsection.

Second step: the complete welfare formula We are interested in the change in lifetime welfare, *inclusive of post-separation future*, of a current employee at our firm. Hence, using the notation of equation (4), we are interested in:

$$\frac{dV^E(M)}{dM} = \frac{dV(M)}{dM} + \frac{dV^U(M)}{dM}.$$

After some algebra presented in Appendix F.1.3, we get:

$$\frac{dV^E(M)}{dM} = \left[\frac{1+r}{1+r-\pi(e^*)}\right]\left[w_M(e^*;M) + \frac{[1-\pi(e^*)]}{r}u_M^U(M)\right]. \tag{9}$$

Naturally, this formula reduces to (8) when $u_M^U(M) = 0$, i.e., when the post-separation future does not depend on the minimum wage.

5.3 Calibrating the welfare formula, by gender

We estimate the following version of the welfare formula (9) *separately by gender*:

$$\frac{dV^E(M)}{dM} = \left[\frac{1+r}{1+r-\pi\left(e^*_{t-1}\right)}\right]\left[w_M\left(e^*_{t-1};M_t\right) + \frac{\left[1-\pi\left(e^*_{t-1}\right)\right]}{r}u_M^U(M_t)\right]. \tag{10}$$

Compared with formula (9), the terms involving e^* are lagged relative to M. This is because, in formula (9), e^* represents the counterfactual effort that the worker would have exerted in the absence of a change in the minimum wage. Empirically, using e_t^* would be incorrect whenever the minimum wage changes at t, because contemporaneous effort is endogenous to the prevailing minimum wage M_t. Therefore, we use e_{t-1}^* to proxy for the *counterfactual effort* that the worker would have exerted had the minimum wage remained at level M_{t-1}.

Formula (10) involves calibrated parameters and estimates. First, for $\pi\left(e_{t-1}^*\right)$ we plug in the average retention rate in month $t-1$, by gender. Hence, within either gender, $\pi\left(e_{t-1}^*\right)$ is the fraction of workers who were retained at t among those who were employed at time $t-1$. Second, we set the monthly discount rate r to 2.5%. This level of discounting is larger than normally assumed in welfare analyses, but is in line with field-experimental evidence on the personal discount factor.[35] Third, we set $u_M^U(M_t)$ to zero. This choice makes sense because the minimum wage has a negligible effect on the worker's future welfare (see Table A.14, discussed in Appendix E.3). We present the robustness of our results to alternative calibrations of π, r and to allowing $u_M^U(M_t)$ to be non-zero in Appendix F.2.

Using the above calibrations, the right-hand side of welfare formula (10) becomes:

$$\frac{1+r}{1+r-\pi\left(e_{t-1}^*\right)} \cdot w_M\left(e_{t-1}^*; M_t\right). \tag{11}$$

To estimate this expression, we proceed in three steps. As a first step, we create the following variable for each worker i:

$$w\left(e_{i,t-1}^*; M_{jt}\right) = \max\left[M_{jt}, \overline{w}\left(Y\left(e_{i,t-1}^*\right)\right)\right], \tag{12}$$

which we refer to as i's "synthetic pay per hour." This variable is the empirical counterpart to expression (3) in the theory. This variable involves a counterfactual: it is the hourly pay that the company would have paid worker i in store j in a month t when the minimum wage increases, had the worker made the same sales as in the

[35]Yearly personal discount rates are estimated at 28% in a representative sample of the Danish population (Harrison et al., 2002, p. 1612) and as large as 35% for enlisted military personnel (Warner & Pleeter, 2001, p. 49).

pre-increase regime.[36] As a second step, we use (12) to create the following variable:

$$\frac{1+r}{1+r-\pi\left(e_{t-1}^{*}\right)}\cdot w\left(e_{i,t-1}^{*};M_{jt}\right), \tag{13}$$

which we refer to as i's "discounted synthetic pay per hour." The third step is to estimate (11) by regressing (13) on M_{jt} (note that (11) has the subscript M but (13) does not) using specification (1).

5.4 Results: Welfare effect of the minimum wage, by gender

The estimates of the welfare effect of the minimum wage (expression 11) by gender are presented in Table 1, column 6. We find that a \$1 increase in the minimum wage increases men's lifetime welfare by 10.4% and women's lifetime welfare by 5.8%, where the percentage is expressed relative to the mean of the dependent variable (synthetic pay per hour). This difference, which reveals that men benefit twice as much as women from the minimum wage, is statistically significant at the 5% level.[37]

Table A.15 shows that these welfare results are robust to using different values for r and calibrations of π. They are also robust to relaxing the assumption that the minimum wage does not impact the outside option, i.e., to allowing $u_M^U(M_{jt})$ to be positive in equation (10). Finally, Table A.9, column 4, explores differences between departments A and B and shows that, in both departments, the welfare effect of the minimum wage is larger for men than for women. These robustness results are discussed in Appendix F.2.[38]

Expression (11) reveals that, when effort is endogenous, the contemporaneous increase in pay (i.e., $w_M\left(e_t^*;M_t\right)$ where, as opposed to (11), effort is not lagged)

[36] Specifically, $w\left(e_{i,t-1}^{*};M_{jt}\right)$ is computed as the total pay in the pre-increase regime. plus any top-up if that amount is below the new minimum wage.

[37] It may be worth emphasizing that these coefficients do not express the increase in welfare *as a percentage of baseline welfare*. Indeed, the latter is unobservable: the theory does not provide a method for recovering welfare *levels* due to the presence of the unobservable term $c(e)$.

[38] Table A.15, column 1, reports the *wrong* estimates of the welfare effects where, *incorrectly*, contemporaneous pay replaces synthetic pay in formula (11): the estimates are very different from Table 1, column 6. Hence, using synthetic pay is essential to estimate the welfare effects correctly.

is not necessarily a good measure of welfare. So, the fact that the minimum wage happens to increase women's and men's pay by nearly the same amount (Table 1, column 1) is uninformative about welfare. This discrepancy between pay and welfare is due to the fact that women earn their pay boost through a greater effort response. In sum, pay is not the same as welfare.

6 *Non-Ceteris-Paribus* Impact of the Minimum Wage by Gender

So far, we have compared female and male workers *ceteris paribus*, i.e., in the same working conditions within the firm (although women's outside option could be, and often was, less favorable). We now turn to the *non-ceteris-paribus* estimates that take on board the fact that, within our firm, women are disproportionately represented in the low-paying department relative to men. Based on this same fact (that women tend to be overrepresented in lower-paying positions) across the entire labor market, Caliendo & Wittbrodt (2022); Blau et al. (2023); Paul-Delvaux (2023) find that, in several countries, a higher minimum wage reduces the gender pay gap. We find the same result within our firm and, in addition, offer novel findings about welfare.

To get at the *non-ceteris paribus* impact of the minimum wage, we estimate a variant of specification (1) replacing department×store and worker fixed effects with store fixed effects only. The results are presented in Table A.16. Across the firm, total pay per hour increases more sharply for women than for men when the minimum wage increases – this is expected due to the women's disadvantaged positions across the firm, and contrasts with the equal pay increase observed in the *ceteris paribus* analysis.[39] The results regarding productivity and retention remain consistent with our previous findings: the women's productivity response is more pronounced than that of men, and their retention increases by more. In this *non-ceteris-paribus* specification, therefore, women benefit more than men from the minimum wage not only

[39]Women's regular pay still increases more than that of men, but women are now equally (rather than less) likely to be "topped-up" after the minimum wage increase.

in terms of retention but, also, in terms of pay. Not surprisingly then, in the *non-ceteris-paribus* welfare analysis, the minimum wage benefits women more than men: welfare increases by 7.4% for women vs. 4.0% for men. Thus, once we move beyond *ceteris paribus*, the welfare gap flips: now, women benefit more than men from the minimum wage. This flip reflects a mechanical effect: because women tend to work in the low-pay department, the minimum wage tops up the women's wages more often. The gender wage gap is a precursor of this mechanical effect.

In the US economy, the gender wage gap is even larger than within our firm (ten to twenty%, depending on the estimates vs. 4.5% in our firm),[40] implying that the mechanical effect of the minimum wage on the welfare of female workers must be even larger than in our sample. This suggests that, unless the endogenous effort response is much larger in the US economy than in our analysis, the minimum wage must increase the welfare of US female workers more so than male workers.

7 Conclusion

This paper has examined an important fairness question: when the minimum wage increases, do male and female workers benefit equally? To address this question empirically, we have studied the differential effect of the minimum wage on pay and welfare by gender among more than 10,000 hourly paid salespeople whose pay is partly based on performance, and who are employed by a large US retailer that operates more than 2,000 stores. The sample is broadly representative of U.S. hourly-paid workers – who make up nearly 60% of the workforce – in terms of pay, termination rates, and gender composition.

We have shown that, in *ceteris-paribus* working conditions, women benefit less from the minimum wage in welfare terms than men despite experiencing a similar

[40]From 2012 to 2015, CPS data shows that women's median hourly pay is 12% lower than men's. Studies attribute much of this gap to women sorting into lower-paying establishments and occupations (e.g., Bayard et al., 2003; Goldin et al., 2017; Barth et al., 2021). In our context, men and women share the same occupation and establishment – resulting in a smaller gender pay gap than at the national level – but the gap persists because women are assigned to lower-paying departments.

pay increase. Hence, the paper's first contribution: to demonstrate empirically that a "facially neutral" improvement (minimum wage increase) in *ceteris-paribus* working conditions can lead to differential worker response by gender, and to a disparate welfare impact. This disparity, we argue, is due to baseline "systemic disparities" that are not under the employer's control – in our case, gender differences in the outside option. Extrapolating from our minimum wage setting, this paper demonstrates empirically (for the first time, to our knowledge) that the same improvement in the pay scheme can have disparate welfare impacts on two identically-situated co-workers who differ only in their outside options.

A secondary yet policy-relevant observation concerns the welfare effect of the minimum wage across the entire economy. If women earn less than men (gender wage gap), then, in the absence of an endogenous effort response, an increase in the minimum wage mechanically benefits women more than men. Because the gender gap is even larger in the US economy than in our *non-ceteris-paribus* analysis, this mechanical effect must be even larger in the US economy, compared to our estimates. This observation implies that, unless the endogenous effort response is substantially greater in the US economy than in our estimates, the minimum wage likely increases the welfare of US female workers more than that of US male workers. These *non-ceteris-paribus* estimates support the notion that the minimum wage acts as a force for gender equalization *even in welfare terms* because female workers typically occupy lower-paid positions. Importantly, however, our findings show that among similarly situated workers, a higher minimum wage disproportionately benefits men in welfare terms.

This paper's approach and empirical findings have emphasized that, when effort is endogenous, differences in pay do not necessarily track welfare differences. Indeed, empirically, we found that, *ceteris paribus*, the minimum wage benefits men strictly more than women in welfare terms – but not in pay terms. The cautionary point that "pay is not welfare" speaks to the growing literature on the gender pay gap.

References

Ahammer, A., Fahn, M., & Stiftinger, F. (2023). Outside Options and Worker Motivation. *Working Paper*.

Aiello, D., Baker, S. R., Balyuk, T., Di Maggio, M., Johnson, M. J., & Kotter, J. D. (2024). The Effects of Cryptocurrency Wealth on Household Consumption and Investment. *Working Paper*.

Akerlof, G. A. (1982). Labor Contracts as Partial Gift Exchange. *The Quarterly Journal of Economics*, *97*(4), 543–569.

Akerlof, G. A., & Yellen, J. L. (1990). The Fair Wage-Effort Hypothesis and Unemployment. *The Quarterly Journal of Economics*, *105*(2), 255–283.

Allegretto, S. A., Dube, A., & Reich, M. (2011). Do Minimum Wages Really Reduce Teen Employment? Accounting for Heterogeneity and Selectivity in State Panel Data. *Industrial Relations: A Journal of Economy and Society*, *50*(2), 205–240.

Allegretto, S. A., Dube, A., Reich, M., & Zipperer, B. (2017). Credible Research Resigns for Minimum Wage Studies: A Response to Neumark, Salas, and Wascher. *Industrial and Labor Relations Review*, *70*(3), 559–592.

Antecol, H., Bedard, K., & Stearns, J. (2018). Equal but Inequitable: Who Benefits from Gender-Neutral Tenure Clock Stopping Policies? *American Economic Review*, *108*(9), 2420–2441.

Barth, E., Kerr, S. P., & Olivetti, C. (2021). The Dynamics of Gender Earnings Differentials: Evidence from Establishment Data. *European Economic Review*, *134*(103713).

Bayard, K., Hellerstein, J., Neumark, D., & Troske, K. (2003). New Evidence on Sex Segregation and Sex Differences in Wages from Matched Employee-Employer Data. *Journal of Labor Economics*, *21*(4), 887–922.

Biasi, B., & Sarsons, H. (2022). Flexible Wages, Bargaining, and the Gender Gap. *Quarterly Journal of Economics*, *137*(1), 215–266.

Blau, F. D., Cohen, I., Comey, M. L., Kahn, L., & Boboshko, N. (2023). The Minimum Wage and Inequality Between Groups. *Working Paper*.

Bolotnyy, V., & Emanuel, N. (2022). Why do Women Earn Less than Men? Evidence from Bus and Train Operators. *Journal of Labor Economics*, *40*(2), 283–323.

Caldwell, S., & Danieli, O. (2023). Outside Options in the Labor Market. *Review of Economic Studies*.

Caliendo, M., & Wittbrodt, L. (2022). Did the Minimum Wage Reduce the Gender Wage Gap in Germany? *Labour Economics*, *78*, 102228.

Card, D., & Krueger, A. B. (2000). Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania: Reply. *American Economic Review*, *90*(5), 1397–1420.

Coviello, D., Deserranno, E., & Persico, N. (2022). Minimum Wage and Individual Worker Productivity: Evidence from a Large US Retailer. *Journal of Political Economy*, *130*(9), 2315–2360.

Derenoncourt, E., & Montialoux, C. (2021). Minimum Wages and Racial Inequality. *Quarterly Journal of Economics*, *136*(1), 169–228.

Di Maggio, M., Williams, E., & Katz, J. (2023). Buy Now, Pay Later Credit: User Characteristics and Effects on Spending Patterns. *Working Paper*.

Dube, A., Giuliano, L., & Leonard, J. (2019). Fairness and Frictions: The Impact of Unequal Raises on Quit Behavior. *American Economic Review*, *109*(2), 620–663.

Dube, A., Lester, T. W., & Reich, M. (2010). Minimum Wage Effects Across State Borders: Estimates Using Contiguous Counties. *The Review of Economics and Statistics*, *92*(4), 945–964.

Dube, A., Lester, T. W., & Reich, M. (2016). Minimum Wage Shocks, Employment Flows, and Labor Market Frictions. *Journal of Labor Economics*, *34*(3), 663–704.

Flinn, C. J. (2006). Minimum Wage Effects on Labor Market Outcomes under Search, Matching, and Endogenous Contact Rates. *Econometrica*, *74*(4), 1013–1062.

Gittings, R. K., & Schmutte, I. M. (2016). Getting Handcuffs on an Octopus: Minimum Wages, Employment, and Turnover. *Industrial and Labor Relations Review*, *69*(5), 1133–1170.

Giuliano, L. (2013). Minimum Wage Effects on Employment, Substitution, and the Teenage Labor Supply: Evidence from Personnel Data. *Journal of Labor Economics*, *31*(1), 155–194.

Goldin, C., Kerr, S. P., Olivetti, C., & Barth, E. (2017). The Expanding Gender Earnings Gap: Evidence from the LEHD-2000 Census. *American Economic Review*, *107*(5), 110–114.

Harrison, G. W., Lau, M. I., & Williams, M. B. (2002). Estimating Individual Discount Rates in Denmark: A Field Experiment. *American Economic Review*, *92*(5), 1606–1617.

Ku, H. (2022). Does Minimum Wage Increase Labor Productivity? Evidence from Piece-rate Workers. *Journal of Labor Economics*, *40* (2), 325–359.

Lazear, E. P., Shaw, K. L., & Stanton, C. (2016). Making Do with Less: Working Harder During Recessions. *Journal of Labor Economics*, *34* (S1), S333–S360.

Le Barbanchon, T., Rathelot, R., & Roulet, A. (2021). Gender Differences in Job Search: Trading off Commute against Wage. *Quarterly Journal of Economics*, *136* (1), 381–426.

Lusher, L., Schnorr, G. C., & Taylor, R. L. (2022). Unemployment Insurance as a Worker Indiscipline Device? Evidence from Scanner Data. *American Economic Journal: Applied Economics*, *14* (2), 285–319.

Neumark, D., Salas, J. I., & Wascher, W. (2014). Revisiting the Minimum Wage—Employment Debate: Throwing Out the Baby with the Bathwater? *Industrial and Labor Relations Review*, *67* (3_suppl), 608–648.

Paul-Delvaux, L. (2023). Minimum Wage and Gender Gaps: Evidence from Morocco. *Working Paper*.

Rebitzer, J. B., & Taylor, L. J. (1995). The Consequences of Minimum Wage Laws Some New Theoretical Ideas. *Journal of Public Economics*, *56* (2), 245–255.

Ruffini, K. (2024). Worker Earnings, Service Quality, and Firm Profitability: Evidence from Nursing Homes and Minimum Wage Reforms. *Review of Economics and Statistics*, *106* (6), 1477–1494.

Schubert, G., Stansbury, A., & Taska, B. (2024). Employer Concentration and Outside Options. *Working Paper*.

Shapiro, C., & Stiglitz, J. E. (1984). Equilibrium Unemployment as a Worker Discipline Device. *American Economic Review*, *74* (3), 433–444.

US Bureau of Labor Statistics. (2015a). *Characteristics of Minimum Wage Workers.* https://www.bls.gov/opub/reports/minimum-wage/2015/home.htm.

US Bureau of Labor Statistics. (2015b). *Labor Force Characteristics by Race and Ethnicity.* https://www.bls.gov/opub/reports/race-and-ethnicity/2015/.

US Census Bureau. (2020). *Retail Jobs Among the Most Common Occupations.* https://www.census.gov/library/stories/2020/09/profile-of-the-retail-workforce.html.

Warner, J. T., & Pleeter, S. (2001). The Personal Discount Rate: Evidence from Military Downsizing Programs. *American Economic Review*, *91* (1), 33–53.

Online Appendix

A Appendix Figures and Tables

Figure A.1: Distribution of Total Pay per Hour by Gender

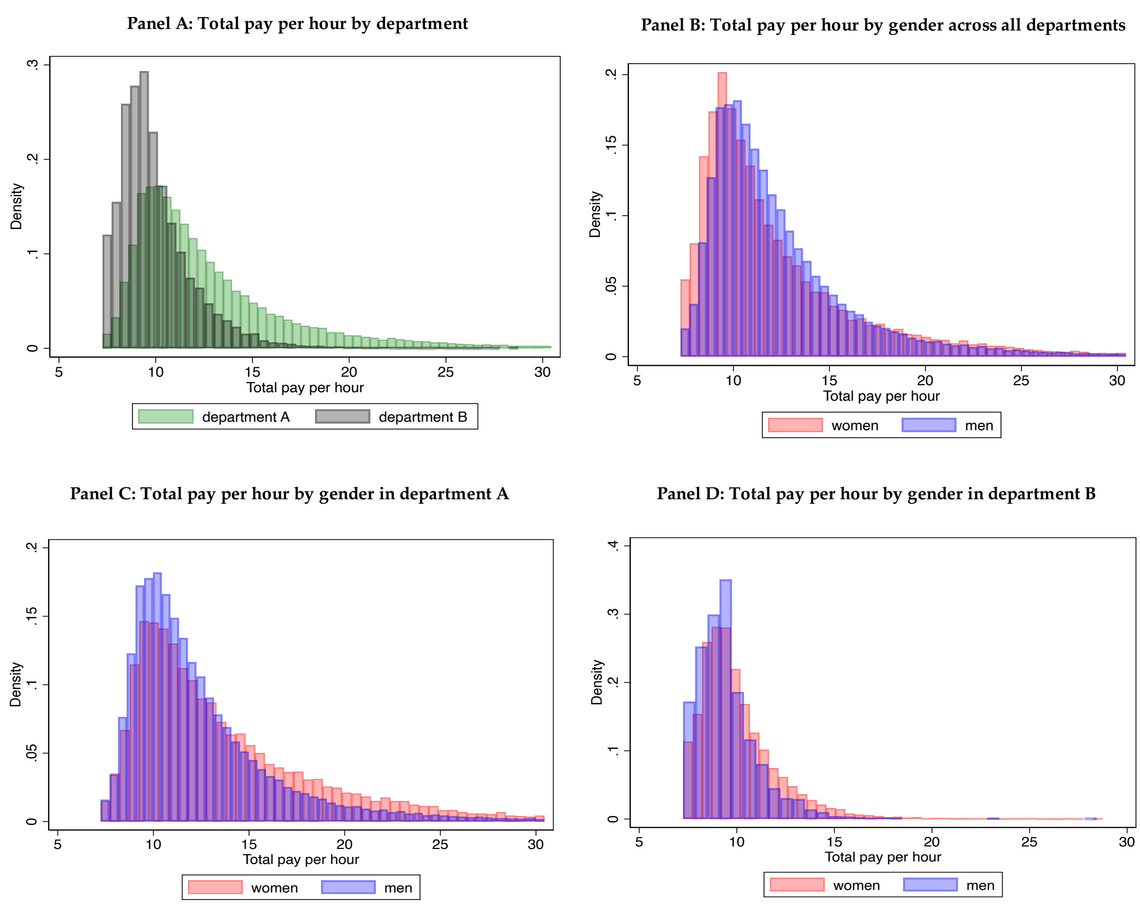

Notes: This figure presents the distribution of total pay per hour by department (panel A), by gender (panel B), by gender within department A (panel C), by gender within department B (panel D). For visual clarity, we remove observations in which total pay per hour is below the minimum wage (0.6% of the sample).

Figure A.2: Distribution of Sales per Hour by Gender

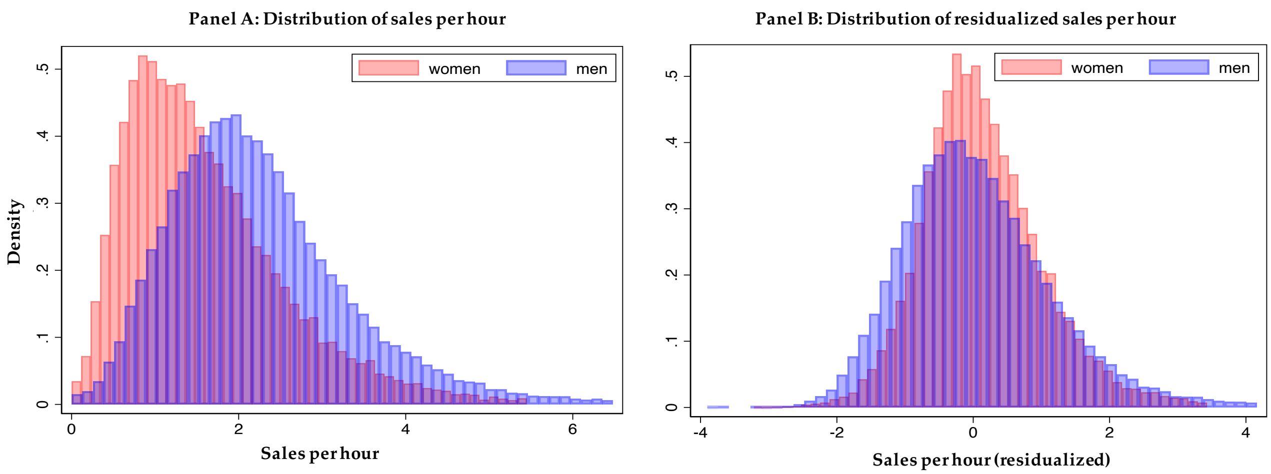

Notes: The figure on the left shows the distribution of sales per hour by gender, while the figure on the right displays the residuals from a regression of sales per hour on the worker's department and store, also by gender. Each observation represents a worker-month. For visual clarity, the top 1% of observations are excluded in the left panel, and residuals below -4 are excluded in the right panel (1% of the sample).

Table A.1: Summary Statistics by Gender

	(1)	(2)	(3)	(4)	(5)	(6)
Sample		**Women**			**Men**	
	Mean	S.D.	Median	Mean	S.D.	Median
# observations [# workers-months]		76,336			141,410	
Worker characteristics, tenure and termination						
Age (in years)	36.37	16.74	28.20	35.54	17.15	27
Home-work distance (in km)	9.172	13.860	6.660	10.100	19.250	7.104
Tenure (in months)	57.74	73.09	27	44.16	59.65	22
Terminated = {0, 1}	0.041	0.199	-	0.048	0.214	-
Department allocation						
Department A (vs. department B)	0.601	0.490	-	0.977	0.148	-
Compensation structure						
Base hourly rate (in $)	6.097	1.235	6	6.144	1.112	6
Commission rate (in %)	2.871	1.763	2.435	2.441	1.447	2.065
Pay: total, regular and top-up						
Total pay per hour (in $)	12.14	4.177	10.78	12.34	3.786	11.27
Regular (fixed+variable) pay per hour (in $)	11.82	4.727	10.58	12.17	4.126	11.17
MinW top-up per hour (in $)	0.319	2.069	0.048	0.175	1.540	0
Minimum wage top-up frequency						
MinW top-up at least one week of the month = {0, 1}	0.534	0.499	-	0.359	0.480	-
MinW top-up all weeks of the month = {0, 1}	0.049	0.215	-	0.021	0.142	-
Number of weeks with minW top-up (0 to 4)	1.019	1.197	-	0.595	0.954	-
Hours worked						
Number of hours per week	27.62	4.847	25	27.570	4.802	25
Part-time = {0, 1}	0.620	0.485	-	0.593	0.491	-
Productivity						
Sales per hour (shrouded units)	1.665	1.353	1.411	2.311	1.477	2.094

Notes: This table reports summary statistics for women and men separately, across all departments. Each observation represents a worker-month. "Terminated" is a dummy variable that equals one if the worker is terminated that month (i.e., not retained). "Supervised" is a dummy variable that equals one if the worker is assigned a direct supervisor. "Base hourly rate" is the monthly base rate per hour worked (in $ per hour). "Commission rate" is the earnings from commissions divided by sales (in %). "Total pay per hour" is the monthly total pay (in $ per hour). "Regular pay per hour" is the total amount earned from the base hourly rate and variable pay (commission rate × sales per hour), without the top-up. "MinW top-up per hour" is the monthly total minimum wage adjustment paid by the company to the worker (in $ per hour). "Number of weeks with minW top-up" is the number of weeks over the months in which the worker is paid a positive minimum wage adjustement by the firm (0 to 4). "Sales per hour" are the sales per hour rescaled by a factor between 1/50 and 1/150 relative to its $ value.

Table A.2: Correlates of Termination

	(1)	(2)	(3)	(4)	(5)	(6)
Dep.Var.	Terminated					
Sample	*Full sample*		*Supervised*		*Unsupervised*	
Sales per hour/100	-0.215***	-0.284***	-0.387***	-0.449***	0.019	-0.011
	(0.076)	(0.089)	(0.085)	(0.091)	(0.135)	(0.161)
Sales per hour/100 × Woman		0.002		0.002		0.000
		(0.001)		(0.001)		(0.002)
Observations	217,746	217,746	162,837	162,837	41,644	41,644
Mean Dep.Var.	0.046	0.046	0.046	0.046	0.046	0.046

Notes: Each observation represents a worker-month. All the regressions include store×department fixed effects and pair×month fixed effects. Even columns also control for a dummy variable for being a woman. Standard errors are two-way clustered at the state level and at the border-segment level. *** p<0.01, ** p<0.05, * p<0.1. "Terminated" is a dummy variable that equals one if the worker is terminated that month (i.e., not retained). "Sales per hour" are the sales per hour rescaled by a factor between 1/50 and 1/150 relative to its $ value, divided by 100 and lagged by one month. Columns 3–4 (resp., columns 5–6) restrict the sample to workers with (resp., without) a direct supervisor assigned during that month.

Table A.3: Summary Statistics by Gender and Department

	(1)	(2)	(3)	(4)	(5)	(6)
Sample	**Women**			**Men**		
	Mean	S.D.	Median	Mean	S.D.	Median
Panel A: Department A						
# observations [# workers-months]		45,878			138,226	
Base hourly rate (in $)	5.809	1.323	6	6.146	1.104	6
Commission rate (in %)	2.898	1.766	2.375	2.437	1.443	2.060
Total pay per hour (in $)	13.57	4.615	12.17	12.41	3.796	11.33
Panel B: Department B						
# observations [# workers-months]		30,458			3,184	
Base hourly rate (in $)	6.523	0.944	6.500	6.041	1.386	6.500
Commission rate (in %)	2.836	1.758	2.528	2.629	1.592	2.326
Total pay per hour (in $)	10.03	2.074	9.549	9.477	1.651	9.209

Notes: This table reports summary statistics for women and men separately in different samples: workers in department A (panel A) and workers in department B (panel B). Each observation represents a worker-month. "Base hourly rate" is the monthly base rate per hour worked (in $ per hour). "Commission rate" is the earnings from commissions divided by sales (in %).

Table A.4: Correlates of Sales per Hour

	(1)	(2)	(3)	(4)
Dep.Var.		Sales per hour		Supervised
Department A	0.683***			
	(0.085)			
Outside option (OOI)		-0.007***		-0.002
		(0.002)		(0.001)
Supervised			0.088***	
			(0.027)	
Observations	217,746	212,443	205,247	205,007
Mean Dep.Var.	2.085	2.085	2.085	0.793

Notes: Each observation represents a worker-month. All regressions include worker fixed effects and pair×month fixed effects. Column 1 also includes store fixed effects, while columns 2 and 3 include store×department fixed effects. Standard errors are two-way clustered at the state and border-segment level. *** p<0.01, ** p<0.05, * p<0.1. OOI is the outside option index (see Appendix D.1 for details on how the OOI variable is constructed). "Supervised" is an indicator for whether the worker has a direct supervisor assigned to them. The variable is missing for some observations, explaining why the sample size drops in columns 3 and 4.

Table A.5: Impact of the Minimum Wage on Top-up, Compensation Scheme and Involuntary Terminations by Gender (Ceteris-Paribus Analysis)

	(1)	(2)	(3)	(4)	(5)
	Top-up		**Compensation scheme**		**Termination**
Dep.Var.	MinW top-up at least one week of the month	Number of weeks with minW top-up	Base hourly rate	Commission rate	Involuntary terminated
MinW	0.189***	0.539***	-0.093	0.044	0.003
	(0.013)	(0.030)	(0.061)	(0.031)	(0.005)
MinW × Woman	-0.063***	-0.125***	0.032	0.017	-0.014***
	(0.015)	(0.029)	(0.108)	(0.030)	(0.003)
Observations	215,558	215,558	215,558	192,016	217,746
Mean Dep.Var.	0.423	0.743	6.128	2.583	0.018

Notes: Each observation represents a worker-month. All regressions include pair×month fixed effects, worker fixed effects and control for MinW×department. Standard errors are two-way clustered at the state level and at the border-segment level. *** p<0.01, ** p<0.05, * p<0.1. "Number of weeks with minW top-up" is the number of weeks over the months in which the worker is paid a positive minimum wage adjustement by the firm (0 to 4). "Base hourly rate" is the monthly base rate per hour worked (in $ per hour). "Commission rate" is the earnings from commissions divided by sales (in %). The value is missing for workers with zero sales per hour (hence, the smaller sample size). "Involuntary termination" is a dummy variable that equals one if the worker is terminated that month and the termination is categorized as "non-voluntary."

Table A.6: Impact of the Minimum Wage on Firm Organizational Adjustments by Gender (Ceteris-Paribus Analysis)

	(1)	(2)	(3)	(4)	(5)
	Hours		**Moves**		**Monitoring**
Dep.Var.	Hours per week	Part-time worker (worse shifts)	Move to high-pay department within same store	Move to another store	Supervised
MinW	0.273	-0.021	0.001	0.005	-0.089
	(0.255)	(0.017)	(0.001)	(0.008)	(0.064)
MinW × Woman	0.065	-0.024	0.001	-0.004	0.019
	(0.056)	(0.020)	(0.001)	(0.004)	(0.017)
Observations	217,746	217,746	217,746	217,746	205,247
Mean Dep.Var.	27.590	0.603	0.084	0.086	0.793

Notes: Each observation represents a worker-month. All regressions include pair×month fixed effects, worker fixed effects and control for MinW×department. Regressions in columns 1 and 2 also include store×department fixed effects. The regression in column 3 includes store fixed effects. Standard errors are two-way clustered at the state level and at the border-segment level. ***p<0.01, ** p<0.05, * p<0.1. "Hours per week" is the average number of hours worked in a week. "Move to high-pay department within same store" is a dummy variable for whether a worker moved from the low- to the high-pay department (B to A) within the same store. "Move to another store" is a dummy variable for whether a worker moved to another store, regardless of whether she/he moved to the same or a different department. "Assigned a supervisor" is an indicator for whether the worker is supervised, i.e., assigned a direct supervisor. The variable is missing for some observations, explaining why the sample size drops in the last column.

Table A.7: Impact of the Minimum Wage on the Termination Rule by Gender (Ceteris-Paribus Analysis)

Dep.Var.	Retained
MinW	-0.004
	(0.006)
MinW × Woman	0.021***
	(0.005)
Sales per hour	0.011***
	(0.001)
Sales per hour × Woman	-0.002***
	(0.001)
MinW × Sales per hour	-0.000
	(0.001)
MinW × Sales per hour × Woman	-0.001
	(0.002)
Observations	217,746
Mean Dep.Var.	0.954

Notes: Each observation represents a worker-month. The regression includes store×department fixed effects, pair×month fixed effects, worker fixed effects and controls for MinW×department and sales per hour×department. Standard errors are two-way clustered at the state and border-segment level. *** p<0.01, ** p<0.05, * p<0.1.

Table A.8: Impact of the Minimum Wage on Pay, Productivity, and Retention by Gender, with Additional Controls (Ceteris-Paribus Analysis)

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Dep.Var. = Sales per Hour						
MinW		0.046	0.048	0.060	0.059	0.055
		(0.041)	(0.041)	(0.040)	(0.041)	(0.041)
MinW × Woman	0.058**	0.047**	0.040**	0.052***	0.052***	0.068***
	(0.024)	(0.019)	(0.018)	(0.015)	(0.014)	(0.015)
Observations	217,500	217,746	217,746	217,746	212,427	217,746
Mean Dep.Var.	2.085	2.085	2.085	2.085	2.085	2.085
p-value MinW+MinW×Woman		0.086	0.097	0.024	0.020	0.007
Panel B: Dep.Var. = Total Pay per Hour						
MinW		0.601***	0.572***	0.553***	0.535***	0.547***
		(0.124)	(0.130)	(0.124)	(0.125)	(0.125)
MinW × Woman	0.124	0.117	0.108	0.087	0.083	0.103
	(0.138)	(0.155)	(0.164)	(0.164)	(0.172)	(0.147)
Observations	215,312	215,558	215,558	215,558	210,418	215,558
Mean Dep.Var.	12.271	12.271	12.271	12.271	12.271	12.271
p-value MinW+MinW×Woman		0.010	0.021	0.024	0.035	0.018
Panel C: Dep.Var. = Retained						
MinW		-0.011*	-0.007	-0.003	-0.003	-0.004
		(0.005)	(0.006)	(0.005)	(0.006)	(0.005)
MinW × Woman	0.020***	0.016***	0.016***	0.019***	0.020***	0.018***
	(0.004)	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)
Observations	217,500	217,746	217,746	217,746	212,427	217,746
Mean Dep.Var.	0.954	0.954	0.954	0.954	0.954	0.954
p-value MinW+MinW×Woman		0.361	0.088	0.018	0.023	0.070
<u>Controls in the regression:</u>						
Department×store×month FE	✓					
Tenure (above median) & MinW×Tenure		✓				
Age (above median) & MinW×Age			✓			
Childbearing age & MinW×Childbearing age				✓		
Home-work distance & MinW×Home-work distance					✓	
Share of women & MinW×Share of women						✓

Notes: Each observation represents a worker-month. All regressions include store×department fixed effects, worker fixed effects, pair×month fixed effects, MinW×department and the extra controls indicated at the bottom of the table. Standard errors are two-way clustered at the state and border-segment level. *** p<0.01, ** p<0.05, * p<0.1. The dependent variables are in the panel headings. The independent variables are de-meaned such that the coefficient for "MinW" picks up the effect of the minimum wage for men when the variables are equal to the sample mean, and the results are comparable across columns.

Table A.9: Impact of the Minimum Wage on Pay, Productivity, and Retention by Gender and Department (Ceteris-Paribus Analysis)

	(1) Productivity	(2) Pay	(3) Retention	(4) Welfare
Dep.Var.	Sales per hour	Total pay per hour	Retained	Discounted synthetic pay per hour
Impact of MinW in Department A				
MinW × Department A	0.046	0.562***	-0.002	16.430***
	(0.043)	(0.154)	(0.005)	(3.449)
MinW × Woman × Department A	0.051***	0.077	0.019***	-9.231**
	(0.018)	(0.162)	(0.004)	(3.958)
Impact of MinW in Department B				
MinW × Department B	0.106**	0.400*	-0.007	39.000***
	(0.044)	(0.201)	(0.009)	(10.090)
MinW × Woman × Department B	0.057*	0.078	0.022***	-9.495*
	(0.032)	(0.120)	(0.005)	(4.723)
Observations	217,746	215,558	217,746	197,333
Mean Dep.Var.	2.085	12.271	0.954	195.876
Effects in Department A				
Effect MinW for Men (%)	2.0%	4.5%	-0.2%	8.5%
Effect MinW for Women (%)	4.8%	4.7%	1.8%	3.2%
Effects in Department B				
Effect MinW for Men (%)	8.8%	4.2%	-0.7%	26.7%
Effect MinW for Women (%)	15.1%	4.8%	1.6%	17.2%

Notes: Each observation represents a worker-month. Triple interactd version of specification (1). All regressions include store×department fixed effects, worker fixed effects, pair×month fixed effects, and an uninteracted dummy variable for being in "department A." Standard errors are two-way clustered at the state and border-segment level. *** p<0.01, ** p<0.05, * p<0.1. "Sales per hour" are the sales per hour rescaled by a factor between 1/50 and 1/150 relative to its $ value. "Total pay per hour" is the monthly total pay (in $ per hour). "Retained" is a dummy variable that equals one if the worker is retained that month (i.e., not terminated). "Discounted synthetic pay per hour" is the synthetic pay per hour -- i.e., the hourly pay the company would have paid the worker had they made the same sales as in the month before the minimum wage increase, calculated as the maximum of total pay per hour in *t-1* and the minimum wage in *t* -- multiplied by the discount factor $[(1+r)/(1+r-\pi)]$, where r is the monthly discount rate and π is the average monthly retention rate by gender (lagged). "MinW" is the predominant minimum wage (in $). "Effect MinW for Men (%)" [resp., "Effect MinW for Women (%)"] is the percent effect of a $1 increase in MinW relative to the mean of the outcome variable for men [resp., women] in department A or B.

Table A.10: Impact of the Minimum Wage on Pay and Productivity by Gender with Alternative Specifications (Ceteris-Paribus Analysis)

	(1)	(2)	(3)	(4)
	Poisson pseudo-likelihood regression		**Log-log regression**	
Dep.Var.	Sales per hour	Total pay per hour	Log sales per hour	Log total pay per hour
MinW	0.024	0.043***		
	(0.017)	(0.009)		
MinW × Woman	0.024***	0.005		
	(0.009)	(0.011)		
Log MinW			0.292	0.327***
			(0.304)	(0.118)
Log MinW × Woman			0.188**	0.057
			(0.078)	(0.096)
Observations	217,746	215,558	217,746	215,558
p-value MinW+MinW×Woman	0.027	0.013	0.078	0.004

Notes: Each observation represents a worker-month. All regressions include store×department fixed effects, worker fixed effects, pair×month fixed effects, and control for MinW×department. Standard errors are two-way clustered at the state and border-segment level. *** p<0.01, ** p<0.05, * p<0.1. The first two columns present a Poisson pseudo-likelihood regression, while the last two a log-log regression that logs the dependent and independent variables.

Table A.11: Test of Pre-trends by Gender

	(1)	(2)	(3)	(4)
	Productivity	**Pay**		
Dep.Var.	Sales per hour	Total pay per hour = col.(3)+(4)	Regular pay per hour (fixed+variable)	MinW top-up per hour
Panel A: 12-Months Pre-Trend				
Pre-trend (12 months)	-0.081	0.094	0.116	-0.022
	(0.147)	(0.303)	(0.080)	(0.357)
Pre-trend (12 months) × Woman	0.159	0.051	-0.120	0.172
	(0.120)	(0.216)	(0.114)	(0.317)
Observations	113,648	111,933	111,933	111,933
Mean Dep.Var.	2.085	12.271	12.046	0.225
p-value Pre-trend+Pre-trend×Woman	0.428	0.604	0.915	0.612
Panel B: 6-Months Pre-Trend				
Pre-trend (6 months)	-0.029	-0.011	0.049	-0.060
	(0.142)	(0.188)	(0.049)	(0.224)
Pre-trend (6 months) × Woman	0.023	-0.294	-0.051	-0.243
	(0.105)	(0.245)	(0.051)	(0.289)
Observations	150,924	149,010	149,010	149,010
Mean Dep.Var.	2.085	12.271	12.046	0.225
p-value Pre-trend+Pre-trend×Woman	0.955	0.236	0.958	0.271
Panel C: 3-Months Pre-Trend				
Pre-trend (3 months)	0.104	0.163	0.063	0.100
	(0.144)	(0.215)	(0.058)	(0.252)
Pre-trend (3 months) × Woman	-0.067	-0.444	-0.010	-0.435
	(0.139)	(0.232)	(0.054)	(0.271)
Observations	180,466	178,394	178,394	178,394
Mean Dep.Var.	2.085	12.271	12.046	0.225
p-value Pre-trend+Pre-trend×Woman	0.671	0.234	0.252	0.200

Notes: Each observation represents a worker-month. "Pre-trend (j months)" corresponds to the estimate of $\eta(1-0) - \eta(j-1)$ in the specification reported in the paper and "Pre-trend (j months)×Woman" to the estimate of $\theta(1-0) - \theta(j-1)$, where j is equal to 12 in panel A, 6 in panel B and 3 in panel C. All regressions include store×department fixed effects, worker fixed effects, pair×month fixed effects and control for MinW×department. Standard errors are two-way clustered at the state and border-segment level. *** p<0.01, ** p<0.05, * p<0.1. See Appendix C.3 for more details.

Table A.12: Impact of the Minimum Wage on Pay, Productivity, and Retention by Gender in Different Subsamples (Ceteris-Paribus Analysis)

	(1)	(2)	(3)	(4)
	Balanced panel	**Gender of team members**	**Distance of centroids**	
Sample restricted to …	Workers with 40 months of tenure or more	Teams composed entirely of male sales associates	Stores located in counties with centroids <37.5 km apart	Stores located in counties with centroids <18.75 km apart
Panel A: Dep.Var. = Sales per Hour				
MinW	-0.078	0.090	0.025	0.086
	(0.083)	(0.215)	(0.045)	(0.063)
MinW × Woman	0.069***		0.054**	0.081**
	(0.016)		(0.020)	(0.025)
Observations	79,714	13,595	152,523	100,212
Mean Dep.Var.	2.063	2.408	2.025	1.973
p-value MinW+MinW×Woman	0.911		0.155	0.090
Panel B: Dep.Var. = Total Pay per Hour				
MinW	0.404***	0.610	0.507**	0.781
	(0.115)	(0.411)	(0.199)	(0.420)
MinW × Woman	0.097		0.042	0.166
	(0.085)		(0.192)	(0.218)
Observations	78,412	13,550	150,750	98,931
Mean Dep.Var.	12.803	11.857	12.387	12.585
p-value MinW+MinW×Woman	0.009		0.146	0.166
Panel C: Dep.Var. = Retained				
MinW	-0.001	-0.035	0.005	-0.002
	(0.010)	(0.039)	(0.006)	(0.001)
MinW × Woman	0.009***		0.016***	0.015***
	(0.003)		(0.003)	(0.004)
Observations	79,714	13,595	152,523	100,212
Mean Dep.Var.	0.979	0.950	0.956	0.957
p-value MinW+MinW×Woman	0.521		0.011	0.002

Notes: Each observation represents a worker-month. Regressions are estimated in a subsample of observations, described in the panel headings. All regressions control for store×department fixed effects, worker fixed effects, pair×month fixed effects, and MinW×department. Standard errors are two-way clustered at the state and border-segment level. *** p<0.01, ** p<0.05, * p<0.1. Column 2 restricts the sample to teams composed entirely of male sales associates (hence, the sample drops).

Table A.13: Impact of the Minimum Wage by Gender and Outside Option – Robustness (Ceteris-Paribus Analysis)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dep.Var.	Sales per hour							
	Robustness to different measures of OOI				Robustness to different transition matrices			
Measure of OOI:	*OOI does not account for unemployment duration*		*OOI accounts for work-home distance*		*BGT transitions*		*CPS transitions (balanced sample)*	
MinW	0.231***	0.218***	0.139***	0.136***	0.212***	0.201***	0.221***	0.212***
	(0.040)	(0.044)	(0.038)	(0.038)	(0.042)	(0.045)	(0.040)	(0.044)
MinW × Woman	-0.016	0.018	0.010	0.020	-0.010	0.024	-0.007	0.019
	(0.014)	(0.069)	(0.013)	(0.018)	(0.013)	(0.071)	(0.013)	(0.065)
MinW × OOI	-0.006***	-0.005***	-0.003**	-0.003**	-0.007***	-0.007***	-0.007***	-0.007***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)
MinW × Woman × OOI		-0.002		-0.001		-0.003		-0.002
		(0.003)		(0.001)		(0.005)		(0.004)
Observations	212,443	212,443	210,258	210,258	212,462	212,462	212,462	212,462
Mean Dep.Var.	2.087	2.087	2.086	2.086	2.087	2.087	2.087	2.087

Notes : Each observation represents a worker-month. All regressions include store×department fixed effects, worker fixed effects, pair×month fixed effects and control for MinW×department. They also control for OOI and OOI×Woman. Standard errors are two-way clustered at the state and border-segment level. *** p<0.01, ** p<0.05, * p<0.1. OOI is the outside option index, which varies based on the worker's zip code and gender, and is lagged by one year. In columns 1-2, the OOI does not include the unemployment duration scaling factor. In columns 3-4, the OOI accounts for worker home-work distance. In columns 5-6, the OOI uses BGT transitions (not gender specific). In columns 7-8, the OOI uses CPS transitions (balanced sample). See Appendix D.1 (paragraph "extensions") for more details.

Table A.14: Impact of the Minimum Wage on the Outside Option
(Zip-code Level Analysis)

	(1)	(2)	(3)	(4)
Dep.Var.	MinW (in t+1)	MinW (in t+1)	OOI of Women	OOI of Men
OOI of Women	0.001			
	(0.002)			
OOI of Men		-0.002		
		(0.001)		
MinW			-0.052	-0.171
			(0.102)	(0.136)
Observations	4,234	6,630	4,234	6,630
Mean Dep.Var.	7.965	7.965	12.61	17.29
Effect MinW (%)	<0.01%	-0.02%	-0.41%	-0.99%

Notes: Each observation represents a zip code-year. All regressions include year fixed effects. The sample in columns 1 and 3 [resp., 2 and 4] is restricted to zip code-years with at least one female (resp., male) worker from our sample. See Appendix D.1 for more details on the OOI. *** p<0.01, ** p<0.05, * p<0.1.

Table A.15: Impact of the Minimum Wage on Welfare by Gender – Robustness (Ceteris-Paribus Analysis)

Dep.Var.: Discounted synthetic pay per hour

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Estimating the wrong welfare results	**Sensitivity of welfare results to aggregation of π**			**Sensitivity of welfare results to r**		**Sensitivity of welfare results to inluding outside option**
Assumptions	Use *contemporaneous* pay per hour rather than *synthetic* pay per hour	π = average retention rate *across all months*, *by gender*	π = average retention rate, measured *each month*, *by gender×department*	π = average retention rate, measured *each month*, *by gender×department×store*	r = 0.41%	r = 1.5%	outside option u^U by gender
	$\frac{1+r}{1+r-\pi(\boldsymbol{e}_t^*)} \cdot w(\boldsymbol{e}_{it}^*; M_{jt})$	$\frac{1+r}{1+r-\boldsymbol{\pi}(\boldsymbol{e}^*)} \cdot w(e_{i,t-1}^*; M_{jt})$	$\frac{1+r}{1+r-\boldsymbol{\pi_d}(\boldsymbol{e}_{t-1}^*)} \cdot w(e_{i,t-1}^*; M_{jt})$	$\frac{1+r}{1+r-\boldsymbol{\pi_{dj}}(\boldsymbol{e}_{t-1}^*)} \cdot w(e_{i,t-1}^*; M_{jt})$	$\frac{1+\boldsymbol{r}}{1+\boldsymbol{r}-\pi(e^*)} \cdot w(e_{i,t-1}^*; M_{jt})$		$\frac{1+r}{1+r-\pi(e^*)} \cdot \left(w(e_{i,t-1}^*; M_{jt}) - \frac{\boldsymbol{1-\pi(e^*)}}{\boldsymbol{r}} \boldsymbol{u^u}(\boldsymbol{M_{jt}})\right)$
MinW	9.714***	20.790***	23.541***	100.716***	29.139***	24.049***	10.582***
	(2.841)	(3.991)	(4.224)	(19.533)	(5.526)	(4.595)	(3.245)
MinW × Woman	1.224	-12.418***	-20.515***	-58.012***	-17.670***	-14.451***	-12.836***
	(5.753)	(1.682)	(5.191)	(16.585)	(2.487)	(1.985)	(1.153)
Observations	215,558	197,333	197,333	197,333	197,333	197,333	192,581
Mean Dep.Var.	191.856	179.451	197.566	387.359	250.200	207.161	179.223
p-value MinW+MinW×Woman	0.066	0.052	0.602	0.080	0.062	0.056	0.560

Notes: All regressions include store×department fixed effects, worker fixed effects, pair×month fixed effects and control for MinW×department. Standard errors are two-way clustered at the state and border-segment level. *** p<0.01, ** p<0.05, * p<0.1. In column 1, "discounted pay per hour (contemporaneous; not synthetic)" computes the welfare formula with the contemporaneous pay [w(e_i,t)] rather than the synthetic pay [w(e_i,t-1)]. In columns 2-7, the dependent variable is the discounted synthetic pay per hour using the synthetic pay [w(e_i,t-1)]. How we measure the discount factor, [(1+r)/(1+r-π)], is explained in the column headings. In column 7, we allow the outside option to vary with the minimum wage. The outside option is measured with OOI. See Appendix F.2 for more details.

Table A.16: Impact of the Minimum Wage on Pay, Productivity, Retention and Welfare by Gender (Non-Ceteris-Paribus Analysis)

	(1)	(2)	(3)	(4)	(5)	(6)
	Productivity	**Pay**			**Retention**	**Welfare**
Dep.Var.	Sales per hour	Total pay per hour = col. (3)+(4)	Regular pay per hour (fixed+variable)	MinW top-up per hour	Retained	Discounted synthetic pay per hour
MinW	-0.061	0.197	-0.081	0.278***	0.000	7.656***
	(0.044)	(0.149)	(0.152)	(0.050)	(0.006)	(1.382)
MinW × Woman	0.129**	0.357***	0.352**	0.005	0.005**	7.473***
	(0.050)	(0.095)	(0.132)	(0.061)	(0.002)	(2.202)
Observations	217,746	215,565	215,565	215,565	217,746	198,033
Mean Dep.Var.	2.085	12.27	12.05	0.225	0.954	195.9
p-value MinW+MinW×Woman	0.209	<0.001	0.048	<0.001	0.399	<0.001

Notes: All regressions include store fixed effects, pair×month fixed effects and control for the uninteracted woman dummy. There are no store×department fixed effects and no worker fixed effects in the regression. Standard errors are two-way clustered at the state and border-segment level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$. "Sales per hour" are the sales per hour rescaled by a factor between 1/50 and 1/150 relative to its \$ value. "Regular pay per hour" is the total amount earned from the base hourly rate and variable pay (commission rate × sales per hour), without the top-up. "MinW top-up per hour" is the monthly total minimum wage adjustment paid by the company to the worker (in \$ per hour). "Retained" is a dummy variable that equals one if the worker is retained that month (i.e., not terminated). "Discounted synthetic pay per hour" is the synthetic pay per hour—i.e., the hourly pay the company would have paid the worker had they made the same sales as in the month before the minimum wage increase, calculated as the maximum of total pay per hour in *t-1* and the minimum wage in *t*—multiplied by the discount factor $[(1+r)/(1+r-\pi)]$, where r is the monthly discount rate and π is the average monthly retention rate by gender (lagged).

B Similarity Between Our Sample Workers and US Hourly-Paid Workers

Given the absence of comprehensive economy-wide data that compare female and male workers at the same level of granularity as we do,[41] we compute summary statistics at our firm's level (i.e., without controlling for department or store) and juxtapose them to statistics for US workers who were "paid by the hour" in 2015 (henceforth, hourly workers for short).

Our worker pool resembles US hourly workers in several dimensions: worker pay, retention, and other factors.

- Among US hourly workers, 4.1% of women and 2.5% of men were paid "at or below" the minimum wage (US Bureau of Labor Statistics, 2015a). These percentages are similar to the fraction of our workers who receive minimum wage top-ups for four weeks in a month in our setting (4.9% of women, 2% of men).
- The bottom decile of hourly earnings is also very similar to our setting: \$9 per hour for male hourly workers and \$8.3 for female hourly workers (compared to \$8.9 and \$8.4 in our setting).[42]
- For US hourly workers, median hourly earnings was \$12.6 for women and \$14.6 for men; for our workers, these figures are somewhat lower but comparable at, respectively, \$10.8 and \$11.3. Hourly earnings might be lower in our setting because workers are somewhat younger: the average and median ages are 36 and 27 years old in our sample, compared to a median and average age of 40 years old among all hourly workers.
- The gender disparity in monthly termination rates among US hourly workers is

[41] Indeed, outside of lab experiments, we are unaware of any economy-wide study in the US that controls for job characteristics finer than establishment×occupation, and Figure A.1 shows that such fine distinctions make a big difference for our purposes.

[42] The statistics on the bottom decile of pay and age are taken from the CPS data, January 2015.

comparable to our setting: 2.9% for women and 3.6% for men,[43] as opposed to 4.1% and 4.8% in our setting.

- One key difference is workforce composition. In our study context, men make up 70% of the workforce, compared to 44% in the broader retail sector (US Census Bureau, 2020). This higher share of men in our sample reflects our focus on retail sales roles rather than lower-wage positions like cashier jobs, which employ more women.

[43]These rates are calculated as the number of female (male) workers who lost their job and have been searching for a new one in the last month, as a share of the number of employed female (male) workers (US Bureau of Labor Statistics, 2015b).

C Minimum Wage Variation and Identification

C.1 Minimum wage variation

Our data contain information on the geographical location of stores (latitude and longitude), which we match with the monthly statutory minimum wage level in that store, extracted from the public dataset maintained by the Washington Center for Equitable Growth. Variations in minimum wage take place at state, county, and city levels; with city and county minimum wages always set to be higher than the state minimum wage.

From February 2012 to June 2015, our sample of stores is affected by 70 variations in minimum wage: 49 variations are at the state level, and 21 are at the county or city level. The exact timing of each minimum wage change is reported in Table C.1 and presented visually in Figure C.1.

Figure C.1: Variations in Minimum Wage from February 2012 to June 2015

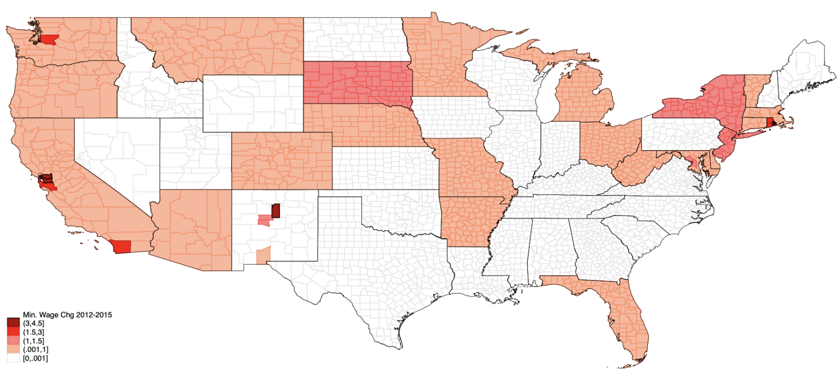

Notes: Store locations are withheld for confidentiality reasons.

Of all the variations in minimum wage present in our sample, the three changes in Florida coincide with state-level variations in unemployment insurance potential benefits duration (see Lusher et al. (2022) Online Appendix 2, page 6). State-level changes in unemployment insurance potential benefits duration in Arkansas, Illinois, Michigan, Georgia, North and South Carolina, and Missouri occurred either before

Table C.1: Changes in Minimum Wages from February 2012 and June 2015

State	State	Date C.1	W_{t-1}	W_t	Date C.2	W_{t-1}	W_t	Date C.3	W_{t-1}	W_t	Date C.4	W_{t-1}	W_t
Alaska	AK	2015m2	7.75	8.75									
Arkansas	AR	2015m1	7.25	7.5									
Arizona	AZ	2013m1	7.65	7.8	2014m1	7.8	7.9	2015m1	7.9	8.05			
California	CA	2014m7	8	9									
Colorado	CO	2013m1	7.64	7.78	2014m1	7.78	8	2015m1	8	8.23			
Connecticut	CT	2014m1	8.25	8.7	2015m1	8.7	9.15						
DC	DC	2014m7	8.25	9.5									
Delaware	DE	2014m6	7.25	7.75	2015m6	7.75	8.25						
Florida	FL	2013m1	7.67	7.79	2014m1	7.79	7.93	2015m1	7.93	8.05			
Hawaii	HI	2015m1	7.25	7.75									
Massachusetts	MA	2015m1	8	9									
Maryland	MD	2015m1	7.25	8									
Michigan	MI	2014m9	7.4	8.15									
Minnesota	MN	2014m8	7.25	8									
Missouri	MO	2013m1	7.25	7.35	2014m1	7.35	7.5	2015m1	7.5	7.65			
Montana	MT	2013m1	7.65	7.8	2014m1	7.8	7.9	2015m1	7.9	8.05			
Nebraska	NE	2015m1	7.25	8									
New Jersey	NJ	2014m1	7.25	8.25	2015m1	8.25	8.38						
New York	NY	2013m12	7.25	8	2014m12	8	8.75						
Ohio	OH	2013m1	7.7	7.85	2014m1	7.85	7.95	2015m1	7.95	8.1			
Oregon	OR	2013m1	8.8	8.95	2014m1	8.95	9.1	2015m1	9.1	9.25			
Rhode Island	RI	2013m1	7.4	7.75	2014m1	7.75	8	2015m1	8	9			
South Dakota	SD	2015m1	7.25	8.5									
Vermont	VT	2014m1	8.6	8.73	2015m1	8.73	9.15						
Washington	WA	2013m1	9.04	9.19	2014m1	9.19	9.32	2015m1	9.32	9.47			
West Virginia	WV	2015m1	7.25	8									

County	State	Date C.1	W_{t-1}	W_t	Date C.2	W_{t-1}	W_t	Date C.3	W_{t-1}	W_t
Bernalillo	NM	2013m7	7.5	8	2014m1	8	8.5	2015m1	8.5	8.65
Johnson	IA	2015m11	7.25	8.2						
Montgomery	MD	2014m10	7.25	8.4						
Prince George's	MD	2014m10	7.25	8.4						
Santa Fe	NM	2014m4	7.5	10.66	2015m3	10.66	10.84			

City	State	Date C.1	W_{t-1}	W_t	Date C.2	W_{t-1}	W_t	Date C.3	W_{t-1}	W_t	Date C.4	W_{t-1}	W_t
Alburquerque	NM	2013m1	7.5	8.5	2014m1	8.5	8.6	2015m1	8.6	8.75			
Berkeley	CA	2014m10	9	10									
Las Cruces	NM	2015m1	7.5	8.4									
Oakland	CA	2015m3	9	12.25	2016m1	12.25	12.55						
Richmond	CA	2015m1	9	9.6	2016m1	9.6	11.52						
San Diego	CA	2015m1	9	9.75									
San Francisco	CA	2013m1	10.24	10.55	2014m1	10.55	10.74	2015m1	10.74	11.05	2015m5	11.05	12.25
San Jose	CA	2013m3	8	10	2014m1	10	10.15	2015m1	10.15	10.3			
Santa Fe	NM	2012m3	9.5	10.29	2013m3	10.29	10.51	2014m3	10.51	10.66	2015m3	10.66	10.84
SeaTac	WA	2013m1	9.04	9.19	2014m1	9.19	15						
Seattle	WA	2013m1	9.04	9.19	2014m1	9.19	9.32	2015m1	9.32	9.47	2015m4	9.47	11
Sunnyvale	CA	2015m1	9	10.3									
Tacoma	WA	2013m1	9.04	9.19	2014m1	9.19	9.32	2015m1	9.32	9.47			
Washington	DC	2014m7	8.25	9.5									

Notes: This table reports all state/county/city variations in statutory minimum wage from 2/1/2012 to 6/30/2015, irrespective of whether there is a store located in that state/county/city. The data are extracted from the public dataset maintained by the Washington Center for Equitable Growth. Our identification strategy effectively leverages only a sub-sample of these changes (70 out of 89), i.e., those that affect at least one store in our sample. We do not report which ones are the 70 variations we leveraged in the paper for confidentiality reasons. W_t (W_{t-1}) refers to the minimum wage level after (before) the change. The states with no change in minimum wage from February 2012 and June 2015 are: AL, GA, IA, ID, IL, IN, KS, KY, LA, ME, MS, NC, ND, NH, NM, NV, OK, PA, SC, TN, TX, UT, VA, WI, WY.

our sample period or sufficiently distant in time from the minimum wage changes employed in our research design.

Table A.14 (columns 1-2) shows that changes in the minimum wage are unrelated to the outside option index (OOI) fo both women and men. See Section D.1 for a precise definition.

C.2 Border discontinuity design

We use a border discontinuity design, as implemented in Card & Krueger (2000), Dube et al. (2010, 2016) and Allegretto et al. (2011, 2017). This approach exploits minimum wage policy discontinuities at the state- or county-border by comparing workers on one side of the border where the minimum wage increased (treatment group) to workers on the other side where the minimum wage did not increase (control group). As shown in Dube et al. (2010), this research design has desirable properties for identifying minimum wage effects since workers on either side of the border are more likely to face similar economic conditions and are likely to experience similar shocks at the same time. The main disadvantage of this design is the risk of cross-border worker movements from control to treated stores (Neumark et al., 2014). We alleviate this concern in Section 3.2.

Following Card & Krueger (2000), Dube et al. (2010, 2016) and Allegretto et al. (2017), we restrict our sample to stores (and their respective workers) located in adjacent counties that share a border. For *state*-level minimum wage variations, we keep stores located in county pairs that: share a *state* border, and whose centroids are within 75 km of each other: see Figure C.2.

For *county*-level minimum wage variations, we "seed" the sample with stores located in those counties that increased their minimum wage, and then add as controls all adjacent counties whose centroids are within 75 km of the seed county. Minimum wage changes at the *city* level are attributed only to stores within the city limits, but not to stores in the county containing that city. Such stores are included as controls, as are stores in all neighboring counties. (In our sample there are no municipalities

Figure C.2: Variations in the Minimum Wage in Bordering Counties

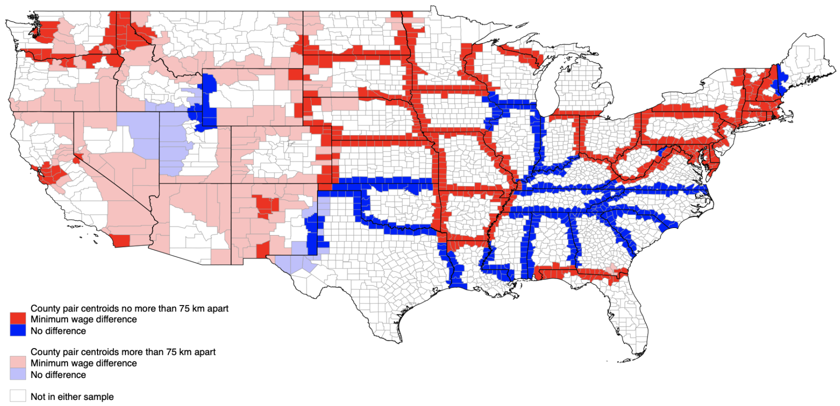

Notes: Store locations are withheld for confidentiality reasons.

that lie in more than one county). For instance, for the city of San Francisco (which increased its minimum wage) we include all the counties that share a county-border with San Francisco County and whose centroids are within 75 km of its centroid (i.e., the counties of Marin, Alameda, and San Mateo).

C.3 Pre-trends

We test for differential pre-trends by gender in the twelve months preceding the minimum wage change using an autoregressive distributed lag model. This model, which has been commonly used in the minimum wage literature (Dube et al., 2010), has the advantage of taking into account the sequential occurrence of changes in the minimum wage level. We estimate:

$$\begin{aligned} Y_{idjpt} \quad = \quad & \alpha + \eta_{12-1}(M_{j,t+12} - M_{j,t+1}) + \eta_{1-0}(M_{j,t+1} - M_{j,t}) + \theta_{12-1}(M_{j,t+12} - M_{j,t+1}) \times Woman_i \\ & + \theta_{1-0}(M_{j,t+1} - M_{j,t}) \times Woman_i + \rho M_{j,t} + X_{idjt}\eta + \delta_i + \zeta_{dj} + \phi_{pt} + \varepsilon_{ijpt}. \end{aligned} \tag{14}$$

Here, $M_{j,t+m}$ is the minimum wage m months after month t, and all other variables are defined as in equation (1). η_{12-1} (η_{1-0}) is a leading coefficient that captures variations in sales per hour during the months -12 to -1 (resp., -1 to 0) from each change in the minimum wage for men. θ_{12-1} (θ_{1-0}) are the corresponding differences across gender. We assess whether men and women are on different trends before the minimum wage increase by estimating whether $\theta_{1-0} - \theta_{12-1}$ is statistically different than zero.

The results are presented in Table A.11. The specification is estimated for the sample of 110 thousand workers-months who are continuously employed for 12 months before the minimum wage event (panel A). We find no gender differential pre-trends preceding changes in the minimum wage. There are also no differential pre-trends in the 6 and 3 months preceding the minimum wage change, using the larger sample of 150 and 180 thousand workers-months who are continuously employed for 6 and 3 months before the minimum wage event (panels B-C). The estimate for $\theta_{1-0} - \theta_{12-1}$ decreases from panel A to C, but it is never statistically significant and a joint test never rejects the lack of pre-trends.

D Measuring the Outside Option

D.1 The outside option index (OOI)

Construction of the OOI Our primary measure of the outside option is the OOI, defined as:

$$OOI_{gzy} = \nu_{gy} \sum_o \theta_{og} \cdot \frac{s_{ogzy}}{s_{ogy}} \cdot w_{ogzy},$$

where g represents the gender, z represents the zip code, y represents the year, and o represents the occupation. The OOI is constructed using data from the American Community Survey (ACS) and the Current Population Survey (CPS) for 2012-2015.

Hourly wages (w_{ogzy}): Calculated using ACS yearly data by dividing median annual earnings (varies by occupation, gender, zip code, year) by total hours worked, based on midpoint estimates of bins for weeks worked and hours per week.[44] Both shares vary by gender, zip code, year.[45]

Occupation availability ($\frac{s_{ogzy}}{s_{ogy}}$): Calculated using ACS yearly data by dividing the share of workers in an occupation (varies by occupation, gender, zip code, year) by the gender-specific national average share for that occupation (varies by occupation, gender, and year).

Nationwide occupation transition shares (θ_{og}): Derived from CPS data by calculating the probability of transitioning from sales occupations (SOC 2-digit = 41-000) to other occupations (with or without an employer change) or changing employers within the same occupation; aggregated for 2012-2015 (varies by gender and occupation). The origin occupation is restricted to sales, while all destination occupations are considered.[46] Table D.2 presents the transition shares.

[44] For any given gzy, we know the number of individuals in each hours-weeks bin (e.g., those working 35-45 hours per week and 50-52 weeks per year). We compute the midpoint of each bin (e.g., 40 and 51 in the example above) and multiply them (40*51) to estimate total hours for each category. Finally, we take a weighted average across categories (now expressed in hours), using each category's population share as weights.

[45] ACS contains info on ZCTA, which typically correspond to zip codes. We map ZCTA to zip codes using the Missouri Census Data Center's (MCDC) Geocorr 2014 crosswalk.

[46] Like Schubert et al. (2024), the nationwide occupation transition shares vary across space but

To ensure that the OOI is interpreted in dollars per hour, we rescale the weight $\theta_{og} \cdot \frac{s_{ogzy}}{s_{ogy}}$ for each ogy so that the sum across all occupations ($\sum_o$) is equal to 1.

Table D.2: Transition Shares by Gender

	Origin: Sales and Related Occupations (41-0000)	
Destination	Women	Men
Management Occupations (11-0000)	8.3%	14.4%
Business and Financial Operations Occupations (13-0000)	4.5%	4.8%
Computer and Mathematical Occupations (15-0000)	0.7%	2.1%
Architecture and Engineering Occupations (17-0000)	0.2%	1.1%
Life Physical and Social Science Occupations (19-0000)	0.2%	0.4%
Community and Social Service Occupations (21-0000)	0.8%	0.4%
Legal Occupations (23-0000)	0.3%	0.3%
Educational Instruction and Library Occupations (25-0000)	3.1%	1.1%
Arts Design Entertainment Sports and Media Occupations (27-0000)	1.8%	1.6%
Healthcare Practitioners and Technical Occupations (29-0000)	2.6%	0.8%
Healthcare Support Occupations (31-0000)	2.5%	0.3%
Protective Service Occupations (33-0000)	0.6%	1.3%
Food Preparation and Serving Related Occupations (35-0000)	9.4%	4.6%
Building and Grounds Cleaning and Maintenance Occupations (37-0000)	2.5%	2.6%
Personal Care and Service Occupations (39-0000)	4.7%	1.3%
Sales and Related Occupations (41-0000)	30.2%	30.1%
Office and Administrative Support Occupations (43-0000)	21.2%	10.7%
Farming Fishing and Forestry Occupations (45-0000)	0.4%	0.5%
Construction and Extraction Occupations (47-0000)	0.3%	4.8%
Installation Maintenance and Repair Occupations (49-0000)	0.3%	5.0%
Production Occupations (51-0000)	3.1%	4.5%
Transportation and Material Moving Occupations (53-0000)	2.5%	7.2%

Notes : This table presents transitions shares by gender derived from CPS data.

Unemployment duration scaling factor (ν_{gy}): Calculated as (employment - unemployment duration) / employment duration, where employment duration is estimated based on our workers' tenure at the time they exit their position, and unemployment duration is sourced from BLS Labor Force Statistics (varies by gender and year).[47]

remain constant over time. Unlike Schubert et al. (2024), we use CPS data as our primary data instead of Burning Glass Technology data, as the latter is not gender-specific. We will show that the results are robust to using Burning Glass Technology data.

[47]Instead of scaling down by unemployment duration, one could alternatively assume that, when assessing their outside option in NPV, employed workers give greater weight to near-term unemployment than to future employment with positive wages. However, this would require taking a stand on the discount factor, which we avoid for transparency. This implies that our measure of OOI likely overestimates the true NPV of outside options while employed.

Extensions of the OOI Table A.13 shows that the results in Table 2 (columns 1-2) are robust to the following adjustments of the OOI measure:

- *Not accounting for unemployment.* In Table A.13 (columns 1-2), we exclude the unemployment duration scaling factor (ν_{gy}) in the OOI measure:

$$OOI_{gzy} = \sum_o \theta_{og} \cdot \frac{s_{ogzy}}{s_{ogy}} \cdot w_{ogzy}.$$

- *Incorporating commuting preferences.* In Table A.13 (columns 3-4), we extend the primary OOI measure to include η_{izy}, the ratio of a worker's commuting distance to the average commuting distance for their zip code:

$$OOI_{igzy} = \eta_{izy} \cdot \nu_{gy} \cdot \sum_o \theta_{og} \cdot \frac{s_{ogzy}}{s_{ogy}} \cdot w_{ogzy}.$$

  η_{izy} is calculated as the ratio of worker i's home-work distance (d_{izy}) to the average home-work distance of workers residing in the same zip code (d_{zy}). Worker-specific distances (d_{izy}) are derived from our primary firm dataset, while zip code averages (d_{zy}) are calculated using LEHD Origin-Destination Employment Statistics (LODES).[48]

  This measure of the OOI varies at the worker (rather than zip code) level. It captures both variation across zip codes and within zip codes across workers, under the assumption that workers with longer commutes have a higher tolerance for commuting and, thus, a higher OOI.

- *Using different transition matrices.* Our preferred measure of θ_{og} is based on CPS data, which allows us to focus on sales workers as the origin occupation and, importantly, provides gender disaggregation. In Table A.13 (columns 5-6), we show that the results are robust to using transition shares from Schubert et al. (2024), based on Burning Glass Technologies (BGT) resume data. (Since

[48]LODES provides data on the number of workers commuting from a given home zip code to all other zip codes (including the home zip itself), which we weight by the distances between the home and work zip codes (final variable varies by zip code and year).

BGT transition matrices are not gender-specific, we construct a non-gender-specific measure of θ_o.) In Table A.13 (columns 7-8), we show that the results are robust to restricting transitions to occupation changes over two consecutive months in CPS data, and aligning cross-occupation transitions with within-occupation transitions for consistency.

Empirical specification for Figure 2 Figure 2 presents results using a more non-parametric specification in which we allow the effect of the minimum wage by gender to vary across different "bins" of OOI by estimating:

$$\begin{aligned} Y_{idjpt} &= \alpha + \sum_{k=1}^{5} \theta_k \mathbb{1}(OOI)^k_{gzy-1} + \sum_{k=1}^{5} \mu_k \mathbb{1}(OOI)^k_{gzy-1} \times Woman_i + \beta M_{jt} + \gamma M_{jt} \times Woman_i \\ &+ \sum_{k=1}^{5} \lambda_k M_{jt} \times \mathbb{1}(OOI)^k_{gzy-1} + \sum_{k=1}^{5} \psi_k M_{jt} \times Woman_i \times \mathbb{1}(OOI)^k_{gzy-1} \\ &+ X_{idjpt}\eta + \delta_i + \zeta_{dj} + \phi_{pt} + \varepsilon_{idjpt}. \end{aligned}$$

$\mathbb{1}(OOI)^k_{gzy-1}$ represents different OOI bins. Panel A uses the full sample of workers. Panel B and C restrict the sample to worker-months in which the worker is supervised or unsupervised, respectively.

D.2 A different proxy for the outside option based on financial transactions data

To construct an alternative measure of our workers' outside option that is different from the OOI, we use financial transaction data sourced from a large financial aggregation and analytics firm. The data includes anonymized bank, credit, and debit card transactions from over 60 million Americans. These data have been employed in previous studies (e.g., Aiello et al. 2024; Di Maggio et al. 2023) and shown to be representative of the broader economy (see Section 2.1.1 of Aiello et al. 2024).

For each (anonymous) account (effectively a member ID), we observe individual transactions with details such as the date, amount, and, when the transaction is in person, the "city" in which the transaction takes place ("city" is a geographic unit used by the analytics firm that is coarser than zip code) and the merchant's name. Transactions are categorized into 43 types (e.g., salary, ATM withdrawal, groceries, mortgage payments, medical spending) based on textual descriptions. While the dataset lacks direct demographic details such as gender or residence zip code, we impute an account owner's gender based on spending in gender-specific merchant categories (e.g., cosmetic stores, women's ready-to-wear stores). The account owner's home city is imputed by the analytics firm using an algorithm that infers the city and state of residence based on the physical locations of frequently visited merchants, which is available from 2014 onward. Before 2014, we impute the account's home city ourselves based on the city where the majority of "location-sensitive" spendings are made (e.g., mortgage, rent, utilities).

To derive a financial-transaction based measure of the outside option, we focus on those accounts within our dataset which, in a given month, received a salary *from our firm* during 2012-2015. We conceptualize each of these accounts as one of our firm's employees. In total, this includes 53k worker/accounts. For each worker/account, we define the termination date as the last month in which a payment from our firm is recorded. In total, this includes 33k terminations from our firm. After each termination, we compute "next-job earnings" by skipping to the first full month in which the worker/account receives a salary from a new employer. These next-job earnings are then averaged across workers at the city and year-of-termination level to obtain a financial-transaction based proxy of outside options that varies by year×gender×city. To approximate "next-job *hourly* earnings", we divide this monthly next-job earnings by the average number of hours worked by our workers in that year×gender×city.

This proxy of our worker's outside option has some advantages over the OOI. First, it measures specifically the post-termination wages earned by (some of) our workers after the termination event, as opposed to average wages earned by salespeople in a given zip code. Second, there is no need to rely on a job-to-job transition matrix for "salespeople in general." However, there are serious disadvantages, too. The main disadvantage is that coverage is majorly spotty: because our firm's stores are present in many more cities than are present in our financial-transaction dataset, the financial-transaction based proxy can only be constructed for approximately one-third of the zip codes covered by the OOI. Second, variation is only at the city level, not at the zip code level. Third, unlike in ACS, we do not see *hourly* wages, and thus need to inpute hours. Nevertheless, if the two proxies can be shown to be highly correlated, this represents a cross-validation of the OOI.

To be able to correlate our two proxies, we need to define them at the same level of geographic granularity. Accordingly, we project the financial-transaction based proxy "next-job hourly earnings," which is defined at the city level, onto zip codes.[49] We then regress this new variable on the OOI (which varies at the zip code level; sample size: 4,564). The estimated coefficient is 0.605 (s.e. 0.208; statistically significant at the 1% level), which indicats that a $1/hour increase in the OOI is associated with a $0.605 increase in next-job hourly earnings. This high correlation level is, we believe, a reassuring cross-validation of the OOI.

[49]So, when a city has several zip codes, all of them get the same value.

E Ancillary Analysis In Support of the Outside Option Mechanisms

E.1 No impact of the minimum wage on firm organizational adjustments by gender

We examine the effect of the minimum wage on various potential organizational adjustments that the firm may have implemented, and test whether these vary by gender. Our focus is on adjustments that could explain why women's productivity increased more than men's in response to a minimum wage increase.

Hours: The minimum wage has no effect on the hours worked by women or men (Table A.6, column 1). This rules out the possibility that the firm disproportionately reduced hours for women, which could have led to higher productivity (sales per hour) due to reduced "fatigue per hour."

Shifts: The minimum wage does not disproportionately increase or decrease the share of women transitioning to full-time work relative to men (Table A.6, column 2). Since better shifts (i.e., busier shopping hours) are systematically allocated to full-time workers in our firms, this implies that the minimum wage did not move women to better shifts within a department.

Switches across departments or stores: The minimum wage does not differentially affect the likelihood of women vs. men switching departments or stores (Table A.6, columns 3-4). Women are therefore not more likely to be upgraded to "better" stores or departments following the introduction of the minimum wage.

Monitoring and termination rule: The minimum wage does not influence the likelihood of a female or male worker being assigned a supervisor (Table A.6, column 5). The minimum wage also does not affect the termination rule – the function mapping high productivity to lower termination risk – differentially by gender (Table A.7). Thus, the increase in women's retention following a minimum wage hike is consistent with them working harder relative to men, rather than resulting from a change in the termination rule.

Compensation scheme: Table A.5 (columns 3-4) shows that the minimum wage does not affect the compensation scheme for either gender. Thus, within-firm incentives do not change differentially by gender with the minimum wage increase.

E.2 Ruling out mechanisms other than the outside option

This section explores other prominent mechanisms beyond organizational adjustments (discussed in the previous section) and beyond the outside option story, that could explain why women become disproportionately more productive after a minimum wage increase (Table 1, column 1).

Gender differences in innate aptitude for the job, risk aversion, propensity to reciprocate: The gender differential productivity response could be due to gender differences in some innate characteristics such as innate aptitude for the job, risk aversion, propensity to reciprocate (as in a gift-exchange model), or, potentially, cognitive ability to deal with the complexity of employment contracts. For example, if women have different innate aptitude than men, they might respond differently to the minimum wage as in Table 1, column 1. But the gender differential in the response to the minimum wage disappears entirely when we control for the outside option (see Table 2, columns 1-2), and even vanishes *within monitoring regime* (supervised vs. unsupervised, columns 3-6), contradicting the hypothesized difference in aptitude.

Gender difference in job-fit dimensions: The gender differential productivity response could, in principle, partly reflect job-fit dimensions that are potentially correlated with gender, such as women valuing their current job more during child-rearing years or when facing shorter commutes. But the gender differential in Table 1, column 1, is unaffected by controlling for child-rearing age and home-work distance (Table A.8, columns 4 and 5). Moreover, as noted earlier, the gender difference in the productivity response disappears entirely when we control for the OOI (and this remains true even in a version of the OOI that incorporates commuting preferences, see Table A.13, columns 3-4).

Gender-specific demand or price shifts: The fact that sales per hour increased more for women than men could, in principle, partly reflect a disproportionate increase in the demand for or price of goods sold primarily by female representatives. Again, this seems unlikely because, as noted earlier, the gender difference in productivity response disappears entirely when we control for the OOI. Nevertheless, in what follows we drill down into the evidence on "gendered" demand or price changes – and rule them out directly.

A "gendered" demand shift could explain the estimates in Table 1, column 1, if the minimum wage alters the composition of customer demand *within a given department*[50] in a way that causes sales by female workers to increase disproportionately more than the men, quite apart and separately from the workers' outside option. To explore this possibility, we use the anonymized financial transactions data described in Appendix D.2. Every dollar spent for which our retailer appears as the merchandiser is linked to a specific store location in our database. This allows us to track the total number of transactions and spending amounts by "account gender" for each store and month.

Our identifying assumption is that if a "gendered" demand shift occurs after a minimum wage increase – disproportionately increasing sales for women salepeople – it would be driven by a rise in spending from female-gendered accounts. If so, we would expect spending from female-gendered accounts to increase disproportionately following a minimum wage increase. To test this, we estimate the following specification:

$$Y_{jpt} = \alpha + \beta M_{pt} + \eta X_{jpt} + \phi_{pt} + \varepsilon_{jpt},$$

where Y_{jpt} is the share of spending by female-gendered accounts in store j of county-pair p in month t (sample mean of 48%), and the other variables are defined as in specification (1). The coefficient of interest β captures the effect of the minimum wage on the fraction of spending by female-gendered accounts.

[50] Recall that our estimates include department×store fixed effects (ζ_{dj}) and control for possible differential effects of the minimum wage across departments: $M_{jt} * Department_d$.

The estimate for β is -0.013 (s.e. 0.012). That the point estimate is negative implies that after a minimum wage increase spending is *less*, not more likely to come from a female-gendered account. With this being said, the estimate is statistically indistinguishable from zero. Therefore, there is no evidence supporting the hypothesis that spending by female-gendered accounts increase disproportionately following a minimum wage increase.

Finally, the minimum wage could have also increased women's sales per hour more than men's if the firm raised prices on feminine SKUs more than masculine ones within the same department. However, this is unlikely, as the company applies nationwide pricing, which is not gender-specific and does not adjust in response to state or local minimum wage changes.

E.3 Increasing the minimum wage widens the gap between inside and outside option

It is reasonable to expect, speaking generally and without specific reference to our setting, that, for workers who *currently* benefit from the minimum wage, a higher minimum wage should widen the gap between inside and outside option. We expect the outside option (in net present value, NPV) to be less sensitive to the minimum wage than the inside option for the following reasons.

(1) Unemployment durations are long – averaging 39 weeks (median 20 weeks) nationwide (BLS Labor Force Statistics). Neither the flow value of unemployment nor its duration are expected to improve with the minimum wage Dube et al. (2016); Gittings & Schmutte (2016); if anything, they may worsen. Consequently, a significant portion of the NPV of becoming unemployed does not improve (and may even deteriorate) with the minimum wage.

(2) It is unlikely that our terminated workers will find a new job as strongly supported by the minimum wage as their current one. We peg this probability at 3% to 25%. Only 3% of hourly workers earned the minimum wage in 2015 according

to the BLS. Thus, if an unemployed worker randomly drew her next job from the nationwide job supply, the probability of landing a minimum wage job would be very low. Moreover, between 2012-2015, about 25% of workers with similar pay as ours transition to equally- or lower-paying job in the CPS data. This figure likely exaggerates the probability for our younger workers, who presumably have have rising pay trajectories.

Table A.14 (columns 3-4) supports this observation by showing that the minimum wage does not affect with the outside option (OOI) for either gender.

F Welfare Appendix

This section begins with a discussion of the model and concludes with an examination of the robustness of the welfare calibration presented in the paper.

F.1 Model

F.1.1 Modeling details and proof of Proposition 1

The function $c(e)$ is strictly increasing in e. We assume that the marginal cost of effort vanishes at zero and is infinite at 1; these assumptions help ensure that optimal effort is interior to $[0, 1]$. Worker performance (in our case, sales per hour) is a non-negative random variable $Y(e)$ that is uniformly bounded from above across all e. Its density $f_Y(y; e)$ has interval support, is twice continuously differentiable in both its arguments, and enjoys the strict monotone likelihood ratio property (MLRP) in e.[51]

Assumption 1 (concavity). $u_{ee} < 0$ *and* $\pi_{ee} \leq 0$.

Under Assumption 1, the worker's optimal effort $e^*(M)$ is the unique solution to the first-order condition:

$$u_e(e; M) + \frac{1}{(1+r)}\pi'(e)V(M) = 0. \tag{15}$$

Proof of Proposition 1

Proof. Fix M. The function $u(e; M)$ shifts down if $u^U(M)$ increases, and thus also if $V^U(M)$ increases. Coviello et al (2022, Lemma 3 part 2 in Online Appendix B) shows that if the function $u(e; M)$ shifts down, the worker's net value $V(M)$ decreases. Therefore, as $u^U(M)$ increases, both functions of e in (15) shift down, hence $e^*(M)$ decreases. ■

[51]This means that the ratio $f_Y(y; e)/F_Y(y; e)$ is strictly increasing in e whenever $f > 0$.

F.1.2 Computing expression (8)

Differentiating (7) with respect to M yields:

$$\frac{dV(M)}{dM} = (1{+}r)\left[\frac{u_e(e^*; M)\,[1+r-\pi(e^*)] + u(e^*; M)\cdot\pi'(e^*)}{[1+r-\pi(e^*)]^2}\cdot\frac{de^*}{dM} + \frac{u_M(e^*; M)}{1+r-\pi(e^*)}\right].$$

This formula simplifies because the numerator of the first fraction inside the brackets is zero. Indeed, substituting (7) into the first-order conditions yields:

$$u_e(e^*; M) + \frac{\pi'(e^*)}{1+r-\pi(e^*)}\cdot u(e^*; M) = 0.$$

Therefore equation (8) holds.

F.1.3 Computing expression (9)

We have:

$$\begin{aligned}
\frac{dV^E(M)}{dM} &= \frac{dV(M)}{dM} + \frac{dV^U(M)}{dM} \\
&= \left[\frac{(1+r)}{1+r-\pi(e^*)}\right] u_M(e^*; M) + \left[\frac{(1+r)}{r}\right] u_M^U(M) \\
&= \left[\frac{(1+r)}{1+r-\pi(e^*)}\right] \left[w_M(e^*; M) - u_M^U(M)\right] + \left[\frac{(1+r)}{r}\right] u_M^U(M) \\
&= \left[\frac{(1+r)}{1+r-\pi(e^*)}\right] w_M(e^*; M) + \left[\frac{(1+r)}{r} - \frac{(1+r)}{1+r-\pi(e^*)}\right] u_M^U(M),
\end{aligned}$$

where the second line used the definition

$$u^U(M) = [r/(1+r)]V^U(M).$$

Substituting:

$$
\begin{aligned}
& \frac{(1+r)}{r}-\frac{(1+r)}{1+r-\pi\left(e^*\right)} \\
= & \frac{(1+r)\left[1+r-\pi\left(e^*\right)\right]-r(1+r)}{r\left[1+r-\pi\left(e^*\right)\right]} \\
= & \frac{(1+r)\left[1-\pi\left(e^*\right)\right]}{r\left[1+r-\pi\left(e^*\right)\right]},
\end{aligned}
$$

we get:

$$
\begin{aligned}
\frac{dV^E(M)}{dM} & =\left[\frac{(1+r)}{1+r-\pi\left(e^*\right)}\right] w_M\left(e^* ; M\right)+\frac{(1+r)\left[1-\pi\left(e^*\right)\right]}{r\left[1+r-\pi\left(e^*\right)\right]} u_M^U(M) \\
& =\left[\frac{(1+r)}{1+r-\pi\left(e^*\right)}\right]\left[w_M\left(e^* ; M\right)+\frac{\left[1-\pi\left(e^*\right)\right]}{r} u_M^U(M)\right] .
\end{aligned}
$$

F.2 Empirical findings: robustness of the welfare effects of the minimum wage by gender

We assess the sensitivity of our welfare results to using different calculations of π, values for r, and assumptions on the outside option. The results are presented in Table A.15.

Columns 2-4 show that the results are not sensitive to how π is calculated. The results are similar if we calculate $\pi\left(e^*\right)$ across all time periods (rather than month by month), by gender (column 2). They are also similar if we calculate $\pi_d\left(e_{t-1}^*\right)$ each month, and aggregated at the gender$\times$department level (column 3) or at the gender$\times$department$\times$store level (column 4).

Columns 5-6 show that the results are very similar if we use a monthly discount rate r of 0.41% or 1.5%. The former corresponds to an annual rate of 5% and the latter to a quarterly $\beta=1/1+r$ of 0.96, commonly used in macroeconomics calibrations.

Columns 7-9 present the sensitivity of our results to different assumptions on the outside option $u_M^U(M_{jt})$. Recall that in our main welfare calculations we set $u_M^U(M_{jt})$ to zero. Results remain very similar if we remain agnostic about $u_M^U(M_{jt})$ and run specification (1) with the following outcome variable:

$$\left[\frac{(1+r)}{1+r-\pi\left(e^{*}\right)}\right]\left[w_{M}\left(e_{i,t-1}^{*};M_{jt}\right)+\frac{\left[1-\pi\left(e^{*}\right)\right]}{r}u^{U}\left(M_{jt}\right)\right], \tag{16}$$

where u^{U} is measured with the gender-specific outside option index (OOI) – see Appendix D.1 for how we construct this measure.