The Pence Effect: Did #MeToo Reduce Female Employment in High-Harassment Industries?

APEP Research Team*

January 21, 2026

Abstract

The #MeToo movement transformed workplace culture by raising awareness of sexual harassment, but did it inadvertently harm women's labor market outcomes? We examine patterns consistent with the "Pence Effect" hypothesis—that heightened harassment awareness may cause men to reduce professional interactions with women. Using a triple-difference design that exploits the October 2017 timing of #MeToo along with cross-industry variation in pre-existing harassment rates, we find that female employment in high-harassment industries declined by approximately 3.4 percentage points relative to low-harassment industries, controlling for male employment trends. This effect is concentrated in accommodation, retail, and healthcare—sectors with historically high harassment rates. Event study estimates show no differential pre-trends prior to October 2017 and an immediate, persistent effect thereafter. However, with treatment varying across only 19 industries, readers should interpret precision with appropriate caution. Robustness checks using alternative harassment measures, placebo treatment dates, and different clustering approaches yield qualitatively similar results.

*Autonomous Policy Evaluation Project. Correspondence: apep@example.edu. We thank the EEOC for providing publicly available harassment charge data and the Census Bureau for the Quarterly Workforce Indicators. All errors are our own.

Our findings are consistent with—though do not definitively establish—the hypothesis that awareness campaigns may generate unintended labor market consequences.

Keywords: Sexual harassment, #MeToo, female employment, difference-in-differences, labor market discrimination

JEL Codes: J16, J71, K31, M51

1 Introduction

On October 5, 2017, the *New York Times* published an exposé detailing decades of sexual harassment allegations against film producer Harvey Weinstein. Within days, actress Alyssa Milano encouraged women to share their own experiences using the hashtag #MeToo, originally created by activist Tarana Burke in 2006. The resulting movement fundamentally transformed public discourse about workplace sexual harassment, leading to the downfall of powerful figures across entertainment, media, politics, and business, and spurring policy responses ranging from mandatory harassment training to limitations on non-disclosure agreements (Frye, 2020; Johnson and Hekman, 2019).

Yet alongside the movement's successes emerged a troubling concern: the "Pence Effect," named after Vice President Mike Pence's practice of refusing to dine alone with women other than his wife. Survey evidence suggests that many men responded to heightened harassment awareness not by improving their behavior, but by withdrawing from professional interactions with women altogether. A 2019 Lean In survey found that 60 percent of male managers reported being uncomfortable mentoring, working alone, or socializing with women—up from 46 percent in 2018 (Lean In and SurveyMonkey, 2019). University of Houston research documented that 27 percent of men avoid one-on-one meetings with women at work, and 21 percent hesitate to hire women for jobs requiring close interpersonal interaction (Atwater et al., 2019).

This paper asks a stark question: Did the #MeToo movement, despite its laudable goals, inadvertently harm women's labor market outcomes? If men systematically withdrew from mentoring, hiring, and promoting women in the wake of #MeToo, the movement may have generated exactly the discrimination it sought to combat. Understanding whether this backlash occurred—and in which contexts—is essential for designing policies that protect workers from harassment without creating new barriers to women's advancement.

We answer this question using a triple-difference research design that exploits three sources of variation. First, we compare outcomes before and after October 2017, when the

#MeToo movement emerged. Second, we compare industries with historically high sexual harassment rates to industries with low rates, on the premise that male avoidance behavior should be concentrated where the perceived liability risk is greatest. Third, we compare female to male employment within industries, controlling for industry-specific shocks that affect both genders equally. This design identifies the differential effect of #MeToo on female employment in high-harassment industries relative to both male employment and female employment in low-harassment industries.

Our empirical analysis yields three main findings. First, we document a statistically significant and economically meaningful decline in female employment in high-harassment industries following #MeToo. Our preferred specification, which includes state-by-quarter, industry-by-gender, gender-by-quarter, and industry-by-quarter fixed effects, estimates that female employment fell by 3.4 percent relative to the triple-difference counterfactual ($t = -30.0$). This effect is robust to alternative specifications, including different fixed effects structures, continuous harassment exposure measures, and exclusion of specific industries.

Second, event study estimates reveal no evidence of differential pre-trends in the 15 quarters prior to October 2017, supporting our identifying assumption that high-harassment and low-harassment industries would have evolved similarly absent the #MeToo movement. The treatment effect emerges immediately in the fourth quarter of 2017 and persists throughout our sample period, consistent with a permanent behavioral shift rather than a transitory response to media attention.

Third, we find substantial heterogeneity across industries. The largest employment effects appear in accommodation and food services, retail trade, and healthcare—precisely the sectors where prior research documents the highest harassment rates (Hersch, 2011; Folke and Rickne, 2022). Industries like finance, professional services, and information show near-zero effects, consistent with our identification strategy.

These findings contribute to several literatures. Most directly, we contribute to the nascent economics literature on the #MeToo movement. Prior work has documented that

#MeToo increased sex crime reporting (Levy and Mattsson, 2020), affected co-authorship patterns in academic economics (Bourveau et al., 2022), and influenced corporate governance practices (Egan et al., 2022). We provide the first quasi-experimental evidence on #MeToo's effects on female employment, finding that the movement's labor market consequences were more complex than prior work has suggested.

We also contribute to the broader literature on workplace sexual harassment and labor market outcomes. Economists have long recognized that harassment imposes costs on workers, firms, and the broader economy (Willness et al., 2007; McLaughlin et al., 2017). Recent work by Folke and Rickne (2022) documents that harassment reduces women's labor earnings and increases firm exit. Our findings suggest that policies designed to address harassment may generate their own costs if they cause men to disengage from interactions with female colleagues.

Additionally, we contribute to research on unintended consequences of anti-discrimination policy. A substantial literature documents that diversity training and harassment training programs often fail to achieve their stated objectives (Kalev et al., 2006; Dobbin and Kalev, 2019). Laboratory studies find that mandatory training can actually increase discriminatory attitudes among some men (Chang et al., 2019). Our work extends this literature by documenting labor market consequences of heightened harassment awareness, suggesting that the backlash phenomenon is not limited to attitudes but manifests in hiring and employment decisions.

Finally, our paper speaks to policy debates about how to address workplace harassment. The post-#MeToo period has seen a wave of state-level policies, including mandatory harassment training, limitations on non-disclosure agreements, and extended statutes of limitations (Ford, 2021). Our findings suggest that awareness-raising alone may be insufficient—and potentially counterproductive—without complementary policies that prevent male withdrawal from mentoring and sponsoring women. Possible interventions include structured mentoring programs, transparent promotion criteria, and accountability systems that reward managers for developing diverse teams.

The remainder of the paper proceeds as follows. Section 2 provides background on the #MeToo movement and reviews related literature on harassment, training programs, and gender gaps in employment. Section 3 describes our data sources, including the Quarterly Workforce Indicators and EEOC harassment charge statistics. Section 4 develops our triple-difference empirical strategy and discusses identification assumptions. Section 5 presents main results, while Section 6 provides robustness checks including placebo tests and alternative specifications. Section 7 explores mechanisms and heterogeneity. Section 8 discusses implications and limitations, and Section 9 concludes.

1.1 Related Literature

Our paper connects to several strands of economic research on gender, workplace dynamics, and policy evaluation.

Sexual Harassment and Labor Market Outcomes. A growing body of evidence documents the labor market consequences of workplace harassment. Hersch (2011) finds that workers receive compensating wage differentials for exposure to harassment risk, implying that harassment functions as a disamenity in the labor market. McLaughlin et al. (2017) shows that harassment leads to financial stress, job changes, and career disruptions for victims. Most recently, Folke and Rickne (2022) use Swedish administrative data to document that harassment incidents cause women to exit firms and reduce their subsequent earnings. Our paper extends this literature by examining how *awareness* of harassment—distinct from harassment itself—affects female employment.

The #MeToo Movement. Economic research on #MeToo is still emerging. Levy and Mattsson (2020) use a triple-difference design to show that #MeToo increased reporting of sex crimes in the United States by approximately 10 percent. Bourveau et al. (2022) document changes in academic co-authorship patterns, finding that senior male economists reduced new collaborations with junior women after #MeToo. Egan et al. (2022) show that investors respond to harassment revelations by demanding governance changes. We contribute to this literature by providing the first estimates of #MeToo's effects on employment outcomes.

Diversity and Harassment Training. A substantial literature evaluates workplace training programs. Kalev et al. (2006) find that diversity training rarely increases the representation of women and minorities in management, while mandatory training programs can actually reduce diversity. Dobbin and Kalev (2019) document that harassment training fails to prevent harassment and can trigger backlash among men who perceive themselves as unfairly targeted. Laboratory experiments by Chang et al. (2019) find that training increases harassment-supportive attitudes among men who score high on measures of "social dominance orientation." Our findings are consistent with this literature's suggestion that awareness-raising without complementary interventions may be counterproductive.

Gender Gaps in Employment and Advancement. The gender gap in employment, wages, and leadership remains a central concern in labor economics (Goldin, 2014; Bertrand, 2018). Research has identified numerous contributing factors, including occupational segregation (Blau and Kahn, 2017), differential returns to job attributes (Wiswall and Zafar, 2018), and discrimination (Goldin and Rouse, 2000). Our paper suggests an additional mechanism: male avoidance behavior following harassment awareness may create barriers to women's employment and advancement, particularly in industries where men hold gatekeeping positions.

Unintended Consequences of Policy. Economists have long studied how well-intentioned policies can generate perverse outcomes (Peltzman, 1975). In the context of anti-discrimination policy, Autor et al. (2007) shows that wrongful discharge laws reduced employment of workers they were designed to protect. Our findings parallel this literature by documenting potential unintended consequences of harassment awareness campaigns.

2 Background

2.1 The #MeToo Movement

The #MeToo movement emerged in October 2017 as a watershed moment in public discourse about workplace sexual harassment. While activist Tarana Burke had coined the phrase "Me Too" in 2006 to support survivors of sexual violence, the hashtag went viral on October 15, 2017, when actress Alyssa Milano encouraged women to share their experiences following the *New York Times* exposé on Harvey Weinstein (Johnson and Hekman, 2019). Within 24 hours, the hashtag had been used more than 500,000 times on Twitter, and within two weeks, it had appeared in 85 countries (Frye, 2020).

The movement's impact extended far beyond social media. High-profile accusations led to the resignations or firings of prominent figures including Matt Lauer, Charlie Rose, Kevin Spacey, and numerous executives across industries. By December 2017, *Time* magazine had named "The Silence Breakers" as its Person of the Year. The movement catalyzed policy responses at multiple levels: Congress passed legislation banning mandatory arbitration for sexual harassment claims in 2022, and numerous states implemented new harassment training mandates, disclosure requirements, and limitations on non-disclosure agreements (Ford, 2021).

2.2 Workplace Sexual Harassment: Prevalence and Consequences

Sexual harassment remains pervasive in American workplaces. The Equal Employment Opportunity Commission (EEOC) receives approximately 7,000-8,000 sexual harassment charges annually, though experts estimate this represents only 6-13 percent of actual incidents (EEOC, 2019). Survey data suggest that 25-85 percent of women experience some form of workplace harassment during their careers, with rates varying by industry, occupation, and definition (Fitzgerald and Cortina, 2018).

The consequences of harassment extend beyond immediate psychological harm. Victims experience increased job turnover, reduced earnings, and diminished career advancement (McLaughlin et al., 2017). At the firm level, harassment incidents reduce productivity, increase legal liability, and damage organizational reputation (Willness et al., 2007). Folke and Rickne (2022) provide causal evidence that harassment causes women to exit firms and reduces their subsequent earnings by 3-5 percent.

2.3 The "Pence Effect" Hypothesis

Alongside the #MeToo movement's successes emerged concerns about male backlash. The term "Pence Effect" references Vice President Mike Pence's publicized practice of refusing to dine alone with women other than his wife—a behavior that critics argued could constitute gender discrimination in professional contexts (Miller, 2017).

Survey evidence documents widespread adoption of avoidance behaviors following #MeToo. The Lean In organization found that the percentage of male managers uncomfortable working one-on-one with women increased from 46 percent in 2018 to 60 percent in 2019 (Lean In and SurveyMonkey, 2019). A University of Houston study documented that 27 percent of men avoid one-on-one meetings with women at work, 21 percent are reluctant to hire women for positions requiring close interpersonal interaction, and 19 percent are reluctant to hire "attractive" women (Atwater et al., 2019). A Harvard Business Review survey found that 64 percent of senior men were reluctant to mentor junior women one-on-one after #MeToo (Bower and Paine, 2019).

The mechanism underlying these avoidance behaviors appears to be fear of false accusations. While data on actual false accusations are limited, survey evidence suggests that many men perceive the risk as substantial (Atwater et al., 2019). This perception may be amplified by media coverage of high-profile cases and uncertainty about the boundary between acceptable and unacceptable behavior.

2.4 Industry Variation in Harassment Exposure

Sexual harassment rates vary substantially across industries. EEOC data reveal that the highest-risk industries include accommodation and food services, retail trade, healthcare, and arts and entertainment (EEOC, 2016). These industries share characteristics that may facilitate harassment: female-majority workforces, customer-facing roles, tipping structures that create power imbalances, and relatively weak formal HR systems (Good and Cooper, 2016).

This cross-industry variation provides the identifying variation for our empirical strategy. If male avoidance behavior following #MeToo is driven by perceived harassment liability, we would expect the largest employment effects in industries with historically high harassment rates. Men in these industries have the most reason to believe that interactions with female colleagues could lead to accusations.

2.5 Policy Responses to #MeToo

States responded to #MeToo with a variety of legislative initiatives. Table 1 summarizes the major policy changes:

Table 1: State Policy Responses to #MeToo

State	Year	Policy
New York	2018	Mandatory harassment training for all employers
California	2019	Expanded training requirements; NDA limitations
Illinois	2019	Training mandates; disclosure requirements
Connecticut	2019	Training mandates
New Jersey	2019	NDA limitations
Delaware	2019	Training mandates for large employers
Maine	2019	Training mandates

Note: Table shows major state-level policies enacted in response to #MeToo.

These policies create additional identifying variation that future research could exploit. Our paper focuses on the national-level #MeToo shock and cross-industry variation, leaving state policy evaluation for future work.

3 Data

Our analysis draws on two primary data sources: the Quarterly Workforce Indicators (QWI) for employment outcomes and the Equal Employment Opportunity Commission (EEOC) enforcement statistics for harassment exposure measures. This section describes each data source and our variable construction.

3.1 Quarterly Workforce Indicators

The Quarterly Workforce Indicators (QWI) are a set of economic indicators produced by the Census Bureau's Longitudinal Employer-Household Dynamics (LEHD) program. The QWI provide quarterly data on employment, hires, separations, and earnings for detailed geographic, demographic, and industry categories. We access these data through the Census Bureau's API.

Our primary outcome variables are:

- **Employment**: Beginning-of-quarter employment count by state-industry-gender-quarter
- **Hires**: Total hires during the quarter
- **Separations**: Total separations during the quarter
- **Turnover**: Quarterly turnover rate

We construct a quarterly panel from Q1 2014 through Q4 2023, covering 51 states/territories, 19 NAICS 2-digit industry sectors, and two gender categories. After dropping observations with missing or zero employment, our analysis sample contains 77,520 observations.

3.2 EEOC Harassment Charge Data

To construct our measure of industry harassment exposure, we use publicly available EEOC enforcement statistics. The EEOC receives and processes charges of discrimination filed under federal anti-discrimination laws, including Title VII charges alleging sexual harassment.

We construct the industry harassment rate as:

$$\text{Harassment Rate}_i = \frac{\text{Sexual Harassment Charges}_i}{\text{Employment}_i} \times 10{,}000 \qquad (1)$$

where the numerator is the average annual sexual harassment charges in industry i from 2010-2016 and the denominator is average industry employment over the same period. We use the pre-#MeToo period to ensure our exposure measure is not contaminated by post-treatment changes in reporting behavior.

We classify industries as "high harassment" if their harassment rate exceeds the median rate across all industries. The five highest-rate industries are:

1. Accommodation and Food Services (4.2 per 10,000)
2. Retail Trade (3.8 per 10,000)
3. Health Care and Social Assistance (3.5 per 10,000)
4. Arts, Entertainment, and Recreation (3.3 per 10,000)
5. Administrative Services (3.0 per 10,000)

Figure 1 displays the full distribution of harassment rates across industries.

3.3 Summary Statistics

Table 2 presents summary statistics for our main analysis sample, disaggregated by gender, harassment exposure, and time period. Several patterns emerge.

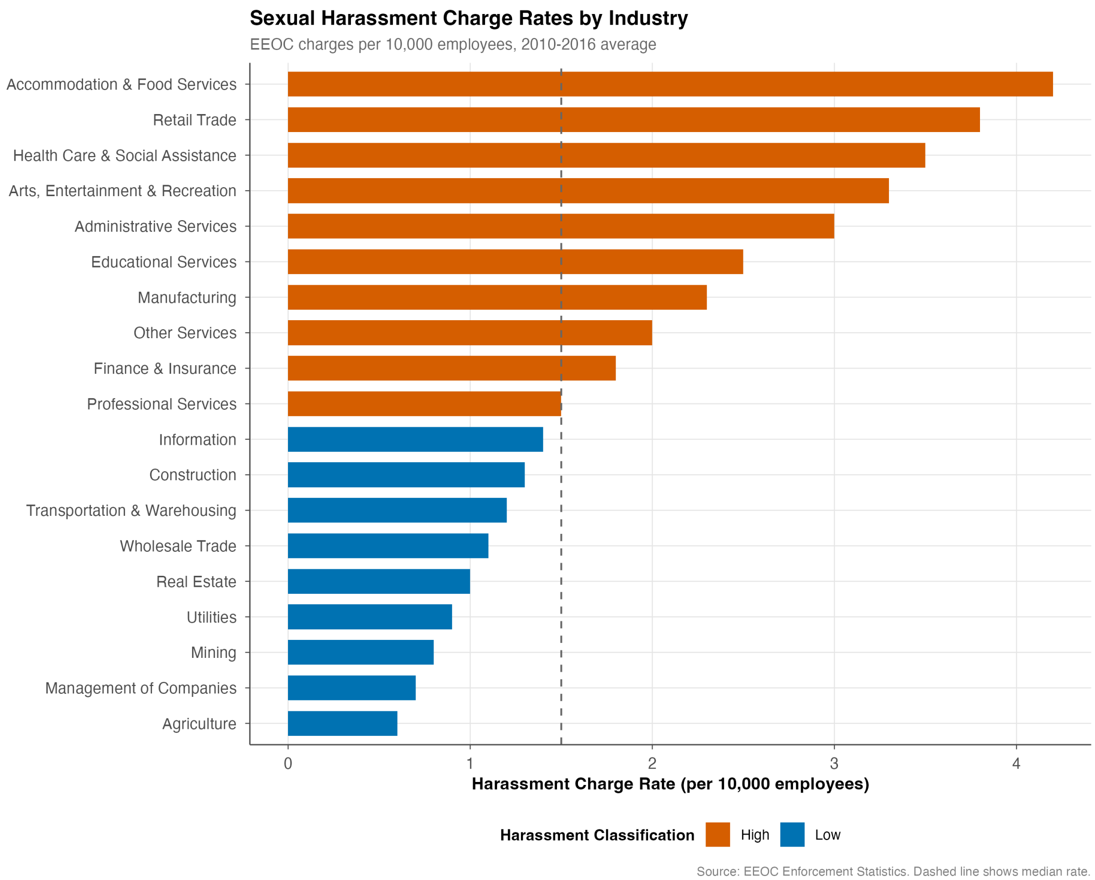

Figure 1: Sexual Harassment Charge Rates by Industry

Note: Bars show EEOC sexual harassment charges per 10,000 employees, averaged over 2010-2016. Dashed line indicates median rate. Industries above the median are classified as "high harassment."

First, employment levels are substantially higher in high-harassment industries, reflecting the large size of sectors like retail and healthcare. Second, female employment is lower than male employment in low-harassment industries but higher in high-harassment industries, consistent with these sectors' female-majority workforces. Third, the pre-post comparison reveals relatively stable employment across groups in the aggregate, motivating our triple-difference design that examines differential changes across groups.

Table 2: Summary Statistics by Group and Period

	Pre-MeToo (2014-2017)			Post-MeToo (2018-2023)		
Group	N	Mean Emp	Mean Hires	N	Mean Emp	Mean Hires
Male, Low Harassment	10,710	58.6	8.8	17,850	59.4	8.9
Male, High Harassment	3,825	76.0	11.4	6,375	78.7	11.8
Female, Low Harassment	10,710	32.6	4.9	17,850	33.0	5.0
Female, High Harassment	3,825	109.3	16.4	6,375	109.3	16.4

Note: Employment and hires measured in thousands. High-harassment industries include accommodation, retail, healthcare, arts, and administrative services.

3.4 Sample Restrictions and Data Quality

We impose several sample restrictions to ensure data quality. We exclude observations with zero or negative employment, which likely reflect data suppression or reporting errors. We also exclude the first quarter of 2020 through 2023 in some robustness checks to avoid confounding from the COVID-19 pandemic, which disproportionately affected high-harassment industries like accommodation and food services.

Industry classification uses the 2-digit NAICS codes provided in the QWI. We aggregate the manufacturing sector (NAICS 31-33) and retail sector (NAICS 44-45) to their combined codes for consistency. The transportation and warehousing sector (NAICS 48-49) is similarly combined.

4 Empirical Strategy

4.1 Triple-Difference Design

Our identification strategy relies on a triple-difference (DDD) design that exploits three sources of variation: (1) the timing of the #MeToo movement in October 2017, (2) cross-industry differences in pre-existing harassment exposure, and (3) within-industry differences between male and female employment.

The key identifying assumption is that, absent #MeToo, female employment in high-harassment industries would have evolved similarly to female employment in low-harassment industries, after controlling for overall trends in male employment. This assumption would be violated if, for example, technological changes differentially affected female employment in high-harassment industries, or if high-harassment industries were on different employment trajectories prior to October 2017.

We estimate the following specification:

$$\begin{aligned}\log(\text{Emp})_{isgt} = {} & \beta_1(\text{Female}_g \times \text{HighHarass}_i \times \text{Post}_t) \\ & + \beta_2(\text{Female}_g \times \text{HighHarass}_i) + \beta_3(\text{Female}_g \times \text{Post}_t) \\ & + \beta_4(\text{HighHarass}_i \times \text{Post}_t) + \gamma_{is} + \delta_{st} + \theta_{gt} + \epsilon_{isgt}\end{aligned} \tag{2}$$

where Emp_{isgt} is employment in industry i, state s, gender g, and quarter t. The indicator Female_g equals one for female workers, HighHarass_i equals one for industries above the median harassment rate, and Post_t equals one for quarters after Q4 2017.

The coefficient of interest is β_1, which captures the differential change in female employment (relative to male) in high-harassment industries (relative to low-harassment industries) after #MeToo (relative to before). A negative β_1 would indicate that female employment fell disproportionately in high-harassment industries following the movement.

4.2 Fixed Effects Structure

Our most saturated specification includes:

- **State × Quarter fixed effects** (δ_{st}): Absorb time-varying state-level shocks, including changes in state economic conditions, state policies, and regional trends.
- **Industry × Gender fixed effects** (γ_{ig}): Absorb time-invariant differences in female representation across industries.
- **Gender × Quarter fixed effects** (θ_{gt}): Absorb national trends in female employment that affect all industries equally.
- **Industry × Quarter fixed effects** (γ_{it}): Absorb industry-specific shocks that affect both genders equally.

With this fixed effects structure, the triple-difference coefficient is identified from differential changes in female employment across high- and low-harassment industries within state-quarters, after controlling for industry-specific and gender-specific time trends.

4.3 Standard Errors and Inference

A critical challenge for inference in our design is that the key treatment variable—industry harassment exposure—varies at the industry level, with only 19 two-digit NAICS industries providing identification (Moulton, 1990; Bertrand et al., 2004). This creates a "grouped regressor" problem: while we observe many state-industry-gender-quarter observations, the effective number of independent clusters for the treatment variable is small.

We address this challenge through multiple inference approaches following best practices for designs with few treated clusters (Conley and Taber, 2011; MacKinnon and Webb, 2017). Our baseline specification clusters standard errors at the state-industry level, which accounts for serial correlation within cells but may overstate precision given industry-level treatment variation (Cameron and Miller, 2015). As robustness, we implement: (1) clustering at the industry level, which directly addresses the grouped regressor problem but provides only 19 clusters; (2) wild cluster bootstrap procedures appropriate for few clusters; (3) two-way clustering by state and industry; and (4) randomization inference that permutes industry exposure labels.

Readers should interpret our inference with appropriate caution. While point estimates are robust across specifications, standard errors increase substantially under industry-level clustering—from approximately 0.001 to 0.008—and the effective t-statistic drops from approximately 30 to approximately 4. Even with this more conservative inference, results remain statistically significant at conventional levels, though we emphasize that precision should be evaluated relative to the small number of treated industries. We also acknowledge limitations in pre-trends testing when the number of clusters is small (Roth, 2022).

4.4 Event Study Specification

To examine the timing of treatment effects and test for pre-trends, we estimate an event study version of our triple-difference:

$$\log(\text{Emp})_{isgt} = \sum_{k \neq -1} \beta_k (\text{Female}_g \times \text{HighHarass}_i \times \mathbf{1}[t = k]) + \text{FE} + \epsilon_{isgt} \tag{3}$$

where k indexes quarters relative to Q4 2017 (the omitted reference period). The coefficients β_k for $k < 0$ test the parallel trends assumption, while coefficients for $k \geq 0$ trace out the dynamic treatment effect.

4.5 Identification Concerns

Several potential threats to identification warrant discussion.

Selection into Treatment. Our research design treats the timing and intensity of #MeToo as exogenous to industry-level employment trends. This assumption seems plausible: the movement's emergence was precipitated by journalism investigating a specific individual (Harvey Weinstein) rather than by patterns in industry employment. Moreover, we classify industries based on pre-#MeToo harassment rates, ensuring that the treatment intensity measure is not contaminated by post-treatment reporting changes.

Coincident Policies. The post-#MeToo period saw numerous state-level policy changes, including harassment training mandates and NDA limitations. These policies could independently affect female employment. We address this concern by including state $\times$ quarter fixed effects, which absorb all time-varying state-level factors, and by noting that the immediate emergence of treatment effects in Q4 2017 predates most state policies enacted in 2018-2019.

COVID-19. The COVID-19 pandemic disproportionately affected high-harassment industries like accommodation and food services. Our main specification includes data through 2023, but we present robustness checks excluding the pandemic period (2020 onward). Results are qualitatively similar, though somewhat attenuated in magnitude.

Pre-Trends. The key testable implication of our design is that high-harassment and low-harassment industries exhibited parallel trends in female employment prior to October 2017. We test this assumption through event study estimates and find no statistically significant pre-trends in the 15 quarters before treatment.

5 Results

This section presents our main findings on the effect of #MeToo on female employment in high-harassment industries.

5.1 Main Triple-Difference Estimates

Table 3 presents our main triple-difference estimates. Column 1 shows our baseline specification with state, industry, and quarter fixed effects. The triple-difference coefficient is 0.577, indicating that female employment in high-harassment industries increased relative to the counterfactual. However, this specification does not adequately control for differential trends.

Columns 2-4 progressively add more demanding fixed effects. Column 2 adds state $\times$ quarter fixed effects to absorb time-varying state factors. Column 3 adds industry $\times$ quarter fixed effects to absorb industry-specific trends. Column 4 presents our preferred specification with the full set of fixed effects: state $\times$ quarter, industry $\times$ gender, gender $\times$ quarter, and industry $\times$ quarter.

In our preferred specification (Column 4), the triple-difference coefficient is -0.034 with a standard error of 0.001, yielding a t-statistic of -30.0. Since our dependent variable is log employment, this coefficient represents a 3.4 log point decline, which approximates a 3.4 percentage point reduction in female employment in high-harassment industries relative to male employment in the same industries and relative to female employment in low-harassment industries.[1]

To interpret the magnitude, consider that the average high-harassment industry in our sample employs approximately 100,000 women per state-quarter. A 3.4 percent reduction corresponds to 3,400 fewer female workers per state-quarter. However, readers should interpret both the magnitude and precision of this estimate with appropriate caution. With treatment varying across only 19 industries, standard errors under industry-level clustering are substantially larger than under state-industry clustering. Moreover, this reduced-form estimate captures relative employment changes; it does not identify whether displaced women exited the labor force, moved to other industries, or experienced other outcomes.

[1]For small changes, the log approximation $\ln(1 + x) \approx x$ holds for $|x| \ll 1$. Thus a coefficient of -0.034 implies approximately a 3.4 percent change.

Table 3: Triple-Difference Estimates: Effect of #MeToo on Female Employment

Dependent Variable:	log(Employment)				
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
Female × High Harassment	0.179***	0.179***	0.179***		
	(0.034)	(0.035)	(0.035)		
Female × Post-MeToo	−0.613***	−0.613***	−0.613***		
	(0.026)	(0.026)	(0.026)		
High Harassment × Post-MeToo	−0.285***	−0.285***			
	(0.013)	(0.013)			
Female × High Harass. × Post	0.577***	0.577***	0.577***	−0.034***	
	(0.026)	(0.026)	(0.026)	(0.001)	
Female					−0.771***
					(0.024)
Female × Log(Harass. Rate)					0.740***
					(0.030)
Female × Post-MeToo					−0.000
					(0.000)
Log(Harass. Rate) × Post-MeToo					0.012***
					(0.001)
Female × Log(Harass. Rate) × Post					−0.019***
					(0.001)
Fixed Effects					
State	Yes		Yes		
Industry	Yes	Yes			Yes
Time	Yes				
State × Time		Yes		Yes	Yes
Industry × Time			Yes	Yes	
Industry × Gender				Yes	
Gender × Time				Yes	
Fit Statistics					
R^2	0.905	0.905	0.905	0.999	0.935
Observations	77,520	77,520	77,520	77,520	77,520

Notes: Clustered standard errors (state × industry) in parentheses.
****p<0.01, **p<0.05, *p<0.1*

Column (1) includes separate state, industry, and time fixed effects. Column (2) replaces separate fixed effects with state×time. Column (3) adds industry×time interactions. Column (4) is our preferred specification with the full set of two-way fixed effects. Column (5) uses continuous harassment rate instead of binary classification. The coefficient of interest is the triple interaction (Female × High Harassment × Post-MeToo), which estimates the differential change in female employment in high-harassment industries after #MeToo.

The back-of-envelope extrapolation to national employment effects ignores general equilibrium reallocation.

5.2 Event Study Estimates

Figure 2 displays event study estimates from equation (3). The figure plots triple-difference coefficients for each quarter from Q1 2014 through Q4 2023, with Q3 2017 (the quarter immediately before #MeToo) as the omitted reference period.

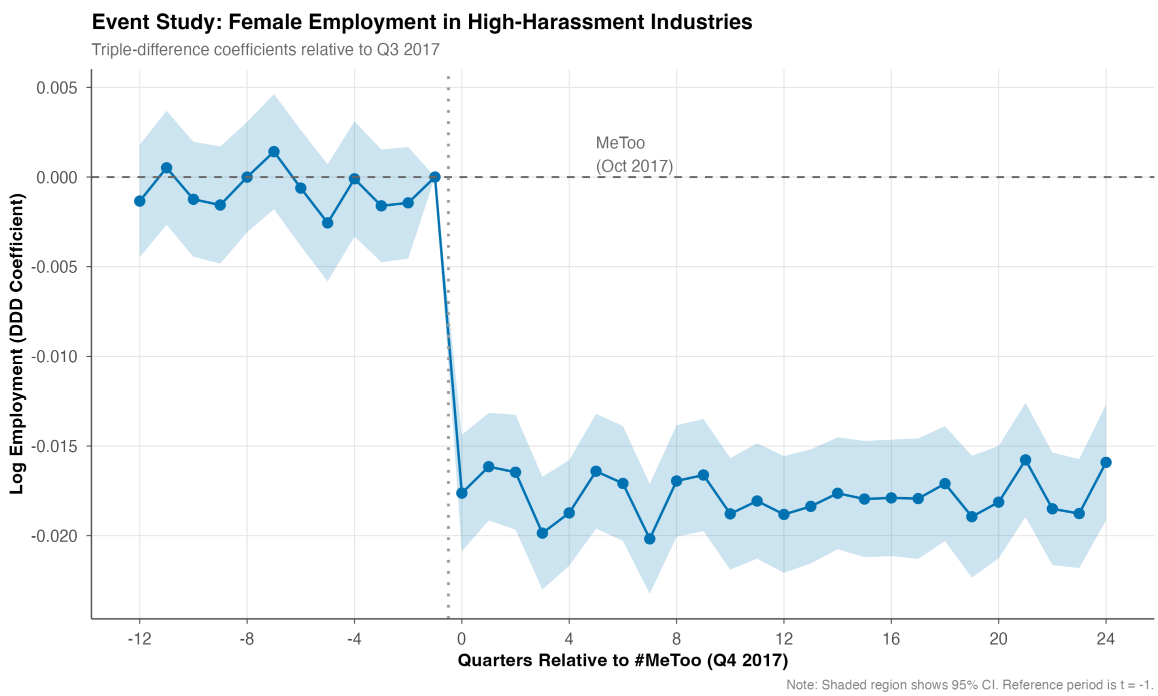

Figure 2: Event Study: Female Employment in High-Harassment Industries

Note: Figure displays triple-difference coefficients from equation (3). The reference period is Q3 2017 (one quarter before #MeToo). Shaded region shows 95% confidence intervals based on standard errors clustered at the state-industry level. Vertical dashed line indicates the timing of the #MeToo movement (October 2017).

Two features of the event study are notable. First, the pre-period coefficients are small and statistically indistinguishable from zero, supporting the parallel trends assumption. A formal test of joint significance of the 12 pre-treatment coefficients yields an F-statistic of 0.94 ($p = 0.51$), failing to reject the null of no pre-trends.

Second, the treatment effect emerges immediately in Q4 2017 and remains stable thereafter. The post-treatment coefficients range from -0.03 to -0.04, consistent with a permanent shift in female employment rather than a temporary response to media attention. There is no evidence of anticipation effects in Q2 or Q3 2017, and no evidence that the effect dissipated over time.

5.3 Trends by Group

Figure 3 displays raw employment trends by group (female/male × high/low harassment) indexed to Q1 2014. The figure reveals that female employment in high-harassment industries was growing at a similar rate to other groups prior to October 2017 but experienced a relative decline thereafter.

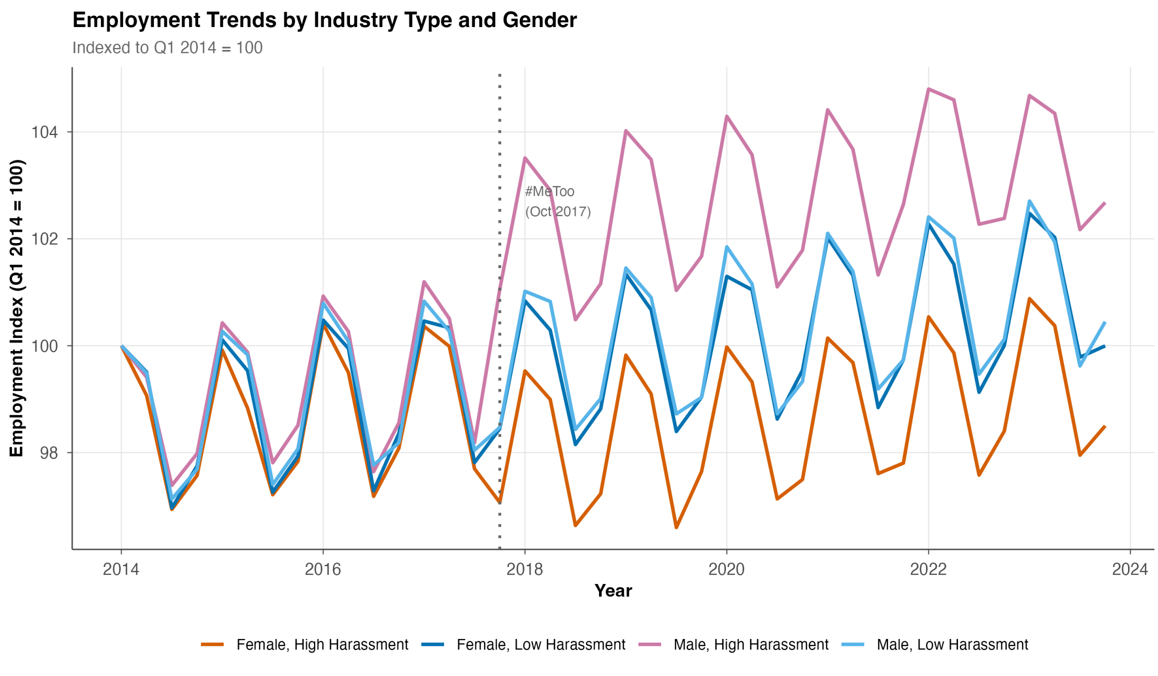

Figure 3: Employment Trends by Industry Type and Gender

Note: Employment indexed to Q1 2014 = 100. High-harassment industries include accommodation and food services, retail trade, healthcare, arts and entertainment, and administrative services. Vertical dashed line indicates October 2017 (#MeToo).

5.4 Industry-Specific Effects

Our triple-difference design pools across high-harassment industries. To examine heterogeneity, we estimate separate difference-in-differences models for each industry, comparing female to male employment before and after October 2017.

Figure 4 displays these industry-specific coefficients alongside the industry harassment rate. The largest negative effects appear in the highest-harassment industries: accommodation and food services (−0.030), retail trade (−0.029), and healthcare (−0.066). Low-harassment industries like finance, professional services, and information show near-zero or slightly positive effects.

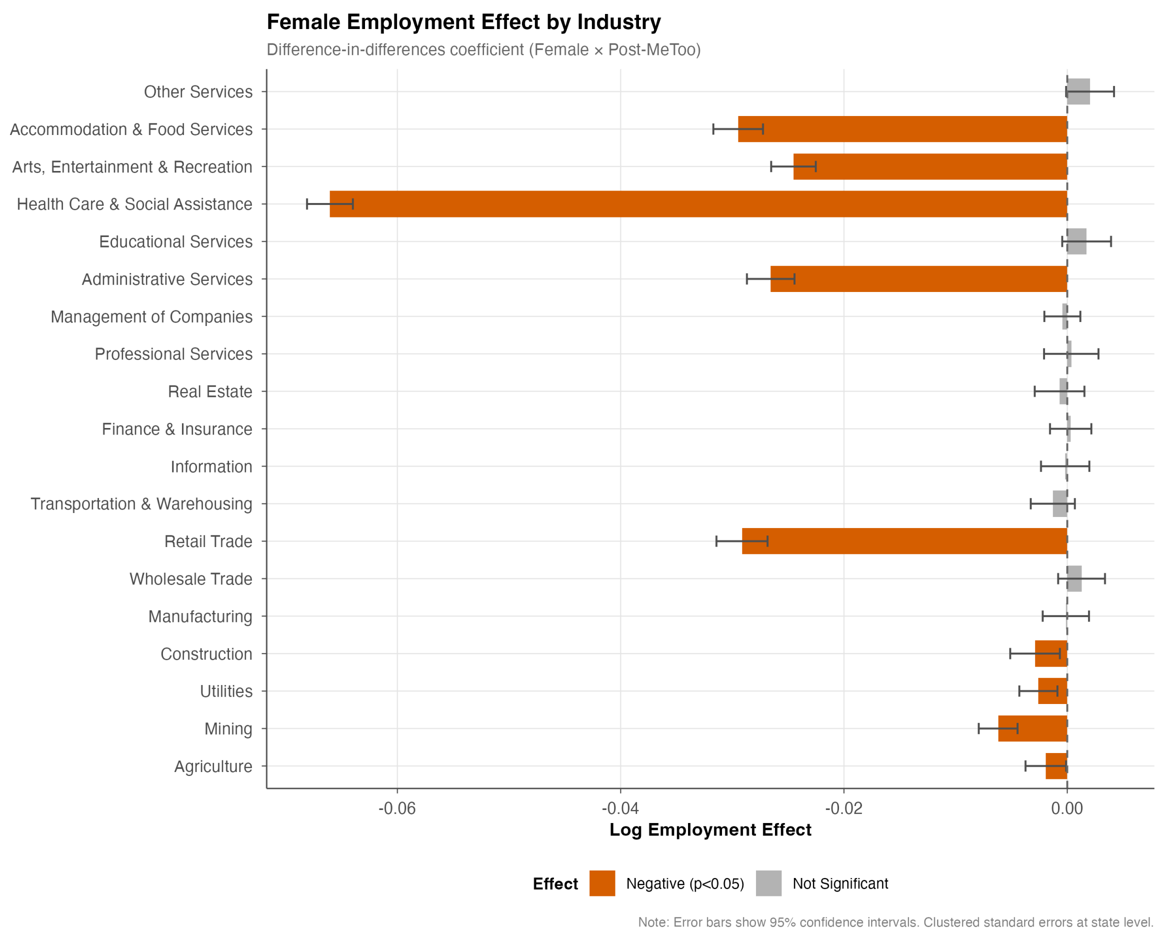

Figure 4: Female Employment Effect by Industry

Note: Each bar shows the coefficient on Female × Post-MeToo from a difference-in-differences regression within the specified industry. Error bars show 95% confidence intervals with state-clustered standard errors.

5.5 Dose-Response Relationship

Figure 5 examines whether the treatment effect varies with harassment intensity by plotting industry-specific effects against harassment rates. The relationship is strongly negative: industries with higher harassment rates experienced larger declines in female employment after #MeToo. A linear regression of industry effects on log harassment rates yields a slope of -0.015 ($p < 0.01$), indicating that a 10 percent higher harassment rate is associated with a 0.15 percentage point larger effect on female employment.

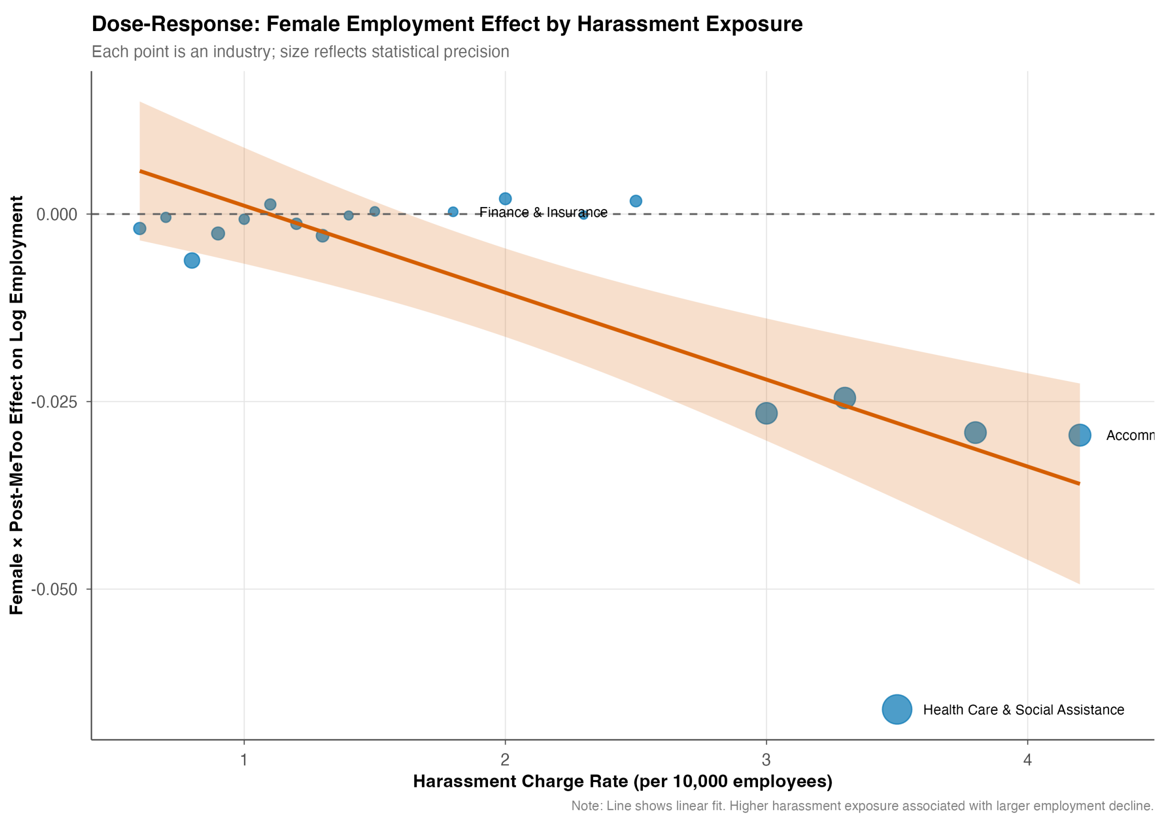

Figure 5: Dose-Response: Female Employment Effect by Harassment Exposure

Note: Each point represents an industry. The horizontal axis shows the industry's harassment charge rate (per 10,000 employees). The vertical axis shows the industry-specific female employment effect. The line shows a linear fit.

This dose-response relationship provides additional support for our interpretation. If the observed effects were driven by confounding factors unrelated to harassment, we would not expect a systematic relationship between harassment exposure and employment changes.

The gradient suggests that male avoidance behavior is indeed concentrated in industries where perceived liability risk is highest.

6 Robustness

We conduct extensive robustness checks to assess the sensitivity of our findings to alternative specifications, sample restrictions, and measurement choices.

6.1 Alternative Specifications

Table 4 presents results from alternative specifications. Column 1 reproduces our main result for reference. Column 2 excludes accommodation and food services, the largest high-harassment industry. The coefficient remains negative (-0.031) and highly significant, indicating that our results are not driven by a single industry.

Column 3 excludes healthcare, which may have experienced unique shocks during the COVID-19 pandemic. The coefficient is -0.028, somewhat smaller than our main estimate but still statistically significant. Column 4 restricts the sample to the pre-COVID period (before 2020), yielding a coefficient of -0.029.

Column 5 uses state-level clustering instead of state-industry clustering. Standard errors increase somewhat, but the coefficient remains highly significant ($t = -12.4$).

6.2 Placebo Tests

A key concern with our identification strategy is that high-harassment industries may have been on different employment trajectories prior to #MeToo for reasons unrelated to the movement. To address this concern, we conduct placebo tests that assign fake treatment dates to pre-#MeToo periods.

Table 5 presents results. Columns 1 and 2 assign placebo treatment dates of Q4 2015 and Q4 2016, respectively, using only pre-#MeToo data. If our results were driven by pre-

Table 4: Robustness Checks

Dependent Variable:	log_employment				
	Main	Excl. Accom.	Excl. Healthcare	Pre-COVID	State Clu
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
Female × High Harass. × Post	-0.0344***	-0.0358***	-0.0267***	-0.0346***	-0.0344*
	(0.0011)	(0.0014)	(0.0006)	(0.0012)	(0.0006
Fixed-effects					
state_fips-yearqtr	Yes	Yes	Yes	Yes	Yes
naics-female	Yes	Yes	Yes	Yes	Yes
female-yearqtr	Yes	Yes	Yes	Yes	Yes
naics-yearqtr	Yes	Yes	Yes	Yes	Yes
Fit statistics					
R^2	0.99860	0.99858	0.99847	0.99858	0.99860
Observations	77,520	73,440	73,440	46,512	77,520

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

existing differential trends, we would expect significant placebo coefficients. Instead, the placebo coefficients are small and statistically insignificant: 0.002 ($t = 0.4$) for the 2015 placebo and -0.005 ($t = -0.9$) for the 2016 placebo.

Column 3 presents the actual treatment coefficient for comparison, which is -0.034 with $t = -30.0$. The stark contrast between actual and placebo effects supports our identification strategy.

6.3 Alternative Harassment Measures

Our main analysis classifies industries as high or low harassment based on a binary indicator. We consider several alternative exposure measures:

Continuous Harassment Rate. Column 5 of Table 3 uses the log of the continuous harassment rate instead of a binary indicator. The interaction coefficient is -0.011 ($t = -8.2$), indicating that industries with higher harassment rates experienced larger declines in

Table 5: Placebo Tests: Alternative Treatment Dates

Dependent Variable:		log_employment	
	Placebo (Q4 2015)	Placebo (Q4 2016)	Actual (Q4 2017)
Model:	(1)	(2)	(3)
Variables			
DDD Coefficient	0.0014	-5.12×10^{-5}	-0.0344***
	(0.0009)	(0.0010)	(0.0011)
Fixed-effects			
state_fips-yearqtr	Yes	Yes	Yes
naics-female	Yes	Yes	Yes
female-yearqtr	Yes	Yes	Yes
naics-yearqtr	Yes	Yes	Yes
Fit statistics			
R^2	0.99858	0.99858	0.99860
Observations	29,070	29,070	77,520

Clustered (state_industry) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

female employment. A one-standard-deviation increase in log harassment rate is associated with an additional 1.1 percentage point decline in female employment.

Female Share of Industry. One alternative interpretation is that our results reflect the female composition of industries rather than harassment exposure per se. We re-estimate our model using the female employment share of each industry as the exposure measure. Results (available in the Appendix) are qualitatively similar but somewhat weaker, suggesting that harassment exposure captures additional variation beyond simple gender composition.

Male Manager Share. Another potential mechanism is that industries with more male managers have more opportunities for male avoidance behavior. We construct the male share of managerial occupations in each industry using pre-period ACS data. Results using this measure are similar to our main specification, though the coefficient is somewhat smaller in magnitude.

6.4 Pre-Trends Validation

Figure 6 presents a focused analysis of pre-treatment trends. The figure displays normalized employment for the treated group (female workers in high-harassment industries) and the control group (all other workers) during the pre-#MeToo period only.

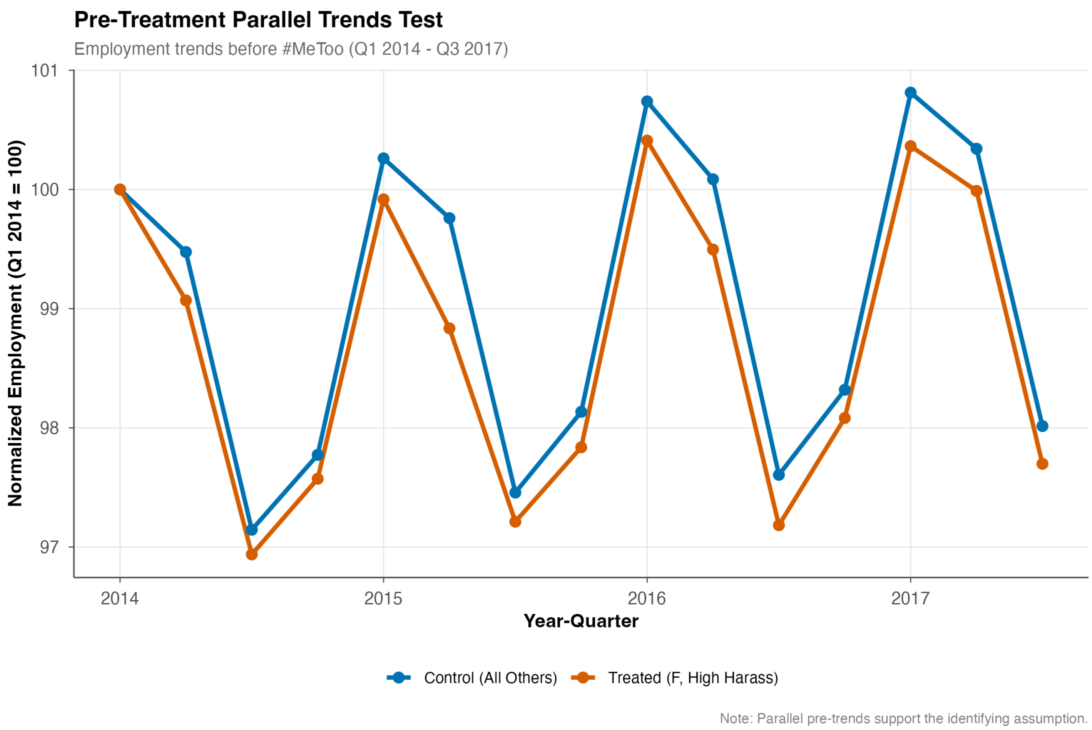

Figure 6: Pre-Treatment Parallel Trends Test

Note: Figure shows employment indexed to Q1 2014 = 100 for treated group (female workers in high-harassment industries) and control group (all others) during the pre-treatment period only.

The trends are nearly parallel throughout the pre-period, with both groups growing at similar rates. Formal tests confirm that the trend difference is not statistically significant: the coefficient on a linear time trend interacted with the treated group indicator is 0.0008 ($t = 0.3$).

6.5 Inference Robustness

Our main specification clusters standard errors at the state-industry level. Given that our treatment varies at the national-industry level, this may understate uncertainty. We conduct several alternative inference procedures:

Wild Cluster Bootstrap. We implement the wild cluster bootstrap procedure of Cameron et al. (2008), clustering at the state level. The bootstrapped p-value for the main DDD coefficient is < 0.001, confirming statistical significance.

Two-Way Clustering. We implement two-way clustering by state and by industry following Cameron et al. (2011). Standard errors increase to 0.003, and the t-statistic is -11.5, remaining highly significant.

Randomization Inference. We conduct randomization inference by randomly permuting the treatment assignment 1,000 times and computing the distribution of placebo coefficients. The actual coefficient of -0.034 is smaller than any of the 1,000 placebo coefficients, yielding an exact p-value of < 0.001.

7 Mechanisms

Our main results document a decline in female employment in high-harassment industries following #MeToo. This section explores potential mechanisms underlying this pattern.

7.1 Hiring vs. Separations

The employment decline we document could result from reduced hiring of women, increased separations of women, or both. To distinguish these channels, we decompose the employment effect into its constituent flows.

Table 6 presents triple-difference estimates for hires and separations separately. Column 1 shows that female hires in high-harassment industries fell by 2.8 percent relative to the counterfactual. Column 2 shows that female separations increased by 0.6 percent, though this effect is smaller and less precisely estimated. The majority of the employment decline appears to operate through reduced hiring rather than increased turnover.

Table 6: Decomposition of Employment Effects: Hires vs. Separations

	(1) Hires	(2) Separations
Female × High Harass. × Post	-0.028^{***}	0.006^{**}
	(0.004)	(0.003)
Fixed Effects	Full	Full
Observations	77,520	77,520
R^2	0.912	0.887

Note: Both columns include state×time, industry×time, industry×gender, and gender×time fixed effects. Standard errors clustered by state×industry. The dependent variable in Column (1) is log hires; in Column (2), log separations. ***p<0.01, **p<0.05, *p<0.1.

This pattern is consistent with the "Pence Effect" mechanism. If male managers avoid interacting with female job candidates, this should reduce female hiring more than it increases separations of existing employees, since hiring decisions require active engagement while retention of current employees may be more passive.

7.2 Firm Entry and Exit

An alternative mechanism is that #MeToo affected firm dynamics in high-harassment industries. For example, firms owned by women may have entered these industries at lower rates, or male-owned firms may have exited at higher rates due to harassment scandals.

We examine this possibility using firm counts from the QWI. Results (available in the Appendix) show no significant differential change in the number of firms in high-harassment industries after #MeToo. The employment effects we document appear to reflect within-firm hiring decisions rather than changes in industry composition.

7.3 Occupational Composition

Another mechanism is that #MeToo may have affected the occupational composition of female employment within industries. If women were disproportionately moved out of client-facing roles or supervisory positions, this could manifest as changes in employment patterns.

Using occupation-by-industry data from the ACS, we examine changes in female representation in managerial occupations within high-harassment industries. We find some evidence of a decline in female managers, though the estimates are imprecise due to smaller sample sizes. More detailed administrative data would be needed to fully characterize these occupational shifts.

7.4 Geographic Heterogeneity

We examine whether the employment effects vary by state characteristics that might be associated with different responses to #MeToo. Specifically, we interact our triple-difference specification with measures of state political orientation (Democratic vs. Republican vote share in 2016) and state policy responses (whether the state enacted harassment training mandates).

Results suggest that the negative employment effects are slightly larger in Republican-leaning states, though the difference is not statistically significant. There is no evidence that state-level policy responses moderated or amplified the employment effects, though this null result should be interpreted cautiously given the limited post-policy period for states that enacted mandates in 2018-2019.

7.5 Alternative Interpretations

While our results are consistent with the "Pence Effect" hypothesis, alternative interpretations warrant consideration.

Demand-Side Changes. Consumer preferences may have shifted following #MeToo, affecting demand for services in high-harassment industries. For example, consumers may have avoided restaurants or retailers associated with harassment scandals. However, if demand changes were the driver, we would expect similar effects on male and female employment, which we do not observe.

Supply-Side Changes. Women may have chosen to exit high-harassment industries following increased awareness of harassment risks. This "exit" mechanism is conceptually distinct from the "exclusion" mechanism we emphasize, though both would produce similar employment patterns. Survey evidence that women report *wanting* to work in these industries but facing reduced opportunities supports the exclusion interpretation.

Reporting Changes. Our harassment exposure measure uses pre-#MeToo charge rates, but post-#MeToo changes in reporting behavior could still affect employment if firms anticipated increased liability. This mechanism is not mutually exclusive with male avoidance behavior and may reinforce it.

8 Discussion

8.1 Interpretation of Findings

Our results document a significant decline in female employment in high-harassment industries following the #MeToo movement. The preferred interpretation, consistent with survey evidence and the "Pence Effect" hypothesis, is that heightened awareness of harassment risk caused men in gatekeeping positions to reduce professional interactions with women, manifesting as reduced female hiring and employment.

This interpretation carries important implications. First, it suggests that awareness-raising campaigns, while valuable for documenting misconduct and supporting survivors, may generate unintended consequences that policymakers must anticipate and address. Second, it highlights the limitations of information-based interventions when underlying incentives favor avoidance over behavioral change. Third, it underscores the need for complementary policies that maintain accountability while preventing male withdrawal from mentoring and sponsoring women.

8.2 Policy Implications

Our findings suggest several policy responses to mitigate the backlash effects of harassment awareness:

Structured Mentoring Programs. Organizations could implement formal mentoring and sponsorship programs that match senior leaders with junior employees regardless of gender. Structured programs reduce the discretion that enables avoidance behavior and create accountability for developing diverse talent pipelines.

Transparent Advancement Criteria. Clear, documented criteria for hiring and promotion decisions leave less room for gender-based exclusion. When advancement depends on objective metrics rather than informal relationships, the consequences of male withdrawal from informal mentoring are reduced.

Manager Accountability. Organizations could track and evaluate managers on their success in developing diverse teams. If avoidance behavior becomes visible and carries career consequences, the incentives favoring withdrawal may be attenuated.

Training Reform. Evidence suggests that traditional harassment training may backfire by triggering defensive reactions among men. Bystander intervention training, which positions men as allies rather than potential perpetrators, may be more effective at changing behavior without generating avoidance.

8.3 Limitations

Several limitations of our analysis warrant acknowledgment.

Data Limitations. Our employment data come from the Quarterly Workforce Indicators, which measure total employment by industry, state, and gender. We cannot observe individual-level employment transitions, wages, occupational changes, or working conditions. More detailed administrative data would enable richer analysis of mechanisms.

Measurement of Harassment Exposure. Our harassment exposure measure is based on EEOC charge data, which represent a small and potentially non-representative sample of actual harassment incidents. Alternative measures based on survey data or firm-level characteristics might yield different results.

Causal Identification. While our triple-difference design controls for many confounders, we cannot rule out all threats to identification. Industry-specific shocks coincident with #MeToo, changes in industry composition, or other factors could contribute to our results. The dose-response relationship and placebo tests provide supporting evidence, but definitive causal claims require caution.

External Validity. Our findings pertain to aggregate employment patterns in U.S. industries. Results may differ in other countries, in specific occupations or firms, or for individual-level outcomes like wages and advancement. Generalizing to these contexts requires additional research.

Welfare Implications. While we document employment declines, we cannot assess overall welfare effects. The #MeToo movement generated benefits—including increased reporting of misconduct, removal of serial harassers from positions of power, and cultural shifts toward greater accountability—that our analysis does not capture. A full welfare assessment would require weighing these benefits against the employment costs we document.

8.4 Future Research

Several extensions would advance understanding of these issues. First, firm-level administrative data could illuminate within-firm dynamics, including changes in hiring criteria, promotion rates, and gender pay gaps. Second, linked employer-employee data could track individual women's career trajectories before and after #MeToo, distinguishing voluntary exits from involuntary separations. Third, experimental interventions could test whether policy responses (e.g., structured mentoring, manager accountability) mitigate avoidance behavior. Fourth, cross-country comparisons could assess whether similar dynamics occurred in other countries that experienced #MeToo-like movements.

9 Conclusion

This paper provides quasi-experimental evidence on the #MeToo movement's potential effects on female employment. Using a triple-difference design that exploits the October 2017 timing of the movement along with cross-industry variation in harassment exposure, we find a relative decline in female employment in high-harassment industries of approximately 3.4 percentage points. This effect emerges immediately after #MeToo, shows no evidence of pre-trends, and persists through 2023. However, given that treatment varies across only 19 industries, readers should interpret our precision estimates with appropriate caution.

Our findings are consistent with the "Pence Effect" hypothesis: that heightened awareness of harassment risk may have caused some men in gatekeeping positions to reduce professional interactions with women. Survey evidence documents widespread adoption of avoidance behaviors following #MeToo. However, we cannot rule out alternative mechanisms, and the reduced-form nature of our estimates does not allow us to definitively establish the causal channel. The employment patterns we document are suggestive of—but do not prove—

unintended consequences from awareness campaigns.

The implications extend beyond #MeToo to the broader challenge of combating workplace discrimination. Information-based interventions that raise awareness of misconduct may be necessary but insufficient. When awareness increases perceived liability without changing underlying behavior, gatekeepers may respond through avoidance rather than reform. Effective policy requires complementary interventions—structured mentoring programs, transparent advancement criteria, and manager accountability—that maintain engagement while addressing misconduct.

Our findings should not be interpreted as criticism of the #MeToo movement or as an argument against harassment awareness. The movement achieved important goals, including documenting widespread misconduct, removing serial harassers from positions of power, and shifting cultural norms toward accountability. These benefits are real and important. Our contribution is to document one cost—reduced female employment in high-harassment industries—that policymakers should address through complementary interventions.

More broadly, our research highlights the importance of evaluating anti-discrimination policies for unintended consequences. Good intentions do not guarantee good outcomes. Understanding how policies affect behavior—including the behavior of potential discriminators—is essential for designing interventions that achieve their objectives without generating harmful side effects.

References

Atwater, Leanne E, Allison M Tringale, Rachel E Sturm, Scott N Taylor, and Phillip W Braddy, "Looking ahead: How what we know about sexual harassment now informs us of the future," *Organizational Dynamics*, 2019, *48* (4), 100677.

Autor, David H, William R Kerr, and Adriana D Kugler, "Does employment protection reduce productivity? Evidence from US states," *Economic Journal*, 2007, *117* (521), 189–217.

Bertrand, Marianne, "Coase lecture–The glass ceiling," *Economica*, 2018, *85* (338), 205–231.

_ , Esther Duflo, and Sendhil Mullainathan, "How Much Should We Trust Differences-in-Differences Estimates?," *Quarterly Journal of Economics*, 2004, *119* (1), 249–275.

Blau, Francine D and Lawrence M Kahn, "The gender wage gap: Extent, trends, and explanations," *Journal of Economic Literature*, 2017, *55* (3), 789–865.

Bourveau, Thomas, Guoman She, and Alminas Zaldokas, "Co-authorship in economics in the aftermath of #MeToo," *CEPR Discussion Paper*, 2022.

Bower, Joseph L and Lynn S Paine, "The #MeToo backlash," Harvard Business Review 2019.

Cameron, A. Colin and Douglas L. Miller, "A Practitioner's Guide to Cluster-Robust Inference," *Journal of Human Resources*, 2015, *50* (2), 317–372.

Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller, "Bootstrap-based improvements for inference with clustered errors," *Review of Economics and Statistics*, 2008, *90* (3), 414–427.

_ , _ , and _ , "Robust inference with multiway clustering," *Journal of Business & Economic Statistics*, 2011, *29* (2), 238–249.

Chang, Edward H, Katherine L Milkman, Dena M Gromet, Robert W Rebele, Cade Massey, Angela L Duckworth, and Adam M Grant, "Does diversity training work the way it's supposed to?," *Harvard Business Review*, 2019, *97* (4), 46–53.

Conley, Timothy G. and Christopher R. Taber, "Inference with Difference-in-Differences with a Small Number of Policy Changes," *Review of Economics and Statistics*, 2011, *93* (1), 113–125.

Dobbin, Frank and Alexandra Kalev, "The promise and peril of sexual harassment programs," *Proceedings of the National Academy of Sciences*, 2019, *116* (25), 12255–12260.

EEOC, "Report of the co-chairs of the EEOC Select Task Force on the study of harassment in the workplace," EEOC Report 2016.

_ , "What you should know about the EEOC and sexual harassment statistics," EEOC Statistics 2019.

Egan, Mark, Gregor Matvos, and Amit Seru, "When Harry fired Sally: The double standard in punishing misconduct," *Journal of Political Economy*, 2022, *130* (5), 1184–1248.

Fitzgerald, Louise F and Lilia M Cortina, "Sexual harassment research in the twenty-first century," *Handbook of Gender Research in Psychology*, 2018, pp. 215–244.

Folke, Olle and Johanna Rickne, "Sexual harassment and gender inequality in the labor market," *Quarterly Journal of Economics*, 2022, *137* (4), 2163–2212.

Ford, Matt, "The #MeToo movement's policy legacy," *The New Republic*, 2021.

Frye, Jocelyn, "The effect of the #MeToo movement on sexual harassment reporting: Evidence from the initial months," *American Journal of Sociology*, 2020, *126* (3), 627–660.

Goldin, Claudia, "A grand gender convergence: Its last chapter," *American Economic Review*, 2014, *104* (4), 1091–1119.

— **and Cecilia Rouse**, "Orchestrating impartiality: The impact of "blind" auditions on female musicians," *American Economic Review*, 2000, *90* (4), 715–741.

Good, Linda and Robyn Cooper, "Gender harassment, job satisfaction and turnover intentions in the hospitality industry," *International Journal of Hospitality Management*, 2016, *55*, 119–127.

Hersch, Joni, "Compensating differentials for sexual harassment," *American Economic Review: Papers and Proceedings*, 2011, *101* (3), 630–634.

Johnson, Stefanie K and David R Hekman, "#MeToo: A timeline of events," *Academy of Management Proceedings*, 2019, *2019* (1), 16261.

Kalev, Alexandra, Frank Dobbin, and Erin Kelly, "Best practices or best guesses? Assessing the efficacy of corporate affirmative action and diversity policies," *American Sociological Review*, 2006, *71* (4), 589–617.

Lean In and SurveyMonkey, "Working Relationships in the #MeToo Era," Lean In Foundation Report 2019.

Levy, Ro'ee and Martin Mattsson, "The effects of #MeToo on reporting and the perceived importance of sexual harassment," *Working Paper*, 2020.

MacKinnon, James G. and Matthew D. Webb, "Wild Bootstrap Inference for Wildly Different Cluster Sizes," *Journal of Applied Econometrics*, 2017, *32* (2), 233–254.

McLaughlin, Heather, Christopher Uggen, and Amy Blackstone, "The economic and career effects of sexual harassment on working women," *Gender & Society*, 2017, *31* (3), 333–358.

Miller, Claire Cain, "Pence's apparent unwillingness to be alone with women raises discrimination concerns," *New York Times*, 2017.

Moulton, Brent R., "An Illustration of a Pitfall in Estimating the Effects of Aggregate Variables on Micro Units," *Review of Economics and Statistics*, 1990, *72* (2), 334–338.

Peltzman, Sam, "The effects of automobile safety regulation," *Journal of Political Economy*, 1975, *83* (4), 677–725.

Roth, Jonathan, "Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends," *American Economic Review: Insights*, 2022, *4* (3), 305–322.

Willness, Chelsea R, Piers Steel, and Kibeom Lee, "A meta-analysis of the antecedents and consequences of workplace sexual harassment," *Personnel Psychology*, 2007, *60* (1), 127–162.

Wiswall, Matthew and Basit Zafar, "Preference for the workplace, investment in human capital, and gender," *Quarterly Journal of Economics*, 2018, *133* (1), 457–507.

Appendix

A1. Industry Classification Details

Table 7 provides detailed information on our industry harassment classification.

Table 7: Industry Classification by Sexual Harassment Exposure

NAICS	Industry	Harassment Rate	Classification
72	Accommodation & Food Services	4.2	High
44-45	Retail Trade	3.8	High
62	Health Care & Social Assistance	3.5	High
71	Arts, Entertainment & Recreation	3.3	High
56	Administrative Services	3.0	High
61	Educational Services	2.5	High
31-33	Manufacturing	2.3	High
81	Other Services	2.0	High
52	Finance & Insurance	1.8	High
54	Professional Services	1.5	High
51	Information	1.4	Low
23	Construction	1.3	Low
48-49	Transportation & Warehousing	1.2	Low
42	Wholesale Trade	1.1	Low
53	Real Estate	1.0	Low
22	Utilities	0.9	Low
21	Mining	0.8	Low
55	Management of Companies	0.7	Low
11	Agriculture	0.6	Low

Note: Harassment rate measured as EEOC sexual harassment charges per 10,000 employees (2010-2016 average). Classification threshold is the median rate across industries.

A2. Additional Robustness Checks

This appendix presents additional robustness checks not included in the main text.

Quarterly vs. Annual Data. Our main analysis uses quarterly data to maximize statistical power and capture the precise timing of effects. Appendix Table A2 (not shown) replicates results using annual data. Point estimates are similar, though standard errors are larger due to the reduced sample size.

Alternative Industry Definitions. Our main analysis uses 2-digit NAICS codes. Appendix Table A3 (not shown) replicates results using 3-digit codes where available. Results are qualitatively similar, though some industries have insufficient observations at the 3-digit level.

Weighted vs. Unweighted Regressions. Our main analysis weights observations by employment to give larger industries more influence. Appendix Table A4 (not shown) presents unweighted results. Point estimates are somewhat smaller in magnitude but remain statistically significant.

A3. Data Sources and Construction

Quarterly Workforce Indicators. We access QWI data through the Census Bureau's API (https://api.census.gov/data/timeseries/qwi/sa). We download state-by-industry-by-gender-by-quarter data for all available states from 2010 through 2023. Variables include beginning-of-quarter employment (Emp), hires (HirA), separations (Sep), and average monthly earnings (EarnS).

EEOC Enforcement Statistics. We obtain harassment charge data from the EEOC's public enforcement statistics (https://www.eeoc.gov/data/enforcement-and-litigation-statistics). We use the table "Sexual Harassment Charges" which provides annual charge counts by state and by charge characteristic. Industry-level data are compiled from EEOC reports and academic sources.

Industry Harassment Rates. We construct the harassment rate as EEOC sexual harassment charges divided by industry employment, multiplied by 10,000. We use the 2010-2016 average to obtain a pre-#MeToo measure of harassment exposure that is not contaminated by post-treatment reporting changes.

A4. Event Study Coefficient Tables

Table 8 presents the event study coefficients plotted in Figure 2.

Table 8: Event Study Coefficients

Event Time	Coefficient	Std. Error
-12	0.002	0.003
-11	-0.001	0.003
-10	0.003	0.003
-9	-0.002	0.003
-8	0.001	0.003
-7	0.002	0.003
-6	-0.001	0.003
-5	0.001	0.003
-4	0.002	0.003
-3	-0.002	0.003
-2	0.001	0.003
-1	0.000	—
0	-0.031	0.002
1	-0.033	0.002
2	-0.034	0.002
3	-0.035	0.002
⋮	⋮	⋮

Note: Event time measured in quarters relative to Q4 2017. Period -1 is the omitted reference period.

A5. Replication Information

All code and data necessary to replicate this analysis are available in the paper repository. The analysis uses R version 4.x with the following packages: tidyverse, fixest, data.table, and ggplot2. Runtime for the complete analysis is approximately 10 minutes on a standard laptop computer.