The Costs of Hidden Workplace Harassment

Abstract

Using natural language processing techniques and millions of anonymous employee reviews, I propose a novel measure of workplace harassment in U.S. firms and investigate its causal effects on innovation and other firm-level outcomes. I exploit the quasi-experimental reduction in workplace harassment caused by changes in non-disclosure agreement (NDA) laws across U.S. states and document a negative impact of workplace harassment on innovation. Firms with previously higher levels of workplace harassment experience a significant increase in their innovation output following these regulations. The documented effect is significantly more pronounced for firms with minority representation in their inventor teams along the gender, race, and ethnicity dimensions. Underlying these effects are improvements in team capital and collaborative dynamics, as the positive changes in the workplace climate lead to significant increases in inventor productivity, as well as retention and entry rates of skilled workers.

Keywords: Workplace harassment, Non-disclosure agreement, Innovation, Productivity, Team capital, Machine learning, Natural language processing

JEL: G34, G38, O3, J24, J28, M14

1 Introduction

Human capital is a key input in various firm activities, especially in innovation (Romer, 1990), which plays an integral role in the success and growth of firms (Aghion and Howitt, 1992; Kogan et al., 2017). Developing competitive products or cutting-edge technologies often depends on close collaboration between teams or individual team members with diverse backgrounds (Glaeser et al., 1992), each contributing their unique skills and expertise to the collective effort. In such an environment, prejudiced remarks and actions about individuals' gender, religion, ethnicity, or identity can impede this creative synergy, and thus harm firms' success in delivering innovative projects (Coviello et al., 2022).[1] Nonetheless, workplace harassment remains a pervasive problem (WBI, 2021), and there is a burgeoning literature analyzing its consequences for victims and perpetrators.[2] However, how such workplace disruptions would affect firm innovation remains an open question.[3]

In this paper, I investigate the costs of workplace harassment for firms, particularly, for their innovation activity, and further explore the economic mechanisms underlying these costs. I define workplace harassment as unwelcome conduct that is based on race, color, religion, sex (including sexual orientation, gender identity, or pregnancy), national origin, older age, disability, or genetic information (including family medical history), in accordance with the U.S. Equal Employment Opportunities Commission (EEOC). My analysis begins with the development of a novel measure of workplace harassment using semi-supervised machine learning methods applied to anonymous online employee reviews of U.S. firms. The measure differs from existing measures of corporate culture, exhibits both cross-sectional and time-series variation, and has economically significant predictive power for future workplace harassment-related lawsuits. I then exploit the firm-level heterogeneity, coupled with plausibly exogenous variation in state-level anti-harassment regulations, to identify the causal effects of workplace harassment on innovation. I find that exogenous decline in workplace harassment leads to significant

[1] In fact, executives rank workplace culture as the one of the most important factors for innovation and long-term firm value (Graham et al., 2022).

[2] For example, see Rospenda et al. (2009); Raver and Nishii (2010); Okechukwu et al. (2014); Schmitt et al. (2014); Batut et al. (2021); Folke and Rickne (2022); Adams-Prassl et al. (2023).

[3] Jaravel et al. (2018) investigate the impact of the premature death of a close co-inventor on the innovation output of the inventor.

gains in both the quality and quantity of firms' innovation output. These effects are significantly more pronounced for companies with inventor teams that include ethnic, racial, and gender minorities on their inventor teams. I further show that this causal relationship can be to a large extent attributed to the maintenance and growth of the collaborative dynamics and inventor team capital, as improvements in workplace climate lead to higher inventor productivity and higher rates of skilled worker retention and entry.

One of the key challenges in investigating the effects of workplace harassment on innovation is the lack of a meaningful measure of workplace harassment that is available for a large set of companies. While one direct approach could be to rely on news articles or lawsuits, these typically cover high-profile incidents or arrive with a large time lag, and are rare because victims of harassment are often silenced through out-of-court settlements and non-disclosure agreements. More importantly, they primarily capture the resolution of past incidents. I overcome these limitations by developing a novel firm-level and time-varying workplace harassment measure for large publicly traded U.S. firms. To construct the measure, I use approximately 3.6 million employee reviews for the period 2009-2022 from Glassdoor, a widely used job review platform with about 55 million monthly unique visitors. A key feature of these reviews is that they are published anonymously, without any identifying information about the reviewers. This gives employees the freedom to openly discuss not only the positive aspects of their workplace, but also the negative aspects, such as toxicity or potential workplace harassment, without fear of retaliation.[4] Therefore, employees provide a credible and insider view of their employers, which is supported by recent literature (Green et al., 2019; Dube and Zhu, 2021).

To construct my measure, I apply word embeddings and the cosine similarity algorithm to the textual content of employee reviews (Li et al., 2021). Unlike traditional methods such as "bag of words", this approach allows me to capture the semantic information contained in unstructured texts such as employee reviews (Goldstein et al., 2021). Specifically, the algorithm learns the meaning of words based on the context in which they are used, providing unique vectors for each word that can be used to assess whether an employee is discussing workplace

[4] In support of this premise, Boudreau et al. (2023) show that maintaining trust and anonymity could help to increase the reporting of harassment incidents in organizations.

harassment in a given review. To do this, I first derive the dictionary of workplace harassment from the textual corpus of all employee reviews by providing seed words that closely follow the EEOC's definition of workplace harassment. Next, I obtain the score of the most similar word in a given review to this dictionary. Finally, I calculate the workplace harassment score of a given firm by averaging these similarity scores across all reviews in a given year.

I ensure the external validity of my measure by performing several tests. First, I provide a number of concrete examples showing that reviews that score high on the harassment measure do indeed discuss instances of workplace harassment and prejudiced behavior. Second, by aggregating the workplace harassment measure at the firm-year level, I show that firms with high workplace harassment scores are less likely to be among the "best places to work" as defined by Edmans (2011), and tend to have lower ESG scores. For a more direct validation, I use auxiliary data on civil rights lawsuits in employment from the Federal Judicial Center. I refer to these lawsuits as harassment-related lawsuits because they often (88% of the cases) contain phrases such as "retaliation" or "harassment" (Dougal et al., 2022).[5] I document that there is a significant positive relationship between the level of workplace harassment and future harassment-related lawsuits. Firms with a one standard deviation higher workplace harassment score are 29% more likely to face harassment-related lawsuits in the next one to five years, suggesting that my measure captures hidden (unreported) harassment.[6]

Having constructed and validated my measure of workplace harassment, I next investigate its effects on firm outcomes, particularly innovation. To do so, I exploit the staggered adoption of regulations that prohibit the misuse of non-disclosure agreements (NDAs) in cases of workplace harassment in thirteen states across the United States beginning in 2018. NDAs, also known as confidentiality agreements, are commonly used by companies to protect their business or trade secrets.[7] However, with increased public awareness about the harassment cases, it has become known that these confidentiality agreements are also used by companies

[5]For example, one such lawsuit was recently filed against Tesla for racial harassment: `https://www.wsj.com/business/tesla-is-accused-of-tolerating-racial-harassment-in-eeoc-suit-ef72f336`

[6]I provide further robustness tests using harassment-related lawsuit penalties as the outcome variable from the Good Jobs First organization's Violation Tracker database. In placebo tests where the outcome variables are lawsuits related to corporate misconduct in environmental, competition, consumer, or financial matters, I find no significant effect, suggesting that the workplace harassment measure does not capture economic misconduct.

[7]For example, Balasubramanian et al. (2021) document that NDAs are more prevalent in the tech industry, and employees with NDAs earn about 6% more, alluding to the role of valuable business information.

to silence employees about the instances of harassment in order to protect their reputations (Joyce, 2022). These reforms create a reasonably exogenous variation in the workplace harassment environment, because after the passage of these rules, employment contracts (i.e., both future and past) cannot contain provisions to conceal workplace harassment, and all the past confidentiality agreements on the issue of workplace harassment in these states become invalid.[8] Given that firms now have limited ability to conceal workplace harassment, these changes would lead to increased reputational costs for firms, especially those with systemic harassment problems. I provide evidence that workplace harassment declines significantly in firms with previously high harassment scores, and that these firms are significantly less likely to face harassment-related lawsuits following the regulations.

Exploiting this exogenous decline, I test the effect of a reduction in workplace harassment, as facilitated by bans on the use of NDAs in harassment cases, on firms' innovation output. Innovation output refers to the economic value of patents, constructed using the stock market reaction to the announcement of a given patent grant (Kogan et al., 2017). Using a triple-difference research design, where identifying variation comes from comparing changes in innovation output of firms with higher workplace harassment scores to changes in innovation output of firms with lower workplace harassment scores in treated states and that of all firms in never-treated states, I find a significant negative causal relation between workplace harassment and innovation performance (t-stat=8.43). One standard deviation decrease in workplace harassment leads to 10% increase in innovation output. These effects are economically significant and similar in magnitude to the reduction in innovation output due to negative wealth shocks to inventor productivity documented in Bernstein et al. (2021). In addition to the economic value of patents, I find that both the total number of patents and the number of scientifically valuable (i.e., highly cited) patents increase strongly in response to reductions in workplace harassment. These improvements are magnified in the third and subsequent years after the reduction in workplace harassment.

One can argue that never-treated states are not a good comparison group because they may have systemically different socioeconomic characteristics (e.g., political or macroeconomic

[8] In the baseline analysis, the state of the firm is the state in which the firm is headquartered. I find similar evidence when I define state to be the state of employees, which I obtain from Glassdoor reviews.

trends) than the treated states, meaning that being treated is not quasi-random. To address this concern, I restrict the analysis to those states that passed the regulations restricting the use of confidentiality agreements in harassment cases. In this setting, the identifying variation comes from the timing difference of the regulations, which is unlikely to be correlated with firms' innovation performance. The results from this estimation are remarkably similar to the baseline, suggesting that the never-treated states in my analysis represent a suitable control group.

Parallel trends is a key assumption underlying my analysis. It is possible that firms with higher harassment scores in the treated states have different trends in workplace harassment or innovation output than firms in the control group. I address this concern in two ways. First, I estimate a fully dynamic version of the baseline analysis and directly test for possible pre-trends. Second, given that the treatment is staggered, I test for the existence of pre-trends by estimating a fully dynamic version of the estimator obtained in the stacked difference-in-difference setting (Cengiz et al., 2019). Reassuringly, these tests show no violation of the parallel trends assumption.

I then supplement the main sample with the characteristics of the inventors responsible for these successful patents, aggregated to the firm level, collected from PatentsView. The results show that the impact of workplace harassment on innovation output is more pronounced in firms whose inventor teams include minorities in terms of gender, ethnicity, and race. Specifically, the observed effects are particularly pronounced in firms whose inventor teams include female minorities as opposed to all-male teams. Similarly, the effects are more pronounced when inventor teams are composed of individuals from minority ethnic (racial) backgrounds as compared to teams with members belonging to a single majority ethnic (racial) identity (i.e., European (white)).[9]

Finally, I examine the economic mechanism of the effects of workplace harassment on innovation. Exposing teams that develop patents or support these activities to workplace harassment can lead to productivity losses by damaging creative synergy and team-specific capital (Jaravel et al., 2018). The damage to team capital can occur either because team members may

[9]Note that, on average, inventor teams are predominantly male (90%), white (75%), and of European ethnicity (82%).

experience psychological stress due to workplace harassment (e.g., Okechukwu et al., 2014; Oswald et al., 2015) or they, particularly the more skilled ones, may opt to leave. In addition, the effects may be exacerbated if companies with a reputation for workplace harassment experience difficulty attracting skilled employees or replacing lost talent, which is key to sustaining the collaborative dynamics that drive successful innovation projects.

Using the inventor-level data, I analyze inventor productivity in the baseline research design, where I also include inventor fixed effects and compare the changes in patenting productivity of inventors in firms with previously high levels of workplace harassment to those of other inventors in the control group. I find that inventor productivity in delivering successful patents increases by about 6 to 7% in response to a one standard deviation decrease in workplace harassment. The economic magnitudes underscore that the firm-level effects are largely driven by productivity gains.

To test whether firms with higher workplace harassment experience higher turnover, I use online employee resumes from LinkedIn data. It includes employees' work experience history, including the start and end dates of their jobs at different workplaces, as well as their educational backgrounds, such as college entrance and graduation dates and degrees earned, complementing recent work using such data (e.g., Jeffers, 2023; Fedyk and Hodson, 2022).

To construct firm-level measures of talent turnover, I use educational attainment and industry experience as proxies for skills (Custódio et al., 2013; Tambe et al., 2020; Niessen-Ruenzi and Zimmerer, 2021). I then directly test the role of changes in the collaboration dynamics by estimating the baseline triple-difference specification with the proxies for net talent outflow rates as the dependent variables. I find that following the bans on the misuse of NDAs, firms with previously higher rates of workplace harassment experience a substantial reduction in net talent outflow rates. In particular, the net outflow rates of employees with advanced degrees and a Forbes Top 100 college degree decrease by about 19% and 22%, respectively, in response to a one standard deviation decrease in workplace harassment. This effect is even more pronounced for employees with high industry experience, with an approximately 1.9-fold decrease, which may indicate not only expertise but also strong job transferability within the industry.

Having documented significant improvements in workplace climate and firm innovation, I next ask whether firms also perform better. I find that performance in terms of sales and profit growth increases significantly following restrictions on the use of NDAs to conceal workplace harassment. Specifically, one standard deviation decrease in workplace harassment leads to 5.3% increase in sales per employee and 2.2% increase in firm profit growth.

This paper is the first to document an economically significant negative impact of workplace harassment on innovation. Beyond fostering a better workplace climate, regulatory changes aimed at addressing workplace harassment have far-reaching implications, ultimately enhancing firm performance by enriching the firm's talent capital. In addition, this study underscores the importance of social platforms in shedding light on aspects of firms that are otherwise difficult to observe.

Links to literature This study contributes to several strands of literature. First, my work contributes to the longstanding literature that studies innovation (e.g., Kerr and Lincoln, 2010; Moser et al., 2014; Manso, 2011; Aghion et al., 2013; Bernstein, 2015). Specifically, I contribute to the literature that provides a "bottom-up" explanation of firms' innovation performance by studying employees and inventors. For example, Bernstein et al. (2021) examines the impact of negative wealth shocks on inventors' productivity, and Acharya et al. (2014) emphasizes the role of employment protection regulations in promoting innovation. I complement this literature by emphasizing the role of collaboration dynamics and team capital, especially given the firm-centric nature of innovation in the current era (Babina et al., 2023).

Second, this paper relates to the literature on culture and firm performance (e.g., Edmans, 2011; Guiso et al., 2015; Li et al., 2021; Graham et al., 2022). I contribute to this literature in two ways. First, I introduce a novel measure of workplace harassment that complements existing measures of corporate culture, which are often derived from managers' speeches, interviews, or company websites. Second, by demonstrating the impact of workplace harassment on innovation, the results of this paper echo the facts documented by Graham et al. (2022), which show that managers strongly believe that improving corporate culture can increase firm value.

Third, with respect to the effect of workplace harassment on firm performance, my work is related to the contemporary work of Hacamo (2023), who shows that racial harassment in firms leads to consumer boycotts. In addition, Au et al. (2023) documents a negative relationship between sexual harassment and stock performance. My paper complements this work by providing causal evidence on the costs of general workplace harassment on firms' innovation output, and by examining the underpinnings of these effects.

Fourth, my work is related to work on the labor market effects of workplace harassment. Folke and Rickne (2022) examines how sexual harassment plays a role in widening the gender wage gap. Adams-Prassl et al. (2023) examines the consequences of workplace violence against women and its impact on employment outcomes. They show a significant decline in employment for both victims and perpetrators and in the proportion of women in the firm. I complement this literature by showing the costs to worker productivity, the role of skills in determining job transitions, and the implications for team capital.

Fifth, this paper contributes to the literature on measuring unobservable firm characteristics using unstructured textual data (e.g., Hanley and Hoberg, 2019; Chen et al., 2019; Dim, 2020; Goldstein et al., 2021; Li et al., 2021; Sautner et al., 2023). My paper builds on the similar natural language processing technique of Li et al. (2021), who measure corporate culture using earnings call transcripts, and joins Green et al. (2019), Dube and Zhu (2021) by showing the role of rank-and-file employee reviews in understanding firm performance.

Finally, my paper is closely related to Sockin et al. (2022) in terms of shocks to the workplace harassment environment, and contrasts their findings. Their analysis focuses on changes in negative reviews on Glassdoor by comparing industries by their broad NDA use rate (i.e., not specific to harassment), and documents more negative reviews in these industries after harassment-specific NDA bans, while my analysis examines changes in workplace harassment by comparing firms by their pre-ban hidden workplace harassment rate, and documents significant declines in workplace harassment after the bans.

2 Data Sources and Description of Sample Characteristics

In this section, I provide institutional information for each database I use to produce the main variables of interest, summarize the distributional nature of these variables, and give a detailed definition of each variable in Table A1. I discuss the construction of the *Workplace Harassment Score* separately in Section 3.

2.1 Data on Employee Reviews

Measuring workplace harassment is not straightforward, which is a common empirical challenge in quantifying the various dimensions of corporate culture (Grennan and Li, 2022). In particular, workplace harassment goes significantly underreported, and hidden within firms (Dahl and Knepper, 2021). While one approach could be to rely on news articles or lawsuits, these typically cover high-profile incidents or arrive with a large time lag. More importantly, they primarily capture the resolution of past incidents. In this paper, I overcome these limitations by developing a measure of workplace harassment for U.S. public companies using employee reviews from Glassdoor, an online job review platform with approximately 55 million monthly unique visitors. Launched in 2007, Glassdoor began offering reviews in 2008 and operates under a ”give to get” policy, requiring users to post a review in order to access others. The platform monitors reviews for integrity and removes fraudulent submissions.

A key aspect of these employee reviews is that they are published anonymously, without any identifying information about the reviewers.[10] This gives employees the freedom to discuss not only the positive aspects of their jobs, but also the negative ones. Particularly in the context of harassment, such a feature is useful to increase reporting (Boudreau et al., 2023), as victims of harassment tend to underreport due to the risk of retaliation (Dahl and Knepper, 2021).[11]. Indeed, recent literature has used various dimensions of employee reviews to document new insights into firm performance and governance; studies such as Green et al. (2019) and Sheng (2021) highlight the predictive power of employee expectations for stock returns, while Dube and Zhu (2021) shows its role in influencing company policy through employee feedback on company amenities.

[10] Glassdoor states: “Nobody, including your employer, will be able to see your email address, social media profiles, or any other personal information you provide — either in your user profile or your resume.” See here: https://help.glassdoor.com/s/article/Protecting-member-anonymity?language=en_US

[11] For example, Hacamo (2023) studies racial harassment using employee reviews from indeed.com, a platform similar to Glassdoor.

As shown in Panel A in Figure A3, employees rate companies on several dimensions, including overall satisfaction, compensation, benefits, and overall company culture. In addition to generic ratings, employees provide textual reviews about specific aspects of their jobs in three different sections: ”Pros” (discussing positive aspects), ”Cons” (discussing negative aspects), and ”Advice to Management” (providing feedback to management). Of these, the ”Cons” section provides a natural avenue for exploring employees’ experiences of harassment, which I will later use in the textual analysis to construct the workplace harassment measure. For example, in the second review in Panel B in Figure A3, the reviewer shares an experience of sexual harassment and a bullying culture, while the first reviewer shares an experience of slow promotions and salary increases.

To obtain employee reviews, I first collect firm names from the Compustat database via Wharton Research Data Services (WRDS), focusing on the period from 2009 to 2022. To obtain the sample of publicly traded U.S. firms, I focus on those with share codes 10 and 11 (i.e., U.S. common stock) in CRSP and a currency code of USD as indicated in the Compustat database. This results in 7,485 unique firm names.

I then matched these company names to those from the Glassdoor website. At this stage, I first collected the Glassdoor company name suggestions with the corresponding Compustat company names. I then manually verified these matches using additional information such as headquarters location, stock tickers, website links, and using additional information on the web when this information was not available. This comprehensive process resulted in 4,490 matched company names.

In the final stage, I collected employee reviews for these firms, requiring a minimum of five reviews per year for a given firm to ensure data quality. Overall, the final sample covers the years 2009 to 2022, with a total of 3.7 million reviews obtained for 3,288 unique firm names, averaging 1,533 unique firms per year. Table 1 provides summary statistics on the distribution of the number of reviews, showing an average of 169.9 reviews per firm per year, with a median of 39 reviews per firm per year. Scaled per 1,000 employees, this equates to an annual rate of approximately 18.7 reviews per 1000 employee.

2.2 Data on Harassment-Related Lawsuits

Workplace harassment is a violation of federal civil rights laws, specifically Title VII of the Civil Rights Act of 1964, the Age Discrimination in Employment Act of 1967 (ADEA), and the Americans with Disabilities Act of 1990 (ADA), all of which are enforced by the U.S. Equal Employment Opportunities Commission. Although such lawsuits are relatively rare, they serve as a valuable natural setting for assessing the external validity of the workplace harassment measure. Specifically, to examine whether the measure is predictive of future events and thus captures non-public and unreported harassment. Textual analysis of these lawsuits by Dougal et al. (2022) shows that 88% of them contain keywords such as ”retaliation” or ”harassment. Thus, in this paper, I refer to these lawsuits as ”harassment-related lawsuits” and draw data on them from two datasets: (i) the Federal Judicial Center’s (FJC) Integrated Database, accessible through the WRDS, for the baseline analysis, and (ii) the Violation Tracker Database, maintained by the Good Jobs First (GJF) organization, for robustness.

FJC Data The Federal Judicial Center provides comprehensive data on active and closed cases in the federal judicial system from the 1970s to the present. These cases are broadly categorized into four groups: civil, criminal, bankruptcy, and appeals. In my research, I focus on civil cases because harassment claims typically fall into this category. By utilizing the civil cases dataset, I can identify lawsuits that are specifically related to workplace harassment. Specifically, I extract data where the type of lawsuit is coded as ”442” (Civil Rights) and ”445” (Civil Rights ADA Employment), which includes all lawsuits filed for civil rights violations in the workplace. To link this data to my main sample, I use a linking table developed by WRDS to obtain firm identifiers such as CIK/GVKEY. By merging these data, I construct a dummy variable - “Harassment Lawsuit” - that indicates whether a firm (denoted by “i”) faces an harassment-related lawsuit in a given year (denoted by “t”). Panel B of Table 1 shows that about 4% of the firms in my sample experience at least one such lawsuit during the period 2009-2019.[12]

GJF Data Good Jobs First has developed the Violation Tracker, a comprehensive database focusing on corporate misconduct in the United States. This database tracks corporate violations of regulations and laws in a variety of areas, such as bribery and health and safety, since 2000.

[12] The linking table provided by WRDS for obtaining firm identifiers is accessible until September 2019.

A key advantage of using this database is that the Violation Tracker provides information on the fines imposed on companies as a result of their violations. Analogous to the FJS data, I use this data to identify litigation involving civil rights violations in employment. One caveat to these data is that the date stamp on the cases is given only as the date of the enforcement action, not the date of the lawsuit filing. I lead the variables by one year in the estimates to mitigate concerns about "look-back" bias. Overall, this dataset complements the FJC data with the availability of litigation penalties. Panel B of Table 1 shows that firms are fined an average of $305 for employee civil rights violations over the 2013-2022 period.[13]

	Mean	SD	5%	25%	50%	75%	95%
	Panel A: Summary of Employee Reviews from Glassdoor						
Number of Reviews	169.893	782.692	7.000	15.000	39.000	113.000	591.000
Glassdoor Participation (per 1000 Emp.)	18.764	23.696	1.271	4.810	10.873	22.903	65.217
Workplace Harassment	0.222	0.025	0.185	0.206	0.220	0.236	0.265
	Panel B: Harassment Lawsuits from FJS and GJS						
Harassment Lawsuit	0.039	0.194	0.000	0.000	0.000	0.000	0.000
Litigation Penalty	305.191	3828.007	0.000	0.000	0.000	0.000	0.000
	Panel C: Innovation Output from Kogan et al. (2017)						
Economic Value of Patents per Employee	22.115	89.465	0.000	0.000	0.000	0.653	111.492
Number of Patents per Employee	2.014	8.802	0.000	0.000	0.000	0.042	9.771
Number of Highly Cited Patents per Employee	0.822	4.237	0.000	0.000	0.000	0.000	3.486
	Panel D: Workforce Loss Rate (WLR) from LinkedIn						
WLR with or Above Masters Degree	-0.022	0.070	-0.149	-0.065	-0.018	0.020	0.093
WLR with Below Masters Degree	-0.063	0.135	-0.284	-0.154	-0.068	0.024	0.171
WLR with Forbes College degree	-0.012	0.038	-0.086	-0.031	-0.007	0.011	0.048
WLR with Forbes College degree	-0.074	0.164	-0.337	-0.188	-0.080	0.030	0.214
WLR with High Industry Experience	0.021	0.151	-0.228	-0.083	0.017	0.129	0.273
WLR with Low Industry Experience	-0.146	0.153	-0.429	-0.243	-0.128	-0.034	0.079
	Panel E: Summary of Firm Characteristics from Compustat						
Log(Sales Per Employee)	5.809	0.971	4.347	5.329	5.769	6.333	7.350
Profit Growth	0.097	0.507	-0.352	-0.030	0.065	0.187	0.637
Number of Employees (in thousands)	19.150	71.003	0.313	1.404	4.500	13.500	75.000
Cash/Assets	0.173	0.188	0.007	0.037	0.102	0.239	0.597
Debt/Assets	0.294	0.234	0.000	0.100	0.268	0.425	0.731
PP&E/Assets	0.213	0.221	0.009	0.051	0.127	0.304	0.721
CAPEX/Assets	0.035	0.038	0.001	0.010	0.023	0.046	0.110
R&D/Assets	0.040	0.082	0.000	0.000	0.000	0.042	0.201
Market Share (Sales)	0.005	0.012	0.000	0.000	0.001	0.003	0.027
RoA	0.037	0.156	-0.226	0.012	0.062	0.108	0.201
Tobin's Q	1.573	1.839	0.103	0.514	1.001	1.904	5.175

Table 1: Summary Statistics of Main Variables. This table summarizes the main variables used in this study. The sample is restricted to the period 2013-2022 (same as for the baseline analysis), firms with at least five reviews per year, and firms with non-missing observations on the Compustat firm characteristics listed above. Table A1 provides a detailed description of the variables used in this study.

[13]Note that there are approximately 294 firms in my sample that are penalized. Depending on the size of the firm, employee civil rights litigation penalties can range from $50,000 to $300,000.

2.3 Data on Innovation and Inventor Teams

To measure innovation output at the firm level, I use data on firm patent activity from the recently updated innovation output dataset provided by Kogan et al. (2017). The authors obtain patent data from the Google Patent database and link them to firm identifiers using the names of patent assignees. Using a patent-specific productivity measure has several advantages in this context: (i) it inherently relates to human capital productivity due to human creativity and idea generation, as well as the collaborative efforts across and within teams and the human expertise required for continuous inventive work on existing patents, (ii) patenting activity is a common practice in various industries, including apparel, healthcare, technology, and more, making it suitable for cross- and within-firm analysis, (iii) patenting serves as a strong predictor of a firm's future success, as successful innovations often translate into new products and improved operational efficiencies for the firm.

Kogan et al. (2017) compute the economic value of patents based on stock price movements around the patent issuance date (within a 3-day announcement window). Unlike previous literature that relies on patent citations to measure the value of patents to the firm, Kogan et al. (2017) argue that economic value better captures firm-specific prospects and potential for creative destruction, and is a more robust predictor of firm and output growth. Thus, the comparability of this measure across firms and the ease of quantifying its economic importance to firms make it an ideal proxy. Moreover, the problems of truncation bias, which are severe with citation-based metrics because the most recent patents do not accumulate enough citations, are minimal with such a proxy.[14]

In my analysis I consider patent value as of the patent filing date, as suggested in the recent work by Kempf and Spalt (2023). This approach accounts for the time lag between filing and issuance (2.9 years on average) and mitigates concerns about changes in firm performance due to patent issuance success. Analyzing the impact of harassment around the filing period makes sense because these patents already exist at that point.

[14]In the main analysis, when I perform a robustness test to exclude the COVID period, the problem of the inability to observe the most recently filed patents is also automatically addressed, and the results remain qualitatively similar.

For validity and further robustness, I also include the number of patents and citation-weighted number of patents as outcome variables in my analysis. Indeed, Kogan et al. (2017) shows a positive linear relationship between economic patent value and citation-weighted number of patents. To ensure comparability, I normalize all three variables (total economic patent value, number of patents, number of citation-weighted patents) by the number of employees because workplace harassment measures are averaged across reviews. And the number of reviews is strongly correlated with the number of employees. Panel C in Table 1 provides summary statistics on the measures of innovation output. For example, the value of patents per employee in 1982 dollars in my sample is about $22 thousand, firms have 2 patents per 1000 employees, and 0.8 highly cited patents per 1000 employees.

In addition, I supplement this data with the inventor data from PatentsView. Specifically, I use the gender and last name information of the inventors. The surnames are useful to proxy the ethnic and racial background of the inventors. Using the *ethnicolr* package, I group the inventors into three ethnic groups (African, Asian, and European) and four racial groups (Asian, Black, Hispanic, and White). On average, inventor teams are majority male (90%), European (82%), and white (75%). Next, I calculate the racial and ethnic diversity of the inventor teams using Simpson's Diversity Index method.

2.4 Data on Skilled Employee Flows

To be able to study the role of skilled employee turnover in explaining the relationship between workplace harassment and innovation, I obtain the LinkedIn data from BrightData Initiative, a startup company. LinkedIn is a professional networking plaform, launched in 2003, and has more than 900 (200) million users globally (in the US). The unique feature of this data is that it allows me to observe (i) educational history, including start and graduation date, name of university, type of degree program (ii) employment history, including firm names, job start and end date, position titles, and locations. More importantly, LinkedIn has standardized all this information making it possible to construct measures across firm and educational attainment.

To merge the LinkedIn data, I initially collect firm identifiers from public pages of companies on LinkedIn to match it to the sample provided by the BrightData. For the sample of firms with Glassdoor review data, I initially search company names and collect all the relevant company info from the LinkedIn and next manually check if the names found is matching Compustat and Glassdoor company identifiers.

Using the data, I construct measures of workforce loss rates, focusing on three key skill characteristics of employees, as outlined in previous research (Neal, 1995; Custódio et al., 2013; Niessen-Ruenzi and Zimmerer, 2021). These skill characteristics include:

(i) *Flow of Employees with (without) Advanced Degree*: This measure captures the difference between the number of employees entering and exiting a firm who possess a master's degree or higher qualification (such as PhD, MD, etc.).

(ii) *Net Flow of Employees with (without) a Degree from the Top 100 Forbes Colleges*: This measure captutes the net movement of employees with a degree from the top 100 Forbes-ranked colleges.

(iii) *Net Flows of Employees with Industry Experience*: This measure identifies employees who have above-median experience compared to all employees in a given year (t), either upon exit or entry into a firm.

For each firm (i) in a given year (t), I calculate the workforce loss rate for employees with these specific skill characteristics using the following formula:

$$Workforce\ Loss\ Rate(type)_{i,t} = \frac{Outflow\ (type)_{i,t} - Inflow\ (type)_{i,t}}{Total\ outflow_{i,t} + Total\ inflow_{i,t}} \tag{1}$$

The numerator in this equation captures the net flow of employees with a specific attribute, while the denominator normalizes the variable by the overall size of employee flow. For instance, when calculating the workforce loss rate for employees with a master's degree or above, the numerator represents the difference between the number of such employees exiting and entering the firm. This allows us to capture both the exit and entry patterns of employees with the specified skill attribute. I provide further details on the identifying the skilled employees in Appendix B.3.

2.5 Supplemental Firm-level Datasets

Other data sources I use are Compustat for company financial information. In addition, I collect the Best Companies To Work For In America list published by the Great Place To Work Institute (Edmans, 2011) from Alex Edmans' website,[15] ESG scores from Refinitiv, and firm culture measures from Li et al. (2021).

3 Firm-Level Workplace Harassment Measure

This section discusses the construction and validation of workplace harassment measure.

3.1 Constructing Workplace Harassment Measure

As I discuss in section 2.1, I use the textual content of anonymous employee reviews from Glassdoor to construct the workplace harassment measure. Specifically, I use the textual content of employee reviews from the "Cons" section of Glassdoor, where employees share negative aspects of their jobs. This section serves as a natural way to examine whether employees discuss instances of workplace harassment. In addition, omitting other textual sections (Pros or Advice to Management) helps mitigate potential mismeasurement issues such as false positives.

To quantify workplace harassment using the textual content of employee reviews, I employ a machine learning methodology known as *word2vec* models, developed by Mikolov et al. (2013b,a) at Google. This approach was recently introduced to the finance literature by Li et al. (2021) in their study of corporate culture using earnings call transcripts. They provide valuable insights into how this method can be used to generate data-driven dictionaries and compute firm-level scores. In my paper, I extend this approach to the context of employee reviews, which are more frequently available than earnings calls. In the following sections, I provide a detailed description of the technique, define the concept of workplace harassment that I use to construct the measure, and discuss the measurement process in detail.

[15] Thanks to Sai Zhang, a pre-doctoral researcher funded by the LBS AQR Asset Management Institute, for updating this dataset to 2020.

3.1.1 Word2vec Models

Previous approaches to understanding the similarity of large numbers of text documents have relied primarily on word frequency, which measures the occurrence of words in two documents without regard to their contextual meaning (Goldstein et al., 2021). While this approach has been successful in certain applications, particularly when the documents have a consistent structure (e.g., 10-K filings), it can lead to significant measurement errors when dealing with texts that contain multiple synonyms or are written by a diverse group of people with different writing styles and potential spelling mistakes. To analyze such text, we need to use methods that more efficiently measure the semantics of the text.

One such model, developed by Mikolov et al. (2013b,a), has been successful in incorporating contextual information into text analysis. The main result of this approach is the generation of word vectors, which efficiently capture the meaning of words and enable the identification of analogous (synonymous) words with fewer errors. Word vectors represent the semantic properties of words by attempting to maximize the classification of a word based on its context within a sentence. The algorithm achieves this by taking into account the words before and after the target word.

To prepare the text corpus for textual analysis, several cleaning steps are necessary to ensure that the words are in a standardized format. First, non-word characters and stop words are removed from the text. Stop words are common words that have no meaning and can introduce noise into the analysis. In addition, since the models are trained on sentences to understand context, the reviews are tokenized into sentences. This allows the algorithm to learn the semantics of each word based on the surrounding words.

To capture context more accurately, word combinations (i.e., bigrams) are useful (Mikolov et al., 2013a). For example, the phrase ”verbal abuse” is more relevant in the context of harassment than the word ”verbal” alone. To identify such word combinations, I use the *Phrases* model. This algorithm identifies phrases by analyzing which words occur more often together and less often with other words. For example, it would suggest ”racial harassment” as a phrase, but not ”manager makes”. The latter would remain as unigrams (i.e. single words). For details on preprocessing textual content, see Appendix B.1.

Once the phrases are constructed, I proceed with training the model to obtain the word vectors. The model parameters, which are described in detail in Appendix B.2, are carefully selected to ensure that the word vectors are estimated optimally for the specific objectives of the analysis.

3.1.2 The Definition of Workplace Harassment, Seed Words and Dictionary

Having the model already trained on the employee review text corpus, next goal is to build a dictionary effectively capturing the workplace harassment, which is a key for calculating the firm-level score. I follow the definition given by U.S. Equal Employment Opportunities Commission (EEOC) to obtain seed words and to construct the dictionary:

> Harassment is unwelcome conduct that is based on race, color, religion, sex (including sexual orientation, gender identity, or pregnancy), national origin, older age (beginning at age 40), disability, or genetic information (including family medical history).

Seed words serve as reference points for the model to identify other words that are most similar in meaning, allowing us to understand the language used to describe harassment in the context of employee reviews. Given the broad definition of harassment, which includes various forms of harassment such as race, age, and disability harassment, I create a list of seed words that cover all of these aspects. I provide the seed words in the top panel of the table 2. The list includes variations of the same words with different syntax to capture expressions of harassment in different ways (verb, noun, etc.).

Following similar steps as in Li et al. (2021), I obtain the 1000 most similar words for each word in the seed word list based on cosine distance, resulting in a total of 27,000 words. To ensure uniqueness, I drop duplicated words, resulting in 9,000 unique words. For the final dictionary, I keep the most similar 500 words [16]. Such selection ensures the depth of language used in the context of harassment. Having a larger pool of 500 words, as opposed to a smaller set of 20 or 30 words, minimizes the potential for misclassification.

[16]To further validate the relevance of these 500 phrases, I also ask ChatGPT if they are directly related to workplace harassment, and it confirms that all the words could be used in the context of workplace harassment.

Seed Words									
harass, harassed, harassing, harasses, harassment, harasser, mobbing, bullying, bully, bullied, racism, racist, colorism, sexual, sex, sexism, sexist, transphobia, homophobia, lgbt, lgbtq, phobia, religion, religious, xenophobia, ethnocentric, ageism, ageist, ableism, ableist, discriminate, discrimination, discriminated, discriminating									
Most Similar 100 Harassmet Words									
1	racism	21	harassment_bullying	41	harrassing	61	belittled	81	ageism
2	discrimination	22	homophobic	42	religious_belief	62	discrimination_harassment	82	nepotism
3	sexual_harassment	23	racism_discrimination	43	unethical_behavior	63	harassment_racial	83	hound
4	harrassment	24	harass	44	inappropriate_behavior	64	sexist_racist	84	abusive
5	harassed	25	harassing	45	gaslighting	65	mysogynistic	85	mistreat
6	bullied	26	bias	46	misconduct	66	constantly_harassed	86	stalking
7	sexism	27	mistreatment	47	verbally_abused	67	discriminating	87	sexism_sexual
8	bullying	28	prejudice	48	intimidation	68	sexual_harrassment	88	ageism_racism
9	harassment	29	bullying_harassment	49	racist_sexist	69	mistreated	89	stalked
10	racial_discrimination	30	inappropriate_sexual	50	bullying_behavior	70	inappropriate	90	racism_homophobia
11	harrassed	31	disrespected	51	sexual_assault	71	harras	91	hostility
12	racist	32	sexually_harassed	52	christian	72	belittle	92	harassment_verbal
13	sexist	33	verbal_abuse	53	homophobia	73	male	93	belittling
14	misogynistic	34	discriminatory	54	unprofessional_behavior	74	discriminated	94	hispanic
15	age_discrimination	35	harassment_discrimination	55	abusive_behavior	75	homophobic_comment	95	chauvinistic
16	misogyny	36	harrasment	56	treated_unfairly	76	favoritism	96	verbal_sexual
17	lgbtq	37	harrass	57	racism_sexism	77	treated_poorly	97	covert_racism
18	lgbt	38	bulling	58	sexism_racism	78	demeaned	98	verbal_harassment
19	workplace_bullying	39	discriminates	59	bigotry	79	passive_aggression	99	racial_slur
20	retaliation	40	bully	60	disrespectful	80	unfair_treatment	100	pedophilic

Table 2: Seed Words and Most Similar (i.e., to seed words) 100 Words in Workplace Harassment Dictionary.

Table 2 shows the 100 most similar words from the workplace harassment dictionary. As we can see from the list, the model is remarkably effective at capturing the context of workplace harassment, and it indicates how much language varies across reviews and how pervasive the problem is.

3.1.3 Aggregating the Workplace Harassment Measure

In order to score the level of workplace harassment at the firm-year level, I follow the steps outlined below, after having constructed a dictionary of words described above:

(i) Firstly, I calculate the average cosine similarity of each word in an employee's review with the harassment word dictionary. Each word is represented by a vector of size 100x1, and the harassment word dictionary consists of 500 words. The calculation is performed using the following formula:

$$Word\ Similarity = \frac{\sum_{1}^{500} Cosine\ Similarity(Word\ Vector^{T}_{100\times 1}, Harass\ Matrix_{100\times 500})}{500}$$

where the cosine similarity between two vectors, a and b, is defined as:

Cosine Similarity$(a, b) = \frac{a \cdot b}{|a| \cdot |b|}$

(ii) Next, I repeat step (i) for all the words in a given employee review. This allows me to obtain the similarity of each word in the review with the harassment dictionary. I then take the maximum similarity value across all the words. The objective here is to capture whether the employee is referring to workplace harassment in the review by identifying the word that has the closest distance to the harassment dictionary. For a review with j words, this can be expressed as:

$$Harass\ Review = max(word\ similarity_1, ..., word\ similarity_j)$$

(iii) Finally, to calculate the workplace harassment score for firm i in year t, I take the average of the previously calculated $Harass\ Review$ values for all employee reviews posted in year t.

This is done using the following formula:

$$Workplace\ Harassment_{i,t} = \frac{\sum_{1}^{N} Harass\ Review_{n,t}}{N}$$

3.2 Validating Workplace Harassment Measure

I validate the workplace harassment measure in several ways. First, after constructing the measure at the firm-year level, I examine the trend in the workplace harassment score and contrast it with the national-level workplace harassment charges shown in Figure 1. In panel A, I show how the cross-sectional distribution of workplace harassment scores evolves over time. This figure shows that there is a large amount of cross-sectional heterogeneity in the workplace harassment score and that it changes over time. To contrast the changes, i.e. the decreasing trend, I use the trend for the national-level workplace harassment charges from the Equal Employment Opportunities Commission. Panel B also shows the similar decreasing trend in workplace harassment-related charges starting in 2010, which gives me confidence that my measure accurately tracks the macro-level workplace harassment environment. In addition, Panel B shows that harassment related to employees' race and gender identities together account for about 60% of all charges.

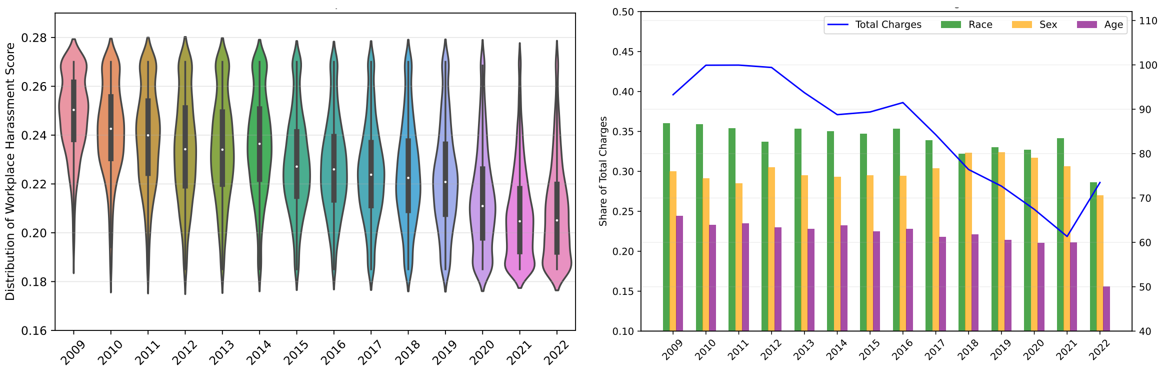

A: Workplace Harassment Score

B: National Workplace Harassment Charges

Figure 1: Validating the Workplace Harassment Measure: Trend in Workplace Harassment Score vs. Total Workplace Harassment Charges. Panel A shows the cross-sectional change in the workplace harassment score constructed using anonymous employee ratings for 2009-2022. Panel B shows the national trend in workplace harassment-related charges for the same period. Workplace harassment charges are from the Equal Employment Opportunities Commission.

Next, I provide concrete examples of employee reviews posted on Glassdoor across companies, industries, and years. To do this, I rank all reviews by their workplace harassment score. I report the review texts with and without workplace harassment complaints in Table 3. In Panel A, I list examples of employee reviews that are in the top percentile of the workplace harassment measure. Reading all the reviews, we can see that employees complain about harassment in various dimensions, sexual, racial, age, and so on. Next, I provide examples of review texts that rank in the bottom percentile of harassment scores. Although these reviews contain complaints about various other things about the jobs, there is no discussion of workplace bullying or harassment. These examples assure us that the workplace harassment measure is high when there is actually some discussion of the topic.

While I have provided examples from employee reviews that convincingly demonstrate the measure's ability to capture instances of abuse and harassment, I next examine the credibility of these harassment complaints on Glassdoor. To test the validity of the workplace harassment measure, I conduct an analysis to test whether the measure can predict workplace harassment lawsuits. Specifically, I estimate the following model:

$$Harassment\ Lawsuit_{i,t+1} = \alpha + \beta\ Workplace\ Harassment_{i,t} + \Gamma X_{i,t} + \gamma_{j,t} + \epsilon_{i,t+1} \quad (2)$$

i, j, and t index the firm, industry, and year, respectively. The variable $Harassment\ Lawsuit_{i,t+1}$ is an indicator variable that takes the value 1 if a harassment lawsuit is filed against firm i in year $t+1$ and otherwise.[17]

[17] $Harassment\ Lawsuit_{i,t+1}$ also equals one if firms face more than one harassment lawsuit. Alternatively, I also run Poisson regressions where $Harassment\ Lawsuit_{i,t+1}$ is a count variable. These results are reported at A3 and confirm the baseline results.

Panel A:			Review Examples from Top 1 percentile
	Review Date	**Company Name**	**Review (Cons)Text**
1	1/20/2022	Oracle	Bully culture, Sexual harassment, ignorance of federal and state labor laws.
2	6/16/2021	Plexus	Only to certain Department in APAC; Poor Management; Dictatorship culture; Prejudice & Bias culture; High turnover rate
3	12/25/2018	Rocky Mountain Chocolate	pay is low for amount of dedication you are supposed to put in, people working there for years promised promotions but promises not fulfilled; owners not the friendliest and can be rude and demanding towards staff ; poor communication with employees, no direct communication from management to employees; schedules only released on weekly basis; sexual harrassment unaddressed when reported to upper management; promotion and benefits based off of favouritism despite lack of competence (having prior relationship to upper management); supplies often not meeting customers demands, such as certain packaging advertised but not provided to store; can ask for days off and have them approved but still be scheduled for those days; when asking for leave of absence have to ask an inappropriate amount of time before (applies to part-timers as well)
4	8/28/2021	Microsoft	Lots of misogyny, racism, and corrupt politics.
5	12/19/2017	Tech Data	Very discriminatory culture when it comes to age and gender. Once you get to a certain number of years you get laid off. Just like Menudo where they get kicked out of the band after hitting a certain age. The faces in the organization are unrecognizable by previous generations laid off; always cleaning house to bring in new people at starting salaries. There are few people there in higher management levels that perpetuate this trend.
6	11/25/2018	SkyWest Airlines	Bullying and harassment is rife from some areas of management; Long hours; Lots of conflict due to inconsistent management
7	8/27/2014	American Woodmark	Management does not care about HARASSMENT, HR is a JOKE
8	9/25/2012	Archer Daniels Midland (ADM)	there is no communication, people falsifying documents, harassment, company rules apply to certain people and not all.
9	10/12/2020	Aaron's	Illegal ethnic discrimination, passive aggressive coworkers who set each other up, managers are only concerned with profits and threaten store managers with termination, no positive management
10	7/12/2015	Meridian Bioscience	Sexual harassment overlooked. No advancement unless you are part of the good ole boys golf club. Upper management (operations) intimidation. Hostile and racist work environment. Document falsification.
11	9/6/2017	Brunswick	This is a very blatantly ageist, racist company. The management in most departments are filled with under qualified, cocky millennials that have been granted a chance because of either who they know or who they have sucked up to.
12	2/1/2016	Unisys	No communication, poor management, no room for advancement unless you are a man and even then not much. No mentoring, no reviews, no HR recourse, rampant sexist attitudes from high up. Too many narrow thinking people in upper management. Unisys does not value people at all and its sad.
13	2/18/2009	Cerner	Bad leaders. Some Distinguished Engineers are pathetic. Very disrespectful and no work ethics. Not a Software Company. If you are a new college graduate then Don' t do the the same mistake as we did. They make you work on a language of their own which will never count to your experience. Some teams are totally confused as to which technology they need to shift to.
14	5/9/2014	America's Car-Mart	Poor leadership, no direction, racist, irrational decisions and direction
15	8/20/2016	The Coca-Cola Company	Anti black executive racism still exist in this organization
Panel B:			Review Examples from Bottom 1 percentile
	Review Date	**Company Name**	**Review (Cons)Text**
1	1/12/2018	Tesla	They do not match 401k.
2	10/19/2020	Telenav	salary not competitive and need expand it's business
3	10/27/2021	Express	Not much room for growth
4	9/19/2017	Primerica	You generate your own leads
5	9/25/2017	Spectrum Brands	Once you're in, you're stuck forever.
6	6/13/2008	IBM	Size and bureaucracy can be stifling.
7	1/12/2010	Agilent Technologies	very process oriented; a bit slow to move; big company
8	2/12/2011	Cisco Systems	too much process, not much growth
9	3/14/2012	Walt Disney Company	Everything is in constant flux.
10	8/6/2012	Papa John's	they're not so great at anything else.
11	3/15/2013	Safeway	they have a mis-appropriated business focus
12	4/22/2015	Broadcom	I have been here too long ;)
13	11/7/2016	Pioneer Natural Resources	Oil and Gas industry is down
14	2/23/2018	Altria	The industry can be tough.
15	3/9/2021	Apple	if staying at home is not what you are looking for it might be not for you

Table 3: Validating Workplace Harassment Measure: Employee Review Examples. This table provides examples of employee reviews with and without workplace harassment complaints. Panel A (B) lists employee review texts that are among the reviews that rank in the top (bottom) 1 percentile of the workplace harassment measure.

The data on harassment lawsuits come from the Federal Judicial Center's Integrated Database, which I describe in Section 2.2. I control for several firm characteristics in the estimation. For example, firms with substantial financial resources may face a higher number of lawsuits.[18] In addition, larger companies may be more vulnerable to harassment incidents and subsequent lawsuits due to the size of their workforce. Furthermore, the measure might be affected by employee participation rates on Glassdoor. Finally, there may be differences in the likelihood of workplace harassment across industries or states. To account for these, I include several control variables in the estimation, such as Cash/Assets, Debt/Assets, PP&E/Assets, CAPEX/Assets, R&D/Assets, Market Share (Sales), RoA, Tobin's Q, Log(# of Employees), Glassdoor Participation, and industry trends (represented by $\gamma_{j,t}$). Finally, to account for cross-period correlation in the standard errors, I cluster them at the firm level.

	Harassment Lawsuit								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Workplace Harassment	0.445***	0.347***	0.465***	0.458***	0.431***	0.397***	0.427***	0.355***	0.202***
	(0.093)	(0.088)	(0.095)	(0.100)	(0.097)	(0.101)	(0.093)	(0.098)	(0.075)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	No	No	No	No	No	No	No	Yes
Industry × Year FE	No	Yes	No	No	No	No	No	No	No
State × Year FE	No	No	Yes	No	No	No	No	No	No
State × Industry × Year FE	No	No	No	Yes	No	No	No	No	No
State Inc. × Year FE	No	No	No	No	Yes	No	No	No	No
State Inc. × Industry × Year FE	No	No	No	No	No	Yes	No	No	No
State Emp. × Year FE	No	No	No	No	No	No	Yes	No	No
State Emp. × Industry × Year FE	No	No	No	No	No	No	No	Yes	No
Firm FE	No	No	No	No	No	No	No	No	Yes
Obs.	13,269	13,263	12,823	11,396	12,073	11,134	13,155	11,775	12,890
R^2	0.090	0.111	0.125	0.284	0.127	0.186	0.120	0.254	0.409

Table 4: Validating Workplace Harassment Measure: Predicting Workplace Harassment Lawsuits. This table reports the results of regressions of future (at year $t+1$) workplace harassment lawsuits on the workplace harassment measure specified in Equation 2. Controls include are Cash/Assets, Debt/Assets, PP&E/Assets, CAPEX/Assets, R&D/Assets, Market Share (Sales), RoA, Tobin's Q, Log(# of Employees), Glassdoor Participation. "State", "State Inc.", and "State Emp." are the states where the firm is headquartered, where the firm is incorporated and from where the majority of employees post reviews in year t, respectively. Industry fixed effects are Fama-French 17 industries. Variable definitions are provided at Table A1. In parentheses below the point estimates are standard errors clustered at the firm level. Full table is given at Table A2.

The estimation results are presented in Table 4. In column 1, the estimation shows that the workplace harassment measure can significantly predict future workplace harassment lawsuits after controlling for time-varying firm characteristics and common trends across firms. In column 2, I compare firms within the same industry and year, distinguishing between those with high and low harassment scores. The results show a similar pattern of predictability, suggest-

[18]For example, the Financial Times highlights that private funds are targeting climate litigation against firms with particularly deep pockets: The money behind the coming wave of climate litigation.

ing that harassment litigation is not limited to specific industries, but rather reflects substantial variation across firms. Additionally, in column 4, I adopt an even more conservative specification by comparing firms within the same industry, state, and year, as harassment lawsuits may be more industry-specific in certain locations. Again, the results are robust. In terms of economic magnitudes, a one standard deviation increase in the harassment score corresponds to about a 29% increase in the likelihood of harassment lawsuits.

Columns 5 through 8 replicate the analysis using alternative state identifiers, such as the firm's state of incorporation[19] and the employees' state.[20] These additional tests address two main concerns. First, the choice of where a firm is incorporated is strategic and would help firms avoid such lawsuits. Second, the headquarter state may not be the state where the majority of the firm's employees are located. The results of these alternative states-based analyses are consistent with those of the headquarters-based analysis, both in terms of economic magnitude and statistical significance, thus ruling out such concerns.

So far, I have exploited between-firm variation in predicting lawsuits, that is, comparing firms within the same industry, year, and state. In column 9, I introduce firm fixed effects to account for time-invariant firm characteristics and exploit within-firm variation in workplace harassment. I find that the change in the workplace harassment score significantly predicts future lawsuits, highlighting the presence of within-firm heterogeneity across time.

Predicting lawsuits a year in advance allows me to argue that the workplace harassment measure captures inside and ex-ante information about these problems. However, underreporting of harassment cases is a common phenomenon (Dahl and Knepper, 2021; Boudreau et al., 2023), the filing of a lawsuit may occur with large time lags. To investigate the prevalence of such a delay, I estimate Equation 2 over a period of up to eight years and plot the coefficient of interest (i.e., β) in Figure 2. The results show that the measure can significantly predict future lawsuits up to five years in advance, and that the predictability decreases from the fifth year onward. In addition, the figure highlights an important pattern, especially later for the iden-

[19] I obtain data on the states in which firms are headquartered and incorporated from Loughran and McDonald's Augmented 10-X Header Data, which is obtained by scraping 10-K filings. Huang et al. (2019) note that Compustat's "state" item only provides the most recent information on the company's headquarters state and does not capture historical information.

[20] On Glassdoor, employees can specify their city and state when they post the review. For company i in year t, I define the state of employees to be the state indicated in the majority of reviews.

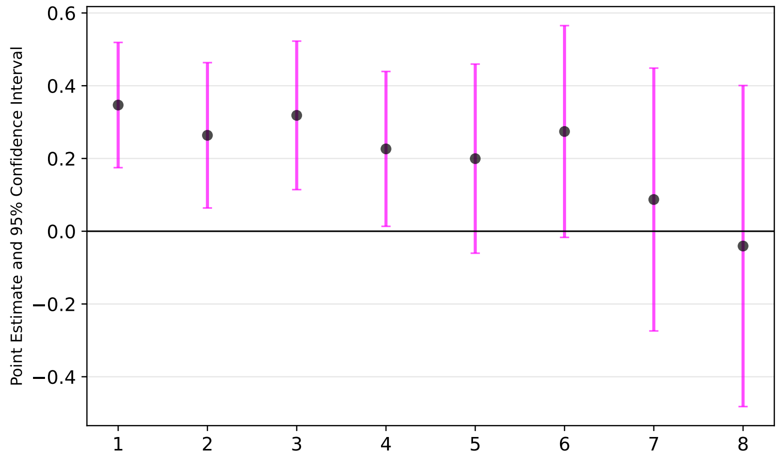

Figure 2: Validating Workplace Harassment Measure: Predicting Workplace Harassment Lawsuits Over the Next N Years. Figure depicts the point estimates and 95% confidence interval bands for the regressions of future (at year $t + n,\ n \in [1, 8]$) workplace harassment lawsuits on workplace harassment measure specified in Equation 2. Controls included are same as in Table 4. Variable definitions are provided at Table A1. Standard errors are clustered at the firm level.

tification setting. The measure of workplace harassment captures ex-ante information about workplace harassment up to five years before it enters the public domain (litigation or public news). Thus, it captures the hidden (internal) aspects of workplace harassment.

Additional Tests: I conduct additional tests to show the validity of the workplace harassment measure. First, using data on harassment lawsuit penalties from the Violation Tracker database (Good Jobs First), I show that the workplace harassment measure has significant power to predict these costs (i.e., in year $t + 1$) shown in Table A4. Second, using the data on the 100 Best Companies to Work for in America from Edmans (2011), I show that firms with higher workplace harassment scores are significantly less likely to be on this list in year $t + 1$ given in Column 1 of Table A5. Third, using data on ESG scores from Refinitiv, I show that firms with higher workplace harassment scores will have significantly lower S, G and ESG scores in year $t + 1$, as shown in Columns 2-5 of Table A5. Next, in the baseline specification given in Eq. 2, I include additional controls related to corporate culture, including the Best Workplace indicator, ESG scores, and corporate culture measures Li et al. (2021), and show in Table A6 that the workplace harassment measure has similar economic and statistical significance in the presence of these variables. That is, the measure captures harassment-specific information that is not directly embedded in these other measures. Finally, to show that the workplace harassment measure does not capture general corporate misconduct behavior, I conduct placebo tests on the other litigation, such as competition, financial, environmental, and consumer protection fraud, and find no significant prediction power, as reported in Table A7.

4 Hidden Workplace Harassment and Innovation

So far, my analysis has focused primarily on the construction and validation of the workplace harassment measure. In this section, I shift the focus to investigating the relationship between workplace harassment and innovation. To do so, I exploit the staggered adoption of regulations prohibiting the misuse of non-disclosure agreements (NDAs) in workplace harassment.

4.1 Identification Strategy

When investigating the impact of workplace harassment on innovation, several econometric challenges need to be addressed. Firstly, a simple regression approach may not suffice as it fails to adequately account for firm-specific time-varying unobservables that could lead to biased estimates. Additionally, firms with lower engagement in innovation activities may have higher levels of workplace harassment, making it difficult to disentangle the causal relationship.

Another challenge is the lack of a systematic measure of workplace harassment to precisely identify if shocks are really affecting innovation outcomes through workplace harassment. A direct way of measuring harassment would be using news articles or lawsuits, however, these do not encounter large set of firms. Moreover, these measures capture the ”ex-post” nature of harassment incidents rather than capturing the hidden harassment within organizations.

I employ an identification strategy that addresses both measurement and endogeneity concerns. First, I develop a systematic measure of workplace harassment, as discussed in Section 3, which has significant predictive power for harassment-specific lawsuits. This predictive power persists even after controlling for industry, location-specific trends, firm characteristics, and other plausibly related measures. In addition, I conduct placebo tests to ensure that the harassment measure does not predict lawsuits related to financial or other economic fraud. Overall, the measure allows me to capture cross-sectional heterogeneity for the quasi-experimental setting.

Having the measure, I exploit recent regulations that prohibit firms from entering into non-disclosure agreements with their employees specifically to conceal harassment issues within the firm and to protect firm's reputation (Joyce, 2022). As workplace harassment has gained significant attention with growing scandals (e.g., the Harvey Weinstein scandal in 2017) or movements to speak up about harassment in the workplaces (e.g, #MeToo), it has become clear that NDAs have been used by companies to hide workplace harassment. Typically, companies would include such confidentiality agreements as a provision in employment contracts to protect valuable information, such as trade secrets and patents, specifically in technology sector (Balasubramanian et al., 2021).

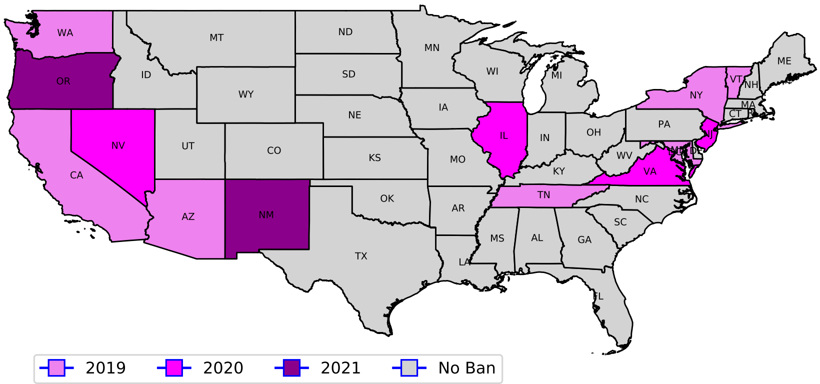

Figure 3: Staggered Passage of NDA-Limiting Regulations by States. The figure shows when and where restrictions on the use of NDAs have been enacted in the United States that prohibit their use to hide harassment in the workplace. More details about these regulations are reported in the Table A12.

NDAs are legally binding contracts, and breaching them can have significant consequences for employees. In the harassment context, the enforceability of such provisions is largely untested (Zhai, 2020), creating uncertainty for victims. This uncertainty may discourage victims from coming forward publicly because they fear the potential consequences of breaching the NDA. Proving a case of harassment can be a difficult task (Clermont and Schwab, 2009), as well as victims may fear retaliation (Dahl and Knepper, 2021), all of which discourage victims from taking legal action. Adopting policies that do not allow companies to hide harassment in the workplace increases their reputational risks, as victims can now openly discuss harassment issues.

Between 2018 and 2022, a total of 13 states passed such legislation, as shown in Figure 3. Arizona took the lead in 2018 by passing House Bill 2020, making it the first state to invalidate nondisclosure clauses that can be used to silence victims of sexual harassment. Since then, 12 other states have followed suit and enacted laws prohibiting the use of such clauses in employment contracts. In particular, California, Illinois, and New Jersey have enacted broader legislation that goes beyond sexual harassment to include other forms of harassment based on factors such as race, age, and disability. Because the measure of harassment used in this study is holistic and encompasses different forms of harassment, the changes in NDA regulation in all 13 states are leveraged.

4.2 Triple Difference Estimation and Baseline Results

To estimate the relationship between workplace harassment and firm innovation activity, I use triple difference estimation. This method allows me to ensure that the observed relationship is solely due to the harassment shocks, rather than possible trends across treated states. That is, comparing changes in firm outcomes in treated states to controls in the simple difference-in-difference setting could be driven by various institutional changes at the state level. Having firms exposed to these state-level regulatory changes adds an additional layer of robustness to the standard difference-in-difference setting to pinpoint causality (Angrist and Pischke, 2009). Specifically, such a setting benefits not only from comparing firms in treated and control states in the post and pre periods, but also from using treated and control firms within both states (Olden and Møen, 2022). More formally, I estimate the following triple difference specification:

$$Y_{i,t} = \beta Post\ L\text{-}NDA_{s(i),t} \times Pre\text{-}Harassment_i + \gamma_{s(i),t} + \nu_{s(i)} \times Pre\text{-}Harassment_i + \gamma_i + \delta_t + \eta_s + \epsilon_{i,t} \quad (3)$$

In this equation, i, s, and t are indices for firm, state headquarters, and year, respectively. Here, $Post\ L\text{-}NDA_{s(i),t}$ takes a value of one for the treated states for the post-period $[t+1, t+4]$ and zero for the pre-period $[t-3, t]$. This variable is zero for all non-treated states for the same period (i.e., $[t-3, t+4]$, where $t-3$ and $t+4$ are 2014 and 2022, respectively). For example, in California, the effective date of the regulation is January 1, 2019. The $Post\ L\text{-}NDA_{s(i),t}$ dummy in this case will get zeros for 2014-2018 and ones for 2019-2022 for all companies headquartered in California. $Pre-Harassment_i$ is the firm-level mean of workplace harassment, which measures the distance to treatment for firms headquartered in state s. The mean is fixed to the pre-estimation period, more specifically to the years 2011-2013, which captures the systematicity of harassment in the firm. That is, firms with higher scores are more exposed to the regulatory changes. Having the three-year (2011-2013) mean of the workplace harassment measure is more conservative than having only 2013, as some idiosyncratic changes in one year may lead to mislabeling the firms as high harassment.[21]

I include $State \times Year$ (i.e., $\gamma_{s(i),t}$) fixed effects, which capture state-level trends that may lead to the differences (i.e., not related to harassment) in the outcome variables. I also include $State$ (i.e., η_s) and $Year$ (i.e., δ_t) fixed effects, which are subsumed by $State \times Year$. In addition, I include $Firm$ (i.e., γ_i) fixed effects, which not only absorb the average harassment rate across firms prior to the analysis period (i.e., subsumes $\nu_{s(i)} \times$ *Pre-Harassment*$_i$), but also control for the idiosyncratic factors that may inf(def)late differences between groups. Finally, I include $Post$ *L-NDA*$_{s(i),t} \times$ *Pre-Log(# of Reviews)*$_i$ as a control to ensure that the effects are not driven by the number of reviews are posted on Glassdoor. Since the number of reviews is highly correlated with the number of employees, this allows me to control for firm size. For the baseline results, I cluster the standard errors at both the firm and state level (Abadie et al., 2023) as the coefficient of interest relies on the variation in these two variables.[22] The coefficient of interest, β, captures the identifying variation in the outcome variable. Specifically, the estimates show the difference in the change in innovation output of previously high harassment score firms in the treated states relative to the control group. The control group includes low harassment score firms in the treated states as well as high and low harassment score firms in the control states.

I begin by studying the impact of regulatory changes on workplace harassment. In the baseline equation above, I use two measures as the outcome variable: workplace harassment score based on anonymous reviews posted on Glassdoor and workplace harassment lawsuits

[21] For robustness, I fix the $Pre-Harassment_i$ to 2013 and find similar results.

[22] I conduct additional sensitivity tests where the standard errors are clustered at the firm level alone, and also two-way clustering at the firm and year level, and the results are statistically similar.

collected from the Violation Tracker database. The results for the workplace harassment score are presented in Panel A of Table 5.

Panel A						
	Workplace Harassment					
	(1)	(2)	(3)	(4)	(5)	(6)
Post N-NDA × Pre-Harassment	-0.103***	-0.102***	-0.103***	-0.112***	-0.075***	-0.082**
	(0.016)	(0.016)	(0.016)	(0.020)	(0.014)	(0.037)
Post L-NDA × Pre-Log(# of Reviews)	-0.000	-0.000	-0.000	0.000	-0.001	0.000
	(0.001)	(0.000)	(0.001)	(0.000)	(0.001)	(0.000)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry × Year FE	No	Yes	No	No	No	No
Obs.	14078	14078	6449	13525	9781	7690
R^2	0.521	0.529	0.537	0.524	0.494	0.740
Sample	Full	Full	Only Treated	Large Firms	Pre-Covid	More Crowd

Panel B						
	Workplace Harassment Lawsuit					
	(1)	(2)	(3)	(4)	(5)	(6)
Post N-NDA × Pre-Harassment	-0.161***	-0.159***	-0.143**	-0.180***	-0.217***	-0.321*
	(0.050)	(0.048)	(0.053)	(0.053)	(0.064)	(0.164)
Post L-NDA × Pre-Log(# of Reviews)	0.005	0.005	0.006	0.005	0.010	0.005
	(0.003)	(0.003)	(0.004)	(0.003)	(0.007)	(0.009)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry × Year FE	No	Yes	No	No	No	No
Obs.	13102	13102	5965	12562	9781	7029
R^2	0.285	0.293	0.248	0.285	0.333	0.310
Sample	Full	Full	Only Treated	Large Firms	Pre-Covid	More Crowd

Table 5: The Effects of NDA-limiting regulations on Workplace Harassment. This table reports results for the triple difference specification given in equation 3, where the outcome variable is Workplace Harassment and Workplace Harassment Lawsuits. Workplace Harassment is a measure of workplace harassment constructed using the textual content of anonymous employee reviews posted on Glassdoor, which is discussed in Section 3. Workplace Harassment Lawsuits is a dummy variable indicating whether the company has faced workplace harassment lawsuits, obtained from the Good Jobs First organization's Violation Tracker database. Industry classification is based on Fama-French 17 industries. The sample of "Large Firms" is those with a book value of assets above $100 million, the "Pre-Covid" sample is estimated for the period 2014-2019, and the "More Crowd" sample is those companies with above the cross-sectional median number of Glassdoor reviews. Variable definitions are provided in table A1. In parentheses are standard errors that are double-clustered at the firm and the state headquarter level.

In column 1, I document a significant decrease in the rate of anonymous harassment reports. The estimate indicates a 0.103 point decrease in workplace harassment scores relative to the control group; 46% relative to its mean (0.222). I then test the robustness of these results in columns 2-6. In Column 2 I control for industry trends, in Column 4 I restrict the sample to large firms (i.e., those with a book value of assets above $100 million), in Column 6 I retain firms with more than the median number of Glassdoor reviews, and I find similar effects. In particular, column 3 narrows the analysis to the treated group of states, as the never-treated group may not be a good control. The results remain economically and statistically consistent. In column 4, the analysis excludes the COVID period to account for variations in work-from-home policies that differ across states. Excluding these years leaves only one treatment year,

2019. However, this restriction does not affect the results, confirming the robustness of the findings.

To confirm the significant drop in workplace harassment, I next examine how harassment-related lawsuits change after the regulations. In Panel B of Table 5, I show that there is a significant drop in the likelihood of a lawsuit following the regulations. The effect is economically meaningful, a firm with previously one standard deviation higher workplace harassment is about 21% less likely to face harassment lawsuits. I then examine the robustness of these results to different sample restrictions and document similar effects.

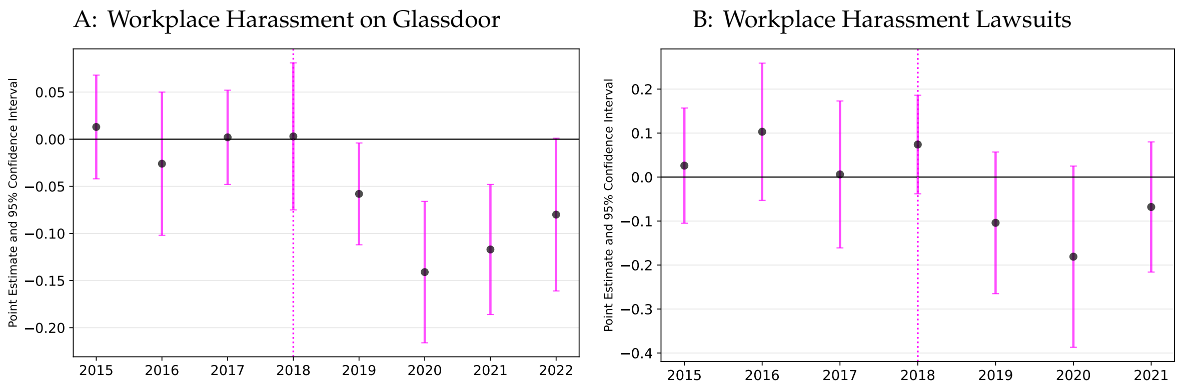

Figure 4: Dynamic Effects of NDA-Limiting Regulations on Workplace Harassment. This figure plots the estimated coefficients of the interaction term in Equation 4, where $Post\ L\text{-}NDA_{s,t}$ is interacted with a series of year indicators before ($[t-4, t]$) and after ($[t+1, t+4]$) the NDA limitation regulations. Panel A shows the results where the outcome variable is Workplace Harassment, while Panel B shows the results where the outcome variable is Workplace Harassment Lawsuits. Workplace Harassment is a measure of workplace harassment constructed using the textual content of anonymous employee reviews posted on Glassdoor, which is discussed in section 3. Workplace Harassment Lawsuits is a dummy variable indicating whether the company has faced workplace harassment lawsuits, obtained from the Good Jobs First organization's Violation Tracker database. The variable definitions can be found in Table A1. The standard errors are double-clustered at the firm and state headquarter level.

To examine the dynamic effects of the regulatory changes on workplace harassment and lawsuits, I interact $Post\ L\text{-}NDA_{s,t}$ with a series of year indicators both before ($[t-4, t]$) and after ($[t+1, t+4]$) the implementation of the NDA restriction regulations, as follows:

$$Y_{i,t} = \sum_{t=t-4}^{t+4} \beta_t Post\ L\text{-}NDA_{s(i),t} \times Pre\text{-}Harassment_i + \gamma_{s(i),t} + \nu_{s(i)} \times Pre\text{-}Harassment_i + \gamma_i + \delta_t + \eta_s + \epsilon_{i,t} \tag{4}$$

The estimated coefficients (i.e., β_t) of these interaction terms are shown in Figure 4. Panel A shows a significant impact of the regulations, beginning in the first year of treatment and continuing for another three years. Similarly, I look at the dynamic changes in harassment-related lawsuits following the restrictions on NDAs in Panel B and observe similar effects. Although the estimates are not significant at each post time point alone, as I show in Table 5, these effects are jointly significant.

The underlying assumption in my estimation is parallel trends. That is, firms with high harassment scores should have similar trends in the outcome variables of interest as firms with low harassment scores. Figure 4 confirms that this assumption is not violated. That is, there is no significant difference in harassment scores or harassment lawsuits between the two groups prior to the implementation of the regulations. Overall, the exogenous decline in workplace harassment suggests that the regulatory changes have had a deterrent effect, as the slight increase in the probability of being reported in the public media or in court would lead to significant reputational damage.

Having documented the exogenous decline in workplace harassment, I next examine how firms' innovation output changes in response to NDA limiting regulations. Specifically, following the large literature on innovation, I use three outcome variables: economic value of patents (Kogan et al., 2017), number of patents, and number of highly cited patents. Highly cited patents are those with above median total citations in the year of filing and their technology class. To ensure comparability across firms of different sizes, all three measures are normalized by the total number of employees.

In Table 6, I report the impact of workplace harassment on innovation. I find that regulations restricting firms' use of NDAs to conceal workplace harassment led to significantly higher innovation output in firms that had high ex-ante exposure to these regulatory changes. In column 1, I show that the effects are significant across all baseline outcome variables. The coefficients on $\beta Post\ \textit{L-NDA}_{s(i),t} \times \textit{Pre-Harassment}_i$ indicate that for those firms with a one standard deviation (0.025) higher pre-harassment rate, valuable innovation output increases by 10%, the number of patents by 5.6%, and the number of highly cited patents by 5%.

These effects are robust to the inclusion of industry trends in column 2, confirming that they are not driven by innovation output across industries. The results are robust to excluding never-treated states in column 3, suggesting that the results are not driven by any significant pre-treatment differences or trends between ever and never-treated states. The results are similar when I restrict the sample to larger firms and firms with more reviews on Glassdoor, implying that the effects are not driven by small firms. The effect is smaller in economic magnitude in column 5, but this only indicates the effect in the first treatment period, and also confirms that the effect is significant after omitting the Covid period.

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A						
	Log(economic value of patents per employee)					
Post L-NDA × Pre-Harassment	4.010***	3.301***	3.984***	4.706***	1.740***	5.097***
	(0.476)	(0.388)	(0.502)	(0.528)	(0.405)	(1.520)
Post L-NDA × Pre-Log(# of Reviews)	-0.075***	-0.075***	-0.077***	-0.073***	0.021***	-0.108***
	(0.012)	(0.006)	(0.013)	(0.011)	(0.007)	(0.026)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry × Year FE	No	Yes	No	No	No	No
Obs.	14078	14078	6449	13525	9781	7690
R^2	0.862	0.873	0.864	0.864	0.947	0.887
Sample	Full	Full	Only Treated	Large Firms	Pre-Covid	More Crowd
Panel B						
	Log(# of patents per employee)					
Post L-NDA × Pre-Harassment	2.260***	1.825***	2.255***	2.938***	1.480**	3.876***
	(0.364)	(0.342)	(0.386)	(0.418)	(0.700)	(0.864)
Post L-NDA × Pre-Log(# of Reviews)	-0.020***	-0.018***	-0.020**	-0.022***	0.011*	-0.053***
	(0.007)	(0.004)	(0.007)	(0.006)	(0.006)	(0.017)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry × Year FE	No	Yes	No	No	No	No
Obs.	14078	14078	6449	13525	9781	7690
R^2	0.836	0.850	0.834	0.843	0.940	0.859
Sample	Full	Full	Only Treated	Large Firms	Pre-Covid	More Crowd
Panel C						
	Log(# of highly cited patents per employee)					
Post L-NDA × Pre-Harassment	1.960***	1.552***	1.973***	2.392***	1.518***	4.236***
	(0.353)	(0.439)	(0.411)	(0.360)	(0.219)	(1.064)
Post L-NDA × Pre-Log(# of Reviews)	-0.027***	-0.024***	-0.028***	-0.027***	-0.014**	-0.054***
	(0.007)	(0.003)	(0.007)	(0.005)	(0.005)	(0.017)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry × Year FE	No	Yes	No	No	No	No
Obs.	14078	14078	6449	13525	9781	7690
R^2	0.742	0.761	0.745	0.750	0.877	0.774
Sample	Full	Full	Only Treated	Large Firms	Pre-Covid	More Crowd

Table 6: The Impact of Workplace Harassment on Innovation Output. This table presents results for the triple difference specification as specified in equation 3, where the outcome variables are *Log(economic value of patents per employee)*, *Log(# of patents per employee)*, and *Log(# of highly cited patents per employee)*. Data on economic value of patents, number of patents, and citation of patents are from Kogan et al. (2017). The sample of "Large Firms" is those with a book value of assets above \$100 million, the "Pre-Covid" sample is estimated for the period 2014-2019, and the "More Crowd" sample is those companies with above the cross-sectional median number of Glassdoor reviews. Variable definitions are provided in Table A1. In parentheses are standard errors that are double-clustered at the firm and the state headquarter level.

Next, I study the dynamic effects of workplace harassment. I estimate Eq. 4 and visualize the estimates on β_t in Figure 5. The figure shows the change in the innovation output measures in three panels four years before and after the regulation. First, the event-study plots show no pre-trends for any of the innovation output measures. The Wald test fails to reject that the coefficients are equal to zero before the regulatory changes, implying that there are no differential pre-trends in high-harassment firms relative to low-harassment firms. Second, innovation output increases steadily in the first two years and the effects become more pronounced in the third and fourth years. This observation intuitively suggests that as the workplace environment continues to improve over time, the positive impact on innovation becomes more pronounced and potentially long-lasting.

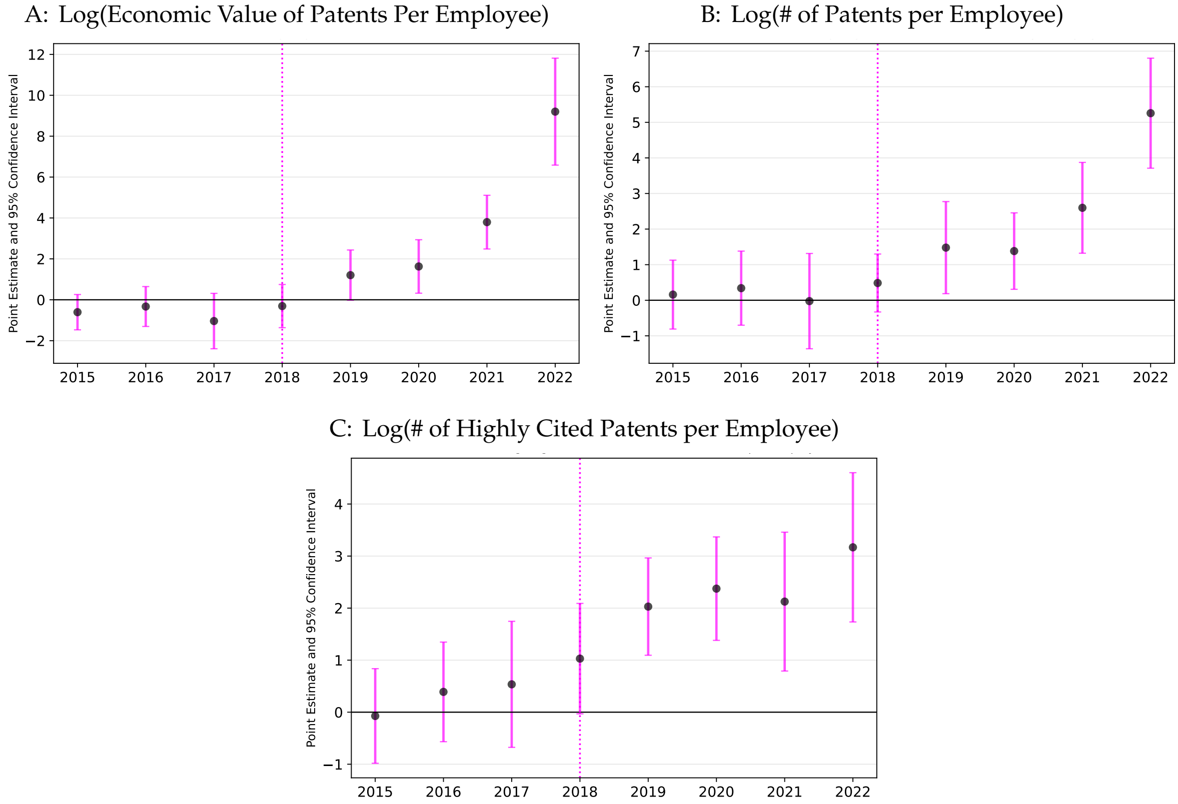

Figure 5: The Impact of Workplace Harassment on Innovation Output - Dynamic Effects. This figure plots the estimated coefficients of the interaction term in Equation 4, where *Post L-NDA*$_{s,t}$ is interacted with a series of year indicators before ($[t-4, t]$) and after ($[t+1, t+4]$) the NDA limitation regulations. Panels A, B, and C show the results where the outcome variable is Log(Economic Value of Patents Per Employee), Log(# of Patents per Employee), and Log(# of Highly Cited Patents per Employee), respectively. Data on economic value of patents, number of patents, and citation of patents are from Kogan et al. (2017). Variable definitions are provided in Table A1. The standard errors are double-clustered at the firm and state headquarter level.

Inventor Team Heterogeneity Given that workplace harassment occurs on the basis of gender, race, ethnicity, and other individual characteristics, I next investigate how the impact of workplace harassment on innovation varies depending on the specific characteristics of innovation teams along these dimensions. Specifically, I study inventor team composition by gender, ethnicity and race. On average, inventor teams are majority male (90%), European (82%) and white (75%). I begin with studying the team by their gender composition, and I report the results in Panel A of Table 7. When the inventor team is composed of only men (i.e, among those who produce patent in the year t), I find no significant effects shown in the first column. Next, I relax the sample and allow for gender heterogeneity in the second column and find that the effect of workplace harassment on innovation is primarily driven by the teams that include at least one woman. I go one step further and test whether the effects are driven by

	Log(economic value of patents per employee)			
Panel A				
Firms with inventor teams:	Only men	At least one woman	With Women Minority	Less 50
Post L-NDA × Pre-Harassment	-0.040	4.211***	4.644***	
	(1.668)	(1.336)	(1.278)	
Post L-NDA × Pre-Log(# of Reviews)	-0.041	0.038	0.034	insufficient
	(0.060)	(0.032)	(0.032)	obs.
Firm FE	Yes	Yes	Yes	
State × Year FE	Yes	Yes	Yes	
Obs.	937	4380	4325	
R^2	0.685	0.844	0.843	
Panel B				
Firms with inventor teams:	No Ethnic Minority	With Ethnic Minority	With Minority of African descent	With Minority of Asian descent
Post L-NDA × Pre-Harassment	-2.954	4.817***	2.669*	4.922***
	(4.197)	(1.374)	(1.486)	(1.381)
Post L-NDA × Pre-Log(# of Reviews)	-0.164***	0.021	0.031	0.021
	(0.035)	(0.025)	(0.036)	(0.026)
Firm FE	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes
Obs.	569	4504	3963	4435
R^2	0.674	0.851	0.860	0.853
Panel C				
Firms with inventor teams:	No Racial Minority	With Racial Minority	With Black Minority	With Asian Minority
Post L-NDA × Pre-Harassment	-0.797	5.102***	2.789*	5.089***
	(2.192)	(1.595)	(1.626)	(1.538)
Post L-NDA × Pre-Log(# of Reviews)	-0.002	0.011	0.035	0.037
	(0.066)	(0.033)	(0.042)	(0.031)
Firm FE	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes
Obs.	758	4063	2129	3943
R^2	0.621	0.845	0.898	0.846

Table 7: The Impact of Workplace Harassment on Innovation: Inventor Team Heterogeneity. This table presents the results for the triple difference specification, as specified in equation 3, in different subsamples by the gender, ethnicity, and race composition of inventor teams. Sample splitting is done by the cross-sectional characteristics of firm inventor teams during the 2014-2022 estimation period. Data on inventor gender is obtained from PatentsView. Data on inventor's race and ethnicity is extracted based on inventor's last name using the ethnicolr algorithm. The variable definitions can be found in the table A1. Standard errors in parentheses are double-clustered at the firm and state headquarter level.

the teams where women are in the minority, meaning that the male composition of the team should be above 50%, and the results show that the effect is driven primarily from the firms which have female minorities. In addition, the sample size indicates that the majority of firms have gender-mixed teams where women are in the minority.

In Panels B and C of Table 7, I report the results by the ethnic and racial composition of the inventor teams. To identify the ethnic/racial background of the inventors, I rely on their surnames provided by the PatensView database. I then use the "ethnicolr" library from Python to estimate the probability that the inventors belong to one of the ethnic/racial groups. In my sample, inventors are grouped into 3 ethnic groups (European, African, and Asian) and 3 racial groups (White, Black, Asian). I find that when active inventor team members do not include minorities (African and Asian in terms of ethnicity, and Black and Asian in terms of race), there are no significant effects. However, when I restrict the sample to teams with the minority ethnic and racial groups, the baseline effects are significantly pronounced. Taken together, these results show that workplace harassment is driven by inventor teams with minority representation.

4.3 Robustness Tests

I conduct several additional analyses to show that the results are robust. So far in the baseline estimation I use the continuous variable to test the exposure of high harassment score firms to the regulatory changes. To explicitly show that the effect is coming from the high harassment score firms, I create the dummy variable *Pre-High Harass* indicating whether the firm is ranked above the *Pre-Harassment* score. Table A8 shows that firms with previously above median harassment score innovate significantly more (i.e., 31% in terms of economic value) compared to firms below the cutoff, supporting the conjecture that the documented effects are driven by firms with previously high harassment environments.

Second, in my analysis, if firms do not innovate in year t, these missing values are filled with zeros following Kogan et al. (2017). It is possible that such firms are not innovators. I first restrict my sample to firms that have ever innovated, but I observe that all firms in my sample have innovated at some point in their lives. I then restrict the sample to firms that are actively innovating. The idea is to test that the effects I document are not driven by the innovation intensity of the firms. Thus, I restrict the sample to firms that have more patents than the median firm and repeat the baseline estimation. The results are reported in table A9 and show that the effects are very similar to the baseline, meaning that the firms in the baseline are comparable.

Third, in the main analysis, I normalize innovation performance by number of employees (Gao and Zhang, 2017) because harassment scores are averaged across all reviews. And the number of reviews is strongly correlated with the number of employees. I check the robustness of this measure by normalizing the economic value of patents by the book value of assets as in Kogan et al. (2017), and find that the result is robust as shown in column 1 of table A10. In addition, I check the robustness of the citation-based measure of innovation output by counting the number of patents in the top decile and quantile among other patents in the same filing cohort and technology class in terms of total citations received, and find similar effects as shown in columns 2 and 3 in Table A10.

Fourth, in the baseline, firms are treated if they are headquartered in the state that passed the NDA-limiting regulations. One of the concerns might be that it might not be a state where the firm's employees are located. Although firms' strategic activities usually centered in their headquarters, such as innovation, I also conduct additional test based on the location of majority of employees. That is, on Glassdoor, employees can post where they are located when they post reviews. I use this information and define the state of the firm from where the majority of reviews are posted in year t. In this case, the firm is treated as if the NDA restricting rules were passed in the state where the majority of its employees are located. I report the results of this test in table A11. Again, I find that the effects are very similar to the baseline in both economic magnitude and statistical power.

Finally, I test the robustness of the parallel trend assumption. To do so, I conduct two separate tests. First, I restrict the sample to firms located in the treated states, excluding firms that are never treated. Here, the plausible exogenous variation comes from time variation in when states enact the laws. I estimate the specification for the dynamic effects given in equation

4 and present the results in Figure A1. The results show that in the pre-period, there is no systematic difference between high and low harassment firms across treated states.

In addition, I test the robustness of the parallel trends assumption in the stacked difference-in-difference setting (Cengiz et al., 2019). For each treatment cohort, I create a subsample that includes the treated cohort and never-treated states. Given that the last treatment starts in 2021, I limit the post-period to 2 treatment years and the pre-period to 4 years across treatment cohorts and re-estimate the equation 4. In this specification, I replace *State* $\times$ *Year FE* and *Firm FE* with *State* $\times$ *Year* $\times$ *Cohort FE* and *Firm* $\times$ *Cohort FE*, respectively. Figure A2 shows similar dynamic effects and provides reassurance that treatment heterogeneity does not lead to violation of parallel trends.

5 Economic Mechanism

In the previous section, I document that an exogenous decrease in workplace harassment leads to significantly higher firm innovation output. A natural question is how a decrease in harassment leads to more innovation.

Teams play an increasingly dominant role in the production of innovation and knowledge (Wuchty et al., 2007). Increasing harassment in the workplace can lead to productivity losses by compromising team capital. For team capital as a conceptual framework, I build on a similar idea to Jaravel et al. (2018), who show that the premature death of the inventor member can lead to persistent losses in the productivity of co-inventors. That is, such collaboration is unique to those team members who co-innovate, and not easily replaceable (Jäger and Heining, 2022). Building on this, I argue that increasing workplace harassment can destroy team capital in the following ways. First, experiencing workplace harassment (directly or as a witness) can lead to psychological distress among team members. Second, victims of harassment or those with a strong aversion to harassment may leave such workplaces as a result, leading to the destruction of creative synergy. Finally, companies with a reputation for workplace harassment may have difficulty attracting the talent they need. I test the role of team capital as an economic mechanism in two steps.

Inventor Productivity: First, I begin by examining whether improvements in workplace harassment actually lead to higher inventor productivity. For the team capital mechanism to work, we should observe a positive response in inventor productivity as workplace harassment decreases in the firm. To do so, I measure inventors' productivity by calculating the number of patents they develop, as well as the number of high quality patents proxied by their economic value and total citations (Bernstein et al., 2021). I estimate the similar regression as in Eq. 3 but the outcome variables are at the inventor level as follows:

$$Inventor\ Productivity_{i,j,s,t} = \beta \text{ Post L-NDA}_{s,t} \times \text{Pre Workplace Harassment}_j + \gamma_i + \omega_j + \delta_{s,t} + \epsilon_i, \quad (5)$$

where i, j, s, t index inventor, firm, state, and year. This estimation allows me to compare the changes in productivity of inventors in firms with previously higher rates of workplace harassment to the others within the treated states, as well as to compare these changes to all other firms in never-treated states. I report the estimation results in Table 8.

In Panel A, I report changes in inventor productivity in terms of highly valued innovation output. A highly valued patent is one that ranks above the median of all other patents in the same filing year and technology class. In column 1, I show that productivity in terms of the number of highly valued patents increases by 5.6% ($0.025 \times 2.241 \times 100$) in response to a one standard deviation lower rate of workplace harassment. One argument could be that productivity growth is driven by industry-specific differences. I control for such trends in column 2, and the effects are similar.

Panel A

	Log(# of high valued patents)			
	(1)	(2)	(3)	(4)
Post L-NDA × Pre-Harassment	2.241***	1.853***		
	(0.798)	(0.598)		
Post L-NDA × Pre-Harassment High			0.063***	0.054***
			(0.022)	(0.019)
Log(# of Reviews)	-0.032**	-0.026**	-0.032**	-0.027**
	(0.015)	(0.013)	(0.015)	(0.013)
Inventor FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes
Industry × Year FE	No	Yes	No	Yes
Obs.	371,428	355,177	371,428	355,177
R^2	0.678	0.685	0.678	0.685

Panel B

	Log(# of patents)			
	(5)	(6)	(7)	(8)
Post L-NDA × Pre-Harassment	1.376**	1.589***		
	(0.574)	(0.585)		
Post L-NDA × Pre-Harassment High			0.058***	0.060***
			(0.010)	(0.014)
Log(# of Reviews)	-0.021	-0.008	-0.021	-0.008
	(0.013)	(0.005)	(0.013)	(0.005)
Inventor FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes
Industry × Year FE	No	Yes	No	Yes
Obs.	371,428	355,177	371,428	355,177
R^2	0.626	0.630	0.626	0.630

Panel C

	Log(# of highly cited patents)			
	(9)	(10)	(11)	(12)
Post L-NDA × Pre-Harassment	1.383**	2.114*		
	(0.664)	(1.116)		
Post L-NDA × Pre-Harassment High			0.062***	0.062***
			(0.016)	(0.020)
Log(# of Reviews)	-0.034	-0.005	-0.034	-0.005
	(0.031)	(0.014)	(0.030)	(0.015)
Inventor FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes
Industry × Year FE	No	Yes	No	Yes
Obs.	371,428	355,177	371,428	355,177
R^2	0.571	0.577	0.571	0.577

Table 8: The Effect of Workplace Harassment on Inventor Productivity. This table shows the results of estimating the Equation 5, where the outcome variable is inventor productivity. In Panel A, inventor productivity is measured as the log of one plus the number of highly valued patents. A patent is highly valued if its economic value is above the median of all other patents in the same filing year and technology class. In Panel B, inventor productivity is measured as the log of one plus the number of patents filed in year t. In Panel C, inventor productivity is measured as the log of one plus the number of highly cited patents. A patent is highly cited if it ranks above the median of all (i.e., forward) citations it has received among all other patents in the same filing year and technology class. Technology class is defined based on the CPC classification. Pre-Harassment High is a dummy variable indicating whether the firm ranks above the median Pre-Harassment level in the cross section of firms. Industry fixed effects are based on the Fama-French 17 classification. State fixed effects are based on the state of residence of the inventor. Variable definitions are provided in table A1. In parentheses are standard errors clustered at the firm and state levels.

Next, to explicitly test whether the effect is indeed driven by the high harassment firms, I rank the firms by the *Pre-Workplace Harassment* level (i.e., the mean of the 2011-2013 workplace harassment score) and create a dummy variable indicating whether the firm has a workplace harassment rate above the median. The results in columns 3 and 4 show that the productivity of inventors in firms above the cutoff harassment rate grows by 5.4-6.3% more.

I repeat these tests using alternative proxies for the inventor's productivity. In Panel B of Table 8, I document that inventor productivity increases similarly in terms of the number of patents. Moreover, in Panel C, I use the number of highly cited patents as a proxy for the inventor's innovation output and find similar effects. These results suggest that a decrease in workplace harassment leads to higher inventor productivity, both in terms of quality and quantity. Moreover, the growth in productivity by economic size is remarkably similar to the firm-level results, underscoring that the growth in firm innovation output in response to lower workplace harassment is largely driven by growth in inventor productivity.

Turnover of Skilled Employees: Having shown that inventor productivity increases in response to an exogenous reduction in workplace harassment, I next move on to an analysis of skilled worker turnover. To track the workforce, particularly with respect to skill levels, I use a unique dataset provided by the BrightsData Initiative (BI), which aggregates publicly available LinkedIn data. This dataset allows me to observe the employment history of individuals, including company names, start and end dates of employment, and educational attainment. Using this information, I construct measures of worker turnover by education and within-industry experience. The detailed description of the data and variables can be found in section 2.4.

First, I conduct an empirical analysis to show that the proxies for skilled worker turnover are indeed correlated with innovation output. Specifically, I estimate the following regression equation:

$$Innovation_{i,t} = \beta\, Workforce\ Loss\ Rate(type)_{i,t} + \Gamma X_{i,t} + \gamma_i + \delta_t + \epsilon_{i,t} \tag{6}$$

The indices i and t represent the firm and the year, respectively. The dependent variable $Innovation_{i,t}$ captures the dollar value of innovation output per employee. The independent variable *Workforce Loss Rate(type)* serves as a proxy for the loss of employees with specific skill attributes. The control variables $X_{i,t}$ are the same as those used in all baseline estimations, as described in equation 3. In addition, the regression includes firm fixed effects γ_i and year fixed effects δ_t to obtain within-firm estimates. Standard errors are clustered at the firm level to account for potential within-firm correlation.

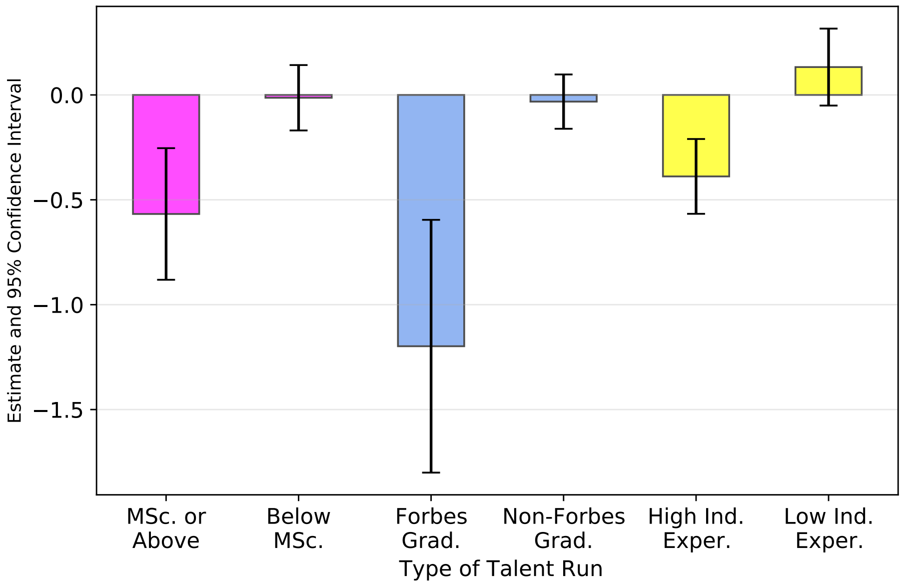

Figure 6: Association Between Innovation and Workforce Loss Rate by Skill. The figure shows the point estimates and 95% confidence bands for the regressions of innovation on the workforce loss rate specified in equation 6. Innovation is the log of the total economic value of patents per employee. The workforce loss rate is calculated over three employee characteristics for skill: advanced degree (master's or higher), Forbes 100 college attendance, and industry experience. The construction of the workforce loss rate is described in equation 1 in section 2.4. The controls included are the same as in Table 4. Variable definitions are provided in Table A1. Standard errors are clustered at the firm level.

The results for the relationship between innovation and the net outflow of workers with different skill attributes are shown in Figure 6. The first two bars show the relationship between innovation and the net flow of employees with an advanced degree (i.e., at least a Master's degree) compared to those with less than a Master's degree. The figure shows that the change in innovation output can be predicted by the change in the flow of employees with advanced degrees, but not otherwise. The next two bars show the effect on innovation of losing employees with degrees from the top 100 Forbes colleges, and the results are consistent with the previous one. Finally, the last two bars illustrate the relationship between innovation and industry expertise. The results show that the relationship is significant only for the net outflow of employees with high industry experience, suggesting that the loss of employees with extensive industry knowledge is relevant to innovation output. Taken together, this exercise suggests that the skill variables used in the analysis can explain variation in innovation activity and gives confidence to use them to further investigate the role of team capital.

	Workforce Loss Rate					
	Advanced degree		Forbes 100 College		Industry Experience	
	Yes	*No*	*Yes*	*No*	*High*	*Low*
Post L-NDA × Pre-Harassment	-0.122**	0.075	-0.104***	-0.026	-0.420***	0.129
	(0.056)	(0.142)	(0.025)	(0.127)	(0.087)	(0.100)
Post L-NDA × Pre-Log(# of Reviews)	-0.002*	-0.008***	-0.001**	-0.009***	-0.011***	-0.002
	(0.001)	(0.003)	(0.001)	(0.003)	(0.003)	(0.002)
Firm	Yes	Yes	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	9792	9937	9366	9946	8977	10001
R^2	0.314	0.390	0.280	0.394	0.623	0.583

Table 9: The Impact of Workplace Harassment on Skilled Workforce Loss Rate. This table presents the results for the triple difference specification, as specified in Equation 3, with the outcome variables being workforce loss rate by three empoyee characteristics: advanced degree (masters or above degree), forbes 100 college attendance and industry experience. The construction of "workforce loss rate" is described in Equation 1 in Section 2.4. The included controls are the same as those in Table 4. The variable definitions can be found in Table A1. The standard errors, presented in parentheses, are double-clustered at the firm and state headquarter level.

To test whether skilled turnover changes in response to the reduction in workplace harassment caused by the regulatory changes, I re-estimate the baseline triple difference equation 3 using (un)skilled flow variables as the outcome. I report the results in table 9. Column 1 shows that in firms with previously higher rates of workplace harassment, the loss of workers with advanced degrees drops significantly after the regulations are enacted. By comparison, the loss of workers without advanced degrees does not show a significant response. In column 3, I compare the impact using the loss of employees with degrees from the top 100 Forbes colleges. Again, I find that the results are significant only for the skilled types. Finally, I run the same estimates for the loss of workers with high and low industry experience (within a given Fama-French 17 industry group). The column with the high industry experience outflow again confirms that the decrease in the loss of workers is mainly driven by the talented group. These results shed some light on the functioning of team capital as an economic mechanism. The retention of skilled workers suggests that innovators have the ability to maintain collaborative dynamics, which consequently leads to higher productivity in inventive work.

6 From Harassment to Overall Firm Performance

In the previous section, I document that firms' innovation output increases significantly in response to decreases in workplace harassment caused by regulatory changes. Next, I ask whether such positive changes in workplace climate and firm innovation output translate into higher firm sales and profits. To do so, I substitute these variables for the outcome variables in the baseline triple difference equation and report the results in Table 10. Columns 1 and 2 show

	Log(Sales Per Employee)		Profit Growth	
	(1)	(2)	(3)	(4)
Post L-NDA × Pre-Harassment	2.122**	2.136**	0.817***	0.880***
	(0.909)	(0.900)	(0.152)	(0.159)
Post L-NDA × Pre-Log(# of Reviews)	-0.010	-0.008	0.002	0.005
	(0.009)	(0.009)	(0.008)	(0.008)
Firm FE	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes
Industry × Year FE	No	Yes	No	Yes
Obs.	14078	14078	13644	13644
R^2	0.920	0.922	0.277	0.323

Table 10: The Impact of Workplace Harassment on Firm Performance. This table shows the results of estimating Equation 3, where the outcome variable is Log(Sales per Employee) and Profit Growth. Profit Growth is defined as growth in sales minus cost of goods sold. Industry classification is based on the Fama-French 17 industries. Variable definitions are provided in Table A1. In parentheses are standard errors clustered at the firm and state headquarters level.

that after the regulatory changes, firms with previously higher levels of workplace harassment generate significantly more sales per employee. Specifically, sales per employee increase by 5.3% for firms with previously one standard deviation higher workplace harassment. I report the estimates for profit growth in columns 3 and 4. I find that a one standard deviation decrease in workplace harassment leads to a 2.2% higher growth rate in firm profits. Taken together, these results suggest that workplace harassment poses significant costs on firm performance and that the regulatory changes that deter firms by eliminating nondisclosure agreements have led not only to lower levels of workplace harassment but also to better firm and employee performance.

7 Conclusion

Workplace harassment is a pervasive problem that affects the lives and work performance of individuals. Studying this problem within organizations and its impact on employee performance and overall organizational performance has been difficult due to the hidden nature of this problem. We only learn about such problems when victims take bold decisions and sue companies, which becomes scandals. In this paper, I use online anonymous employee reviews to measure the systemic extent of the problem and assess its impact on employees' output, especially in innovation activities. Such a measure is effective in capturing the many unreported (i.e., not disseminated in major news media) problems, which allows me to identify the impact of recent regulatory changes that limit the ability of firms to hide workplace harassment using tools such as nondisclosure agreements.

I document an exogenous reduction in workplace harassment and an improvement in job satisfaction as reported on Glassdoor following the passage of the regulations. The results suggest that the regulations induced firms with higher exposure to these regulations to improve their workplace climate to avoid reputational damage. I document that this exogenous shift led to a significant increase in the innovation output of firms with previously high workplace harassment scores. The documented effects are higher for teams that are heterogeneous along the dimensions of gender, race, and ethnicity. I show that higher innovation output in response to decreasing workplace harassment is largely due to improvements in the collaborative work environment and the maintenance of team capital. Overall, this paper demonstrates that workplace harassment poses significant costs on firm performance and that regulatory changes that limit firms' ability to silence victims have far-reaching implications.

References

Abadie, Alberto, Susan Athey, Guido W Imbens, and Jeffrey M Wooldridge, "When should you adjust standard errors for clustering?," *The Quarterly Journal of Economics*, 2023, *138* (1), 1–35.

Acharya, Viral V, Ramin P Baghai, and Krishnamurthy V Subramanian, "Wrongful discharge laws and innovation," *The Review of Financial Studies*, 2014, *27* (1), 301–346.

Adams-Prassl, Abi, Kristiina Huttunen, Emily Nix, and Ning Zhang, "Violence Against Women at Work," *The Quarterly Journal of Economics*, 2023.

Aghion, Philippe and Peter Howitt, "A model of growth through creative destruction," *Econometrica*, 1992, *60* (2), 323–351.

___, John Van Reenen, and Luigi Zingales, "Innovation and institutional ownership," *American economic review*, 2013, *103* (1), 277–304.

Angrist, Joshua D and Jörn-Steffen Pischke, *Mostly harmless econometrics: An empiricist's companion*, Princeton university press, 2009.

Au, Shiu-Yik, Ming Dong, and Andreanne Tremblay, "How much does workplace sexual harassment hurt firm value?," *Journal of Business Ethics*, 2023, pp. 1–23.

Babina, Tania, Asaf Bernstein, and Filippo Mezzanotti, "Financial Disruptions and the Organization of Innovation: Evidence from the Great Depression," *The Review of Financial Studies*, 2023.

Balasubramanian, Natarajan, Evan Starr, and Shotaro Yamaguchi, "Bundling postemployment restrictive covenants: When, why, and how it matters," Technical Report 2021.

Batut, Cyprien, Caroline Coly, Sarah Schneider-Strawczynski et al., "It's a man's world: culture of abuse,# MeToo and worker flows," 2021.

Bernstein, Shai, "Does going public affect innovation?," *The Journal of finance*, 2015, *70* (4), 1365–1403.

___, Timothy McQuade, and Richard R Townsend, "Do household wealth shocks affect productivity? evidence from innovative workers during the great recession," *The Journal of Finance*, 2021, *76* (1), 57–111.

Boudreau, Laura E, Sylvain Chassang, Ada Gonzalez-Torres, and Rachel M Heath, "Monitoring Harassment in Organizations," Technical Report, National Bureau of Economic Research 2023.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, "The effect of minimum wages on low-wage jobs," *The Quarterly Journal of Economics*, 2019, *134* (3), 1405–1454.

Chen, Mark A, Qinxi Wu, and Baozhong Yang, "How valuable is FinTech innovation?," *The Review of Financial Studies*, 2019, *32* (5), 2062–2106.

Clermont, Kevin M and Stewart J Schwab, "Employment discrimination plaintiffs in federal court: From bad to worse," *Harv. L. & Pol'y Rev.*, 2009, *3*, 103.

Coviello, Decio, Erika Deserranno, Nicola Persico, and Paola Sapienza, "Effect of mood and worker incentives on workplace productivity," *The Journal of Law, Economics, and Organization*, 2022.

Custódio, Cláudia, Miguel A Ferreira, and Pedro Matos, "Generalists versus specialists: Lifetime work experience and chief executive officer pay," *Journal of Financial Economics*, 2013, *108* (2), 471–492.

Dahl, Gordon B and Matthew M Knepper, "Why is workplace sexual harassment underreported? the value of outside options amid the threat of retaliation," Technical Report, National Bureau of Economic Research 2021.

Dim, Chukwuma, "Should retail investors listen to social media analysts? Evidence from text-implied beliefs," 2020.

Dougal, Casey, Thomas P Griffin, and Irena Hutton, "A culture of discrimination: Evidence on the causes and consequences of civil rights litigation against US corporations," *Available at SSRN 3967456*, 2022.

Dube, Svenja and Chenqi Zhu, "The Disciplinary Effect of Social Media: Evidence from Firms' Responses to Glassdoor Reviews," *Journal of Accounting Research*, 2021, *59* (5), 1783–1825.

Edmans, Alex, "Does the stock market fully value intangibles? Employee satisfaction and equity prices," *Journal of Financial economics*, 2011, *101* (3), 621–640.

Fedyk, Anastassia and James Hodson, "Trading on Talent: Human Capital and Firm Performance," *Review of Finance*, 2022, *27* (5), 1659–1698.

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

Gao, Huasheng and Wei Zhang, "Employment nondiscrimination acts and corporate innovation," *Management Science*, 2017, *63* (9), 2982–2999.

Glaeser, Edward L, Hedi D Kallal, Jose A Scheinkman, and Andrei Shleifer, "Growth in cities," *Journal of Political Economy*, 1992, *100* (6), 1126–1152.

Goldstein, Itay, Chester S Spatt, and Mao Ye, "Big data in finance," *The Review of Financial Studies*, 2021, *34* (7), 3213–3225.

Graham, John R, Jillian Grennan, Campbell R Harvey, and Shivaram Rajgopal, "Corporate culture: Evidence from the field," *Journal of Financial Economics*, 2022, *146* (2), 552–593.

Green, T Clifton, Ruoyan Huang, Quan Wen, and Dexin Zhou, "Crowdsourced employer reviews and stock returns," *Journal of Financial Economics*, 2019, *134* (1), 236–251.

Grennan, Jillian and Kai Li, "Corporate culture: A review and directions for future research," *Available at SSRN*, 2022.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales, "The value of corporate culture," *Journal of Financial Economics*, 2015, *117* (1), 60–76.

Hacamo, Isaac, "Racial Prejudice in the Workplace and Firm Boycotts," *Available at SSRN 4033827*, 2023.

Hanley, Kathleen Weiss and Gerard Hoberg, "Dynamic interpretation of emerging risks in the financial sector," *The Review of Financial Studies*, 2019, *32* (12), 4543–4603.

Huang, Allen, Kai Wai Hui, and Reeyarn Zhiyang Li, "Federal judge ideology: A new measure of ex ante litigation risk," *Journal of Accounting Research*, 2019, *57* (2), 431–489.

Jäger, Simon and Jörg Heining, "How substitutable are workers? evidence from worker deaths," Technical Report, National Bureau of Economic Research 2022.

Jaravel, Xavier, Neviana Petkova, and Alex Bell, "Team-specific capital and innovation," *American Economic Review*, 2018, *108* (4-5), 1034–1073.

Jeffers, Jessica S, "The Impact of Restricting Labor Mobility on Corporate Investment and Entrepreneurship," *The Review of Financial Studies*, 2023, p. hhad054.

Joyce, Stephen, "States Act to Curb Employer Overuse of Non-Disclosure Agreements," *Bloomberg Law*, 2022. Available at: Link to article.

Kempf, Elisabeth and Oliver Spalt, "Attracting the sharks: Corporate innovation and securities class action lawsuits," *Management Science*, 2023, *69* (3), 1805–1834.

Kerr, William R and William F Lincoln, "The supply side of innovation: H-1B visa reforms and US ethnic invention," *Journal of Labor Economics*, 2010, *28* (3), 473–508.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, "Technological innovation, resource allocation, and growth," *The Quarterly Journal of Economics*, 2017, *132* (2), 665–712.

Li, Kai, Feng Mai, Rui Shen, and Xinyan Yan, "Measuring corporate culture using machine learning," *The Review of Financial Studies*, 2021, *34* (7), 3265–3315.

Manso, Gustavo, "Motivating innovation," *The journal of finance*, 2011, *66* (5), 1823–1860.

Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean, "Distributed representations of words and phrases and their compositionality," *Advances in neural information processing systems*, 2013, *26.*

___ **, Kai Chen, Greg Corrado, and Jeffrey Dean**, "Efficient estimation of word representations in vector space," *arXiv preprint arXiv:1301.3781*, 2013.

Moser, Petra, Alessandra Voena, and Fabian Waldinger, "German Jewish émigrés and US invention," *American Economic Review*, 2014, *104* (10), 3222–3255.

Neal, Derek, "Industry-specific human capital: Evidence from displaced workers," *Journal of labor Economics*, 1995, *13* (4), 653–677.

Niessen-Ruenzi, Alexandra and Leah Zimmerer, "The Value of Skill Signals for Women's Careers," *Available at SSRN 3987238*, 2021.

Okechukwu, Cassandra A, Kerry Souza, Kelly D Davis, and A Butch De Castro, "Discrimination, harassment, abuse, and bullying in the workplace: contribution of workplace injustice to occupational health disparities," *American Journal of Industrial Medicine*, 2014, *57* (5), 573–586.

Olden, Andreas and Jarle Møen, "The triple difference estimator," *The Econometrics Journal*, 2022, *25* (3), 531–553.

Oswald, Andrew J, Eugenio Proto, and Daniel Sgroi, "Happiness and productivity," *Journal of Labor Economics*, 2015, *33* (4), 789–822.

Raver, Jana L and Lisa H Nishii, "Once, twice, or three times as harmful? Ethnic harassment, gender harassment, and generalized workplace harassment.," *Journal of Applied Psychology*, 2010, *95* (2), 236.

Romer, Paul M, "Endogenous technological change," *Journal of Political Economy*, 1990, *98* (5, Part 2), S71–S102.

Rospenda, Kathleen M, Judith A Richman, and Candice A Shannon, "Prevalence and mental health correlates of harassment and discrimination in the workplace: Results from a national study," *Journal of Interpersonal Violence*, 2009, *24* (5), 819–843.

Sautner, Zacharias, Laurence Van Lent, Grigory Vilkov, and Ruishen Zhang, "Firm-level climate change exposure," *The Journal of Finance*, 2023, *78* (3), 1449–1498.

Schmitt, Michael T, Nyla R Branscombe, Tom Postmes, and Amber Garcia, "The consequences of perceived discrimination for psychological well-being: a meta-analytic review.," *Psychological Bulletin*, 2014, *140* (4), 921.

Sheng, Jinfei, "Asset pricing in the information age: Employee expectations and stock returns," *Available at SSRN 3321275*, 2021.

Sockin, Jason, Aaron Sojourner, and Evan Starr, "Non-disclosure agreements and externalities from silence," 2022.

Tambe, Prasanna, Lorin Hitt, Daniel Rock, and Erik Brynjolfsson, "Digital capital and superstar firms," Technical Report, National Bureau of Economic Research 2020.

WBI, "2021 WBI U.S. Workplace Bullying Survey," *Workplace Bullying Institute*, 2021. Available at: Link to report.

Wuchty, Stefan, Benjamin F Jones, and Brian Uzzi, "The increasing dominance of teams in production of knowledge," *Science*, 2007, *316* (5827), 1036–1039.

Zhai, Jingxi, "Breaking the Silent Treatment: The Contractual Enforceability of Non-Disclosure Agreements for Workplace Sexual Harassment Settlements," *Colum. Bus. L. Rev.*, 2020, p. 396.

A Appendix: Additional Figures and Tables

Table A1: Description of Variables used in this Study

Variable	Definition
Firm Level Variables	
Workplace Harassment	A workplace harassment score is constructed using anonymous reviews posted on Glassdoor by applying text analytics techniques. Available for 2009-2022. Winsorized at the 1% level.
Pre-Harassment	Mean of "Workplace Harassment" for 2011-2013 period. Winsorized at the 1% level.
Pre-Harassment High	A dummy variable indicating that the firm ranked above the median of the "Pre-Harassment" score.
Log(# of Reviews)	Log of total number of reviews posted on Glassdoor. Winsorized at the 1% level.
Pre-Log(# of Reviews)	Log of the total number of reviews posted on Glassdoor during the period 2011 - 2013. Winsorized at the 1% level.
Glassdoor Participation	Log of one plus total number of reviews posted on Glassdoor divided by number of employees (Compustat data item EMP). Winsorized at the 1% level.
Harassment Lawsuit	An indicator variable indicating whether a harassment lawsuit was filed against the firm. These are lawsuits with nature of suit codes 442 and 445, as reported by the Federal Judicial Center's (FJC) Integrated Database.
Best Workplace	A dummy variable equals one if the firm is among the 100 Best Companies to Work For in America published by the Great Place To Work Institute, otherwise zero (Edmans, 2011).
S Score	Score of a firm's social performance from Refinitiv ESG.
G Score	Score of a firm's corporate governance performance from Refinitiv ESG.
S&G Score	Average of S and G Scores.
ESG Score	Score of a firm's ESG performance from Refinitiv ESG.

Log(Value of Patent per Employee)	Log of total value of patents as of filing year divided by number of employees (Compustat data item EMP) Source: Kogan et al. (2017)
Log(# of Patents per Employee)	Log of total number of patents as of filing year divided by number of employees (Compustat data item EMP) Source: Kogan et al. (2017)
Log(# of Highly Cited Patents per Employee)	Log of total number of highly cited patents divided by (Compustat data item EMP). Patent is highly cited if its citations rank above the median among other patents in the same filing cohort and technology class. Source: Kogan et al. (2017)
Cash/Assets	Cash and short-term investments (Compustat data item CHE) divided by total assets (Compustat data item AT). Winsorized at the 1% level. Compustat
Debt/Assets	Sum of the book value of long-term debt (Compustat data item DLTT) and the book value of current liabilities (Compustat data item DLC) divided by total assets (Compustat data item AT). Winsorized at the 1% level. Compustat
PP&E/Assets	Property, plant, and equipment (Compustat data item PPENT) divided by total assets (Compustat data item AT). Winsorized at the 1% level. Compustat
CAPEX/Assets	Capital expenditures (Compustat data item CAPX) divided by total assets (Compustat data item AT). Winsorized at the 1% level. Compustat
R&D/Assets	R&D expenditures (Compustat data item XRD) divided by total assets (Compustat data item AT). Missing values set to zero. Winsorized at the 1% level Compustat
Market Share (Sales)	Sales (Compustat data item SALE) divided by total sales in the focal firm's industry classification. Industry classification is based on Fama-French 17 industries. Winsorized at the 1% level. Compustat
RoA	Net Income divided (Compustat data item NI) by total assets (Compustat data item AT). Winsorized at the 1% level. Compustat
Tobin's Q	Sum of the book value of long-term debt (Compustat data item DLTT), the book value of current liabilities (Compustat data item DLC), and common shares outstanding (Compustat data item CSHO) multipled by price per share (Compustat data item PRCC_F) minus total current assets (Compustat data item ACT) divided by total assets (Compustat data item AT). Winsorized at the 1% level. Compustat

Log(# of Employees)	Log of number of employees (Compustat data item EMP).
Log(Sales per Employee)	Log of one plus sales (Compustat data item SALE) divided by employees (Compustat data item EMP).
Profit Growth	Growth of firm profit, calculated as $\frac{profit_t}{profit_{t-1}} - 1$, t indexing year. Profit is sales (Compustat data item SALE) minus cost of goods sold (Compustat data item COGS).
WLR - Advanced Degree	WLR stands for Workforce Loss Rate. Outflow of employees with a master's degree or higher, minus inflow of employees with a master's degree or higher, divided by the sum of total inflow and outflow.
WLR - Forbes 100 College	Outflow of employees with a degree from a college ranked in the top 100 of the Forbes list of American colleges, minus inflow of employees with a degree from a college ranked in the top 100 of the Forbes list of American colleges, divided by the sum of total inflow and outflow.
WLR - Industry Experience	Outflows of highly experienced workers minus inflows of highly experienced workers divided by the sum of total inflows and outflows. Industry experience is defined based on the cumulative experience (in years) of the exiting (entering) worker in the focal firm's industry. Industry experience is high if the experience of the exiting (entering) worker is above the median of the experience of other exiting (entering) workers at the time of exit (entry). Industry classification is based on Fama-French 17 industries.
Racial Diversity	Racial diversity of the firm's inventor team is calculated using Simpson's Diversity Index method. There are four race groups: Asian, Black, Hispanic, and White as extracted based on the inventor's last name using ethnicolr algorithm. The index is calculated as follows: $1 - \sum_{race}((\frac{number\ of\ inventors\ with\ \{race\}\ group}{number\ of\ inventors})^2)$
Ethnic Diversity	Ethnic diversity of the firm's inventor team is calculated using Simpson's Diversity Index method. There are three ethnic groups: African, Asian and European as extracted based on inventor's last name using ethnicolr algorithm. The index is calculated as follows: $1 - \sum_{ethnic}((\frac{number\ of\ inventors\ with\ \{ethnic\}\ group}{number\ of\ inventors})^2)$

Inventor Level Variables

Log(# of High Valued Patents)	Log of one plus the number of highly valued patents. A patent is highly valued if its economic value is above the median of all other patents in the same filing year and technology class. Technology class is defined based on the CPC classification.
Log(# of Patents)	Log of one plus the number of patents filed.
Log(# of Highly Cited Patents)	Log of one plus the number of highly cited patents. A patent is highly cited if it ranks above the median of all (i.e., forward) citations it has received among all other patents in the same filing year and technology class. Technology class is defined based on the CPC classification.

A.1 Additional Validation Tests for the Workplace Harassment Measure

	Harassment Lawsuit								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Workplace Harassment	0.445***	0.347***	0.465***	0.458***	0.431***	0.397***	0.427***	0.355***	0.202***
	(0.093)	(0.088)	(0.095)	(0.100)	(0.097)	(0.101)	(0.093)	(0.098)	(0.075)
Cash/Assets	-0.038**	-0.040**	-0.023	-0.029	-0.034*	-0.038**	-0.035**	-0.022	0.022
	(0.016)	(0.016)	(0.018)	(0.018)	(0.018)	(0.018)	(0.016)	(0.017)	(0.028)
Debt/Assets	-0.028*	-0.024	-0.028*	-0.005	-0.034**	-0.024	-0.024	-0.004	-0.018
	(0.015)	(0.015)	(0.016)	(0.015)	(0.015)	(0.016)	(0.015)	(0.016)	(0.020)
PP&E/Assets	0.041	0.050	0.042	0.011	0.048*	0.067*	0.043	0.040	0.049
	(0.026)	(0.032)	(0.028)	(0.032)	(0.027)	(0.034)	(0.027)	(0.034)	(0.060)
CAPEX/Assets	-0.113	-0.166*	-0.135	-0.123	-0.144	-0.228**	-0.130	-0.174	-0.086
	(0.088)	(0.097)	(0.095)	(0.114)	(0.095)	(0.112)	(0.093)	(0.110)	(0.093)
R&D/Assets	-0.035	-0.017	0.005	0.016	-0.024	-0.011	-0.027	-0.015	0.047
	(0.027)	(0.030)	(0.028)	(0.032)	(0.029)	(0.031)	(0.029)	(0.031)	(0.056)
Market Share (Sales)	1.121*	1.701**	1.231**	1.564**	1.305**	1.887**	1.202**	2.401***	2.474*
	(0.598)	(0.731)	(0.606)	(0.772)	(0.637)	(0.911)	(0.597)	(0.849)	(1.360)
RoA	-0.044***	-0.044***	-0.039**	-0.039*	-0.048***	-0.046***	-0.047***	-0.054***	0.005
	(0.016)	(0.017)	(0.017)	(0.020)	(0.017)	(0.018)	(0.017)	(0.018)	(0.016)
Tobin's Q	0.003	0.003	0.002	-0.000	0.002	0.001	0.003	0.001	0.000
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.003)	(0.002)
Log(# of Employees)	0.050***	0.044***	0.051***	0.043***	0.052***	0.046***	0.050***	0.041***	0.015
	(0.005)	(0.006)	(0.005)	(0.006)	(0.006)	(0.007)	(0.005)	(0.006)	(0.011)
Glassdoor Participation	0.024***	0.020***	0.025***	0.017***	0.025***	0.021***	0.024***	0.016***	0.002
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.006)
Constant	-0.203***	-0.161***	-0.212***	-0.173***	-0.202***	-0.179***	-0.199***	-0.153***	-0.050
	(0.028)	(0.025)	(0.029)	(0.028)	(0.029)	(0.029)	(0.028)	(0.027)	(0.039)
Year FE	Yes	No	No	No	No	No	No	No	Yes
Industry × Year FE	No	Yes	No	No	No	No	No	No	No
State × Year FE	No	No	Yes	No	No	No	No	No	No
State × Industry × Year FE	No	No	No	Yes	No	No	No	No	No
State Inc. × Year FE	No	No	No	No	Yes	No	No	No	No
State Inc. × Industry × Year FE	No	No	No	No	No	Yes	No	No	No
State Emp. × Year FE	No	No	No	No	No	No	Yes	No	No
State Emp. × Industry × Year FE	No	No	No	No	No	No	No	Yes	No
Firm FE	No	No	No	No	No	No	No	No	Yes
Obs.	13269	13263	12823	11396	12073	11134	13155	11775	12890
R^2	0.090	0.111	0.125	0.284	0.127	0.186	0.120	0.254	0.409

Table A2: Validating Workplace Harassment Measure: Predicting Workplace Harassment Lawsuits - Full Table. This table reports the results of OLS regressions of future (at year $t+1$) workplace harassment lawsuits on the workplace harassment measure specified in Equation 2. Controls include are Cash/Assets, Debt/Assets, PP&E/Assets, CAPEX/Assets, R&D/Assets, Market Share (Sales), RoA, Tobin's Q, Log(# of Employees), Glassdoor Participation. "State", "State Inc.", and "State Emp." are the states where the firm is headquartered, where the firm is incorporated and from where the majority of employees post reviews in year t, respectively. Industry fixed effects are Fama-French 17 industries. Variable definitions are provided at Table A1. In parentheses below the point estimates are standard errors clustered at the firm level.

	Harassment Lawsuit							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Workplace Harassment	12.998***	10.201***	13.767***	14.837***	11.353***	9.963***	11.688***	9.384***
	(3.167)	(3.025)	(2.874)	(2.894)	(3.324)	(3.669)	(3.296)	(3.128)
Cash/Assets	-0.597	-0.870	-0.129	-0.790	-0.128	-0.474	-0.470	-0.223
	(0.775)	(0.708)	(0.739)	(0.642)	(0.857)	(0.759)	(0.720)	(0.630)
Debt/Assets	-0.448	-0.456	-0.560	0.099	-0.642	-0.622	-0.357	-0.150
	(0.476)	(0.465)	(0.437)	(0.436)	(0.510)	(0.522)	(0.465)	(0.435)
PP&E/Assets	0.506	0.671	-0.080	-1.110*	0.791	1.308*	0.454	0.841
	(0.556)	(0.624)	(0.608)	(0.603)	(0.619)	(0.697)	(0.510)	(0.565)
CAPEX/Assets	-1.584	-3.174	-0.307	-0.886	-1.763	-3.874	-2.579	-5.588*
	(2.461)	(2.464)	(3.225)	(3.011)	(3.084)	(2.632)	(2.475)	(3.042)
R&D/Assets	-8.924***	-10.932***	-4.309**	-4.683**	-9.648***	-13.385***	-7.570**	-9.028***
	(3.398)	(4.109)	(2.029)	(2.215)	(3.626)	(4.445)	(2.987)	(3.110)
Market Share (Sales)	12.983*	18.577**	8.582*	11.953	16.344**	23.138***	15.873**	20.211***
	(6.807)	(9.255)	(4.954)	(8.003)	(6.860)	(7.318)	(6.572)	(6.401)
RoA	-1.386*	-1.348	-1.532**	-1.220*	-1.238	-1.382	-1.550**	-1.679**
	(0.819)	(0.857)	(0.774)	(0.722)	(0.849)	(0.876)	(0.775)	(0.683)
Tobin's Q	0.172*	0.219**	0.128	0.112	0.129	0.180*	0.190**	0.170*
	(0.099)	(0.095)	(0.082)	(0.100)	(0.111)	(0.100)	(0.089)	(0.100)
Log(# of Employees)	0.932***	0.810***	0.972***	0.928***	0.940***	0.799***	0.894***	0.800***
	(0.061)	(0.097)	(0.072)	(0.081)	(0.063)	(0.085)	(0.061)	(0.073)
Glassdoor Participation	0.229*	0.234*	0.290**	0.256*	0.280**	0.297**	0.212*	0.166
	(0.121)	(0.141)	(0.124)	(0.134)	(0.123)	(0.144)	(0.122)	(0.129)
Constant	-8.561***	-7.339***	-8.404***	-7.571***	-8.138***	-7.178***	-7.856***	-6.375***
	(0.850)	(0.958)	(0.833)	(0.903)	(0.839)	(1.078)	(0.885)	(0.856)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	No	No	No	No	No	No	No
Industry × Year FE	No	Yes	No	No	No	No	No	No
State (Hq.) × Year FE	No	No	Yes	No	No	No	No	No
State (Hq.) × Industry × Year FE	No	No	No	Yes	No	No	No	No
State (Inc.) × Year FE	No	No	No	No	Yes	No	No	No
State (Inc.) × Industry × Year FE	No	No	No	No	No	Yes	No	No
State (Emp.) × Year FE	No	No	No	No	No	No	Yes	No
State (Emp.) × Industry × Year FE	No	No	No	No	No	No	No	Yes
Obs.	13269	12248	10574	4835	10572	7622	10701	5418
Pseudo-R^2	0.296	0.334	0.351	0.457	0.335	0.386	0.333	0.411

Table A3: Validating Workplace Harassment Measure: Predicting Workplace Harassment Lawsuits - Poisson Regression. This table reports the results of Poisson regressions of future (at year $t + 1$) workplace harassment lawsuits on the measure of workplace harassment specified in equation 2. Harassment lawsuit is a count variable that indicates the number of harassment-related lawsuits firms face in year $t + 1$, if none, it is set to zero. Controls included are same as in Table A2. Variable definitions are provided in Table A1. In parentheses below the point estimates are standard errors clustered at the firm level.

	Penalty/Assets			
	(1)	(2)	(3)	(4)
Workplace Harassment	16.467***	15.980***	14.625***	18.257***
	(5.300)	(6.040)	(3.864)	(5.211)
Cash to Assets	-1.067	-1.186	-1.907*	-2.917**
	(1.256)	(1.314)	(1.055)	(1.144)
Debt to Assets	0.050	0.042	-0.278	0.393
	(0.428)	(0.488)	(0.476)	(0.651)
PPE to Assets	0.914	1.051	0.918	0.544
	(0.624)	(0.736)	(0.620)	(0.803)
CAPEX to Assets	-1.738	-3.110	-2.494	-8.059
	(3.160)	(3.290)	(3.157)	(5.117)
R&D to Assets	-4.990*	-6.285*	-3.766	-0.678
	(3.030)	(3.724)	(3.166)	(3.141)
Market Share (Sales)	6.372	-3.818	6.660	-9.656*
	(5.322)	(7.312)	(5.262)	(5.523)
ROA	2.273	1.898	2.465	1.311
	(1.847)	(1.752)	(1.688)	(1.312)
Tobin's Q	-0.195*	-0.145	-0.200**	-0.066
	(0.117)	(0.112)	(0.098)	(0.100)
Log(# of Employees)	0.866***	1.001***	0.830***	1.140***
	(0.079)	(0.104)	(0.078)	(0.119)
Glassdoor Participation	0.176*	0.336***	0.196*	0.396**
	(0.100)	(0.118)	(0.113)	(0.157)
Constant	-8.050***	-8.120***	-6.764***	-7.845***
	(1.419)	(1.653)	(1.061)	(1.580)
Year FE	Yes	No	No	No
Industry ×Year FE	No	Yes	No	No
State ×Year FE	No	No	Yes	No
State ×Industry ×Year FE	No	No	No	Yes
Obs.	15,373	13,062	10,075	3,871
Pseudo-R^2	0.373	0.414	0.444	0.557

Table A4: Validating Workplace Harassment Measure: Predicting Workplace Harassment Lawsuit Penalties. This table reports the results of Poisson regressions of future (at year $t + 1$) harassment-related lawsuit penalties on the workplace harassment measure specified in Equation 2. Penalty/Assets is the total penalties paid by the firm due to harassment-related lawsuits in year $t + 1$, normalized by total assets in year t. The included controls are the same as in Table A2. Variable definitions are provided in Table A1. In parentheses below the point estimates are standard errors clustered at the firm level.

	Best Workplace	S Score	G Score	S&G Score	ESG Score
	(1)	(2)	(3)	(4)	(5)
Workplace Harassment	-0.236***	-64.090***	-38.651***	-51.371***	-59.135***
	(0.067)	(10.710)	(12.172)	(9.258)	(9.385)
Cash/Assets	0.024	1.179	-4.622*	-1.722	-1.980
	(0.024)	(2.537)	(2.740)	(2.080)	(2.188)
Debt/Assets	-0.000	-0.950	-2.034	-1.492	-1.526
	(0.014)	(1.635)	(1.921)	(1.426)	(1.459)
PP&E/Assets	-0.007	3.004	5.030	4.017	5.852**
	(0.018)	(3.297)	(3.325)	(2.777)	(2.909)
CAPEX/Assets	0.111	-27.908**	-27.311**	-27.610***	-34.706***
	(0.086)	(12.285)	(13.767)	(10.662)	(10.838)
R&D/Assets	0.120**	46.423***	20.077***	33.250***	36.038***
	(0.056)	(6.548)	(6.916)	(5.321)	(5.404)
Market Share (Sales)	-0.898***	84.236*	93.655**	88.945***	146.565***
	(0.340)	(43.861)	(42.235)	(31.472)	(33.730)
RoA	0.053***	13.993***	19.407***	16.700***	17.222***
	(0.019)	(2.728)	(2.906)	(2.285)	(2.275)
Tobin's Q	0.004	0.390*	-0.609***	-0.109	-0.009
	(0.002)	(0.203)	(0.212)	(0.163)	(0.168)
Log(# of Employees)	0.032***	10.276***	5.354***	7.815***	9.261***
	(0.006)	(0.454)	(0.496)	(0.379)	(0.394)
Glassdoor Participation	0.025***	3.198***	0.013	1.605***	2.399***
	(0.005)	(0.471)	(0.560)	(0.410)	(0.421)
Constant	-0.054**	28.417***	47.783***	38.100***	29.690***
	(0.024)	(3.097)	(3.551)	(2.695)	(2.715)
Industry x Year FE	Yes	Yes	Yes	Yes	Yes
Obs.	14,989	12,311	12,311	12,311	12,311
R^2	0.061	0.385	0.190	0.358	0.443

Table A5: Validating Workplace Harassment Measure: Predicting Best Workplaces and ESG Scores. This table reports the results of OLS regressions of future (in year $t+1$) Best Workplace indicator and firm ESG scores on the workplace harassment measure specified in Equation 2. Best Workplace is a dummy variable that equals one if the firm is among the 100 Best Companies to Work For in America published by the Great Place To Work Institute (Edmans, 2011), otherwise zero. S, G, S&G (average of S and G scores) and ESG scores are from Refinitiv. The included controls are the same as in Table A2. Variable definitions are provided in Table A1. In parentheses below the point estimates are standard errors clustered at the firm level.

	Harassment Lawsuit											
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Workplace Harassment	0.347***	0.346***	0.356***	0.368***	0.359***	0.358***	0.360***	0.370***	0.532***	0.545***	0.547***	0.543***
	(0.088)	(0.088)	(0.095)	(0.096)	(0.095)	(0.095)	(0.097)	(0.096)	(0.129)	(0.134)	(0.132)	(0.130)
Best Workplace		-0.002										
		(0.030)										
Li et al. (2021) Firm Culture Variables												
Integrity			0.007***									
			(0.003)									
Teamwork				0.000								
				(0.002)								
Innovation					-0.001							
					(0.001)							
Respect						0.003**						
						(0.002)						
Quality							-0.003					
							(0.002)					
Culture								0.001				
								(0.003)				
Refinitiv ESG Scores												
S Score									-0.000			
									(0.000)			
G Score										0.000		
										(0.000)		
S&G Score											0.000	
											(0.000)	
ESG Score												0.000
												(0.000)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	13263	13263	11940	11940	11940	11940	11940	11940	8907	8907	8907	8907
R^2	0.111	0.111	0.110	0.109	0.110	0.110	0.110	0.109	0.129	0.129	0.129	0.129

Table A6: Validating Workplace Harassment Measure: Predicting Workplace Harassment Lawsuits while Controlling for Related Variables. This table reports the results of OLS regressions of future (in year $t+1$) workplace harassment lawsuits on the measure of workplace harassment specified in Equation 2. The included controls are the same as in Table A2. In addition to the baseline controls, I also include controls related to corporate culture, including the Best Workplace indicator, corporate culture measures from Li et al. (2021), and ESG scores. Variable definitions are provided in table A1. In parentheses below the point estimates are standard errors clustered at the firm level.

	Lawsuit			
Offense topic:	Competition	Financial	Enviromental	Consumer
Workplace Harassment	-0.034	-0.026	0.064	-0.008
	(0.036)	(0.028)	(0.095)	(0.057)
Cash/Assets	-0.009	-0.004	-0.018	0.028*
	(0.007)	(0.006)	(0.015)	(0.016)
Debt/Assets	-0.003	-0.007*	-0.009	-0.001
	(0.004)	(0.004)	(0.014)	(0.010)
PP&E/Assets	-0.022**	-0.017***	0.192***	0.001
	(0.009)	(0.006)	(0.039)	(0.020)
CAPEX/Assets	0.016	0.016	-0.248*	-0.031
	(0.029)	(0.028)	(0.129)	(0.070)
R&D/Assets	-0.009	-0.026**	-0.020	-0.003
	(0.015)	(0.011)	(0.032)	(0.021)
Market Share (Sales)	1.141***	0.344**	6.157***	1.969***
	(0.217)	(0.168)	(0.704)	(0.482)
RoA	0.002	-0.019***	-0.029	-0.014
	(0.006)	(0.007)	(0.018)	(0.009)
Tobin's Q	-0.001***	0.001	0.001	-0.002**
	(0.000)	(0.001)	(0.001)	(0.001)
Log(# of Employees)	0.009***	0.005***	0.034***	0.025***
	(0.002)	(0.001)	(0.005)	(0.004)
Glassdoor Participation	0.003*	0.000	0.009*	0.016***
	(0.002)	(0.001)	(0.005)	(0.004)
Constant	0.000	0.009	-0.057**	-0.063***
	(0.010)	(0.008)	(0.029)	(0.021)
Industry $\times$ Year FE	Yes	Yes	Yes	Yes
Obs.	16,449	16,449	16,449	16,449
R^2	0.097	0.054	0.345	0.192

Table A7: Validating Workplace Harassment Measure: Placebo Tests. This table reports the results of OLS regressions of future (at year $t+1$) lawsuits due to corporate misconduct in the areas of competition, finance, environment, and consumer protection on the measure of workplace harassment specified in Equation 2. Data on these lawsuits are from the Good Jobs First Violation Tracker database. The included controls are the same as in table A2. Variable definitions are provided in Table A1. In parentheses below the point estimates are standard errors clustered at the firm level.

A.2 Robustness Tests for Baseline Results

	Log(economic value of patents per employee)	Log(# of patents per employee)	Log(# of highly cited patents per employee)
Post L-NDA × Pre-High Harass	0.314***	0.188***	0.160***
	(0.060)	(0.041)	(0.032)
Post L-NDA × Pre-Log(# of Reviews)	-0.071***	-0.018**	-0.025***
	(0.013)	(0.007)	(0.007)
Firm FE	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes
Obs.	14,078	14,078	14,078
R^2	0.862	0.837	0.742

Table A8: The Impact of Workplace Harassment on Innovation: Robustness - Pre-Harassment as Dummy. This table presents results for the triple difference specification as specified in equation 3, where the outcome variables are *Log(economic value of patents per employee)*, *Log(# of patents per employee)*, and *Log(# of highly cited patents per employee)*. Data on economic value of patents, number of patents, and citation of patents are from Kogan et al. (2017). *Pre-High Harass* is a dummy variable indicating that the firm is ranked above the median of the "Pre-Harassment" score. Variable definitions are provided in Table A1. In parentheses are standard errors that are double-clustered at the firm and the state headquarter level.

	Log(economic value of patents per employee)	Log(# of patents per employee)	Log(# of highly cited patents per employee)
	(1)	(2)	(3)
Post L-NDA × Pre-Harassment	4.448***	2.528***	2.233***
	(0.657)	(0.526)	(0.644)
Post L-NDA × Pre-Log(# of Reviews)	-0.078***	-0.021**	-0.029***
	(0.013)	(0.008)	(0.008)
Firm FE	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes
Obs.	11,074	11,074	11,074
R^2	0.877	0.850	0.759

Table A9: The Impact of Workplace Harassment on Innovation: Robustness - Innovation Intensive Firms. This table presents results for the triple difference specification as specified in equation 3, where the outcome variables are *Log(economic value of patents per employee)*, *Log(# of patents per employee)*, and *Log(# of highly cited patents per employee)*. Data on economic value of patents, number of patents, and citation of patents are from Kogan et al. (2017). The sample is restricted to innovation-intensive firms. Innovation-intensive firms are those above the median number of patents in the cross-section of firms in the base sample. Variable definitions are provided in Table A1. In parentheses are standard errors that are double-clustered at the firm and the state headquarter level.

	Log(economic value of patents to assets)	Log(# of top 10 cited patents per employee)	Log(# of top 25 cited patents per employee)
Post L-NDA × Pre-Harassment	0.190***	2.052***	3.282***
	(0.026)	(0.538)	(0.746)
Post L-NDA × Pre-Log(# of Reviews)	-0.005***	-0.140***	-0.189***
	(0.001)	(0.020)	(0.023)
Firm FE	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes
Obs.	14,078	14,078	14,078
R^2	0.759	0.785	0.791

Table A10: The Impact of Workplace Harassment on Innovation: Robustness - Alternative Measures. This table presents results for the triple difference specification as specified in equation 3, where the outcome variables are *Log(economic value of patents to assets)*, *Log(# of top 10 cited patents per employee)*, and *Log(# of top 25 cited patents per employee)*. Data on economic value of patents, number of patents, and citation of patents are from Kogan et al. (2017). Top 10 (25) cited patents are those that rank in the top decile (quantile) in terms of total citations among all other patents in the same filing year and technology class. The technology class is based on the CPC classification. Variable definitions are provided in Table A1. In parentheses are standard errors that are double-clustered at the firm and the state headquarter level.

	Log(economic value of patents per employee)		Log(# of patents per employee)		Log(# of highly cited patents per employee)	
	(1)	(2)	(3)	(4)	(5)	(6)
Post L-NDA × Pre-Harassment	4.382***	4.927***	2.593***	3.117***	2.513***	3.036***
	(0.649)	(0.655)	(0.486)	(0.508)	(0.257)	(0.302)
Post L-NDA × Pre-Log(# of Reviews)	-0.051***	-0.051***	-0.002	0.003	-0.005	-0.000
	(0.014)	(0.016)	(0.015)	(0.019)	(0.017)	(0.020)
Firm FE	Yes	No	Yes	No	Yes	No
State Emp. × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
State Emp. × Firm FE	No	Yes	No	Yes	No	Yes
Obs.	14193	12993	14193	12993	14193	12993
R^2	0.859	0.866	0.831	0.841	0.731	0.743

Table A11: The Impact of Workplace Harassment on Innovation: Robustness - Employees' State. This table presents results for the triple difference specification as specified in equation 3, where the outcome variables are *Log(economic value of patents per employee)*, *Log(# of patents per employee)*, and *Log(# of highly cited patents per employee)*. Data on economic value of patents, number of patents, and citation of patents are from Kogan et al. (2017). "State Emp." is the state of the employee. It is defined based on the state names provided by the majority of reviewers when they post a review on Glassdoor about company i in year t. Firm is treated if the "State Emp." passes NDA-limiting regulations. Variable definitions are provided in Table A1. In parentheses are standard errors that are double-clustered at the firm and the state headquarter level.

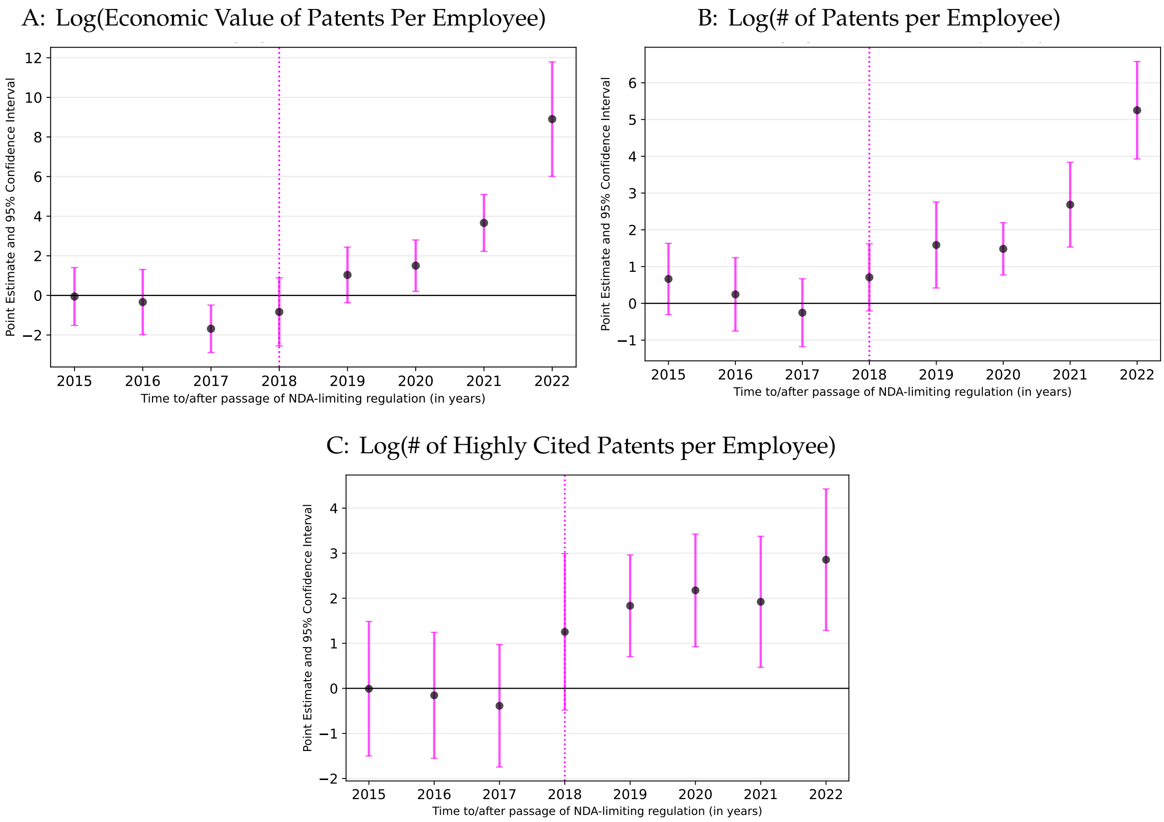

Figure A1: The Impact of Workplace Harassment on Innovation Output: Dynamic Effects - Treated States. This figure plots the estimated coefficients of the interaction term in Equation 4, where $Post\ L\text{-}NDA_{s,t}$ is interacted with a series of year indicators before ($[t-4, t]$) and after ($[t+1, t+4]$) the NDA limitation regulations. Panels A, B, and C show the results where the outcome variable is Log(Economic Value of Patents Per Employee), Log(# of Patents per Employee), and Log(# of Highly Cited Patents per Employee), respectively. Data on economic value of patents, number of patents, and citation of patents are from Kogan et al. (2017). The sample is limited to the thirteen states that pass the NDA-limiting regulation shown in Table A12. Variable definitions are provided in Table A1. The standard errors are double-clustered at the firm and state headquarter level.

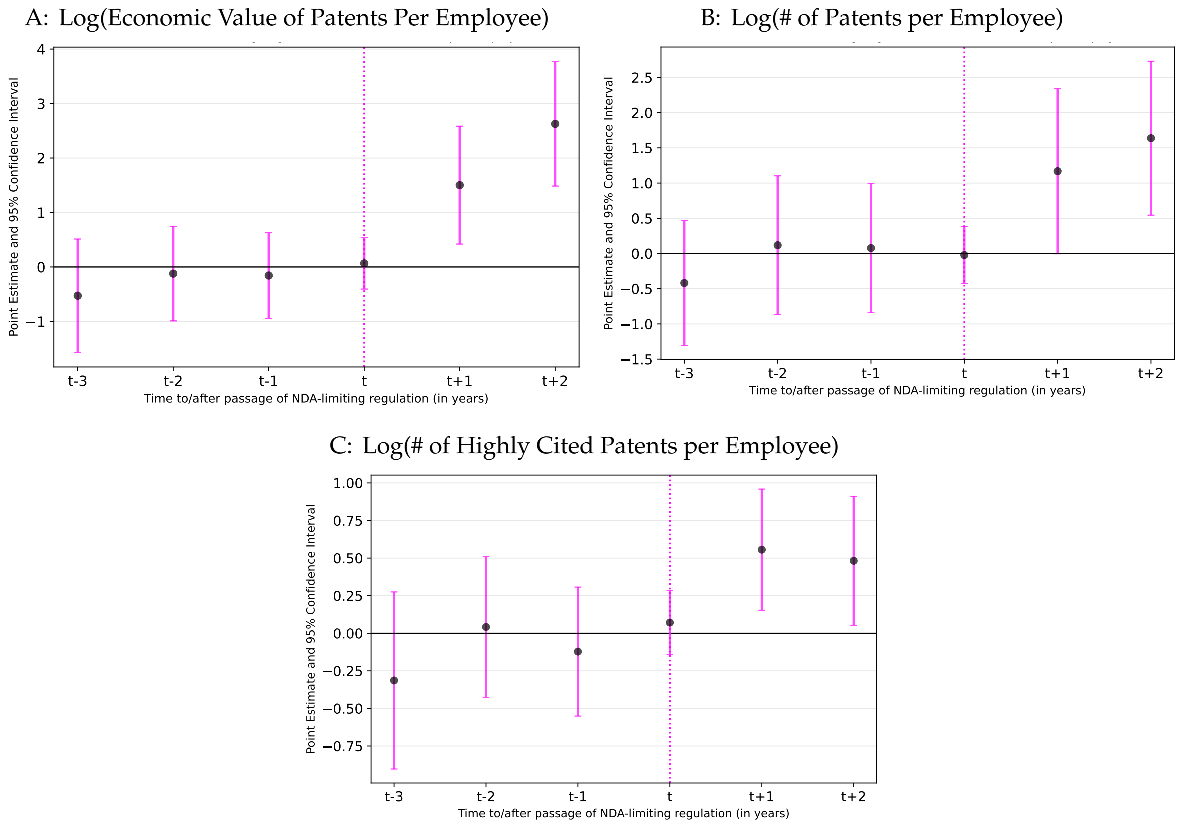

Figure A2: The Impact of Workplace Harassment on Innovation Output: Dynamic Effects - Stacked Difference-in-Difference. This figure plots the estimated coefficients of the interaction term in Equation 4, where $Post\ L\text{-}NDA_{s,t}$ is interacted with a series of year indicators before ($[t-4, t]$) and after ($[t+1, t+4]$) the NDA limitation regulations in stacked difference-in-difference estimation setting (Cengiz et al., 2019). Panels A, B, and C show the results where the outcome variable is Log(Economic Value of Patents Per Employee), Log(# of Patents per Employee), and Log(# of Highly Cited Patents per Employee), respectively. Data on economic value of patents, number of patents, and citation of patents are from Kogan et al. (2017). The sample is limited to the thirteen states that pass the NDA-limiting regulation shown in Table A12. Variable definitions are provided in Table A1. The standard errors are double-clustered at the firm and state headquarter level.

B Appendix on Further Details on Data and Variables

A: Company Reviews Page B: Employee Reviews

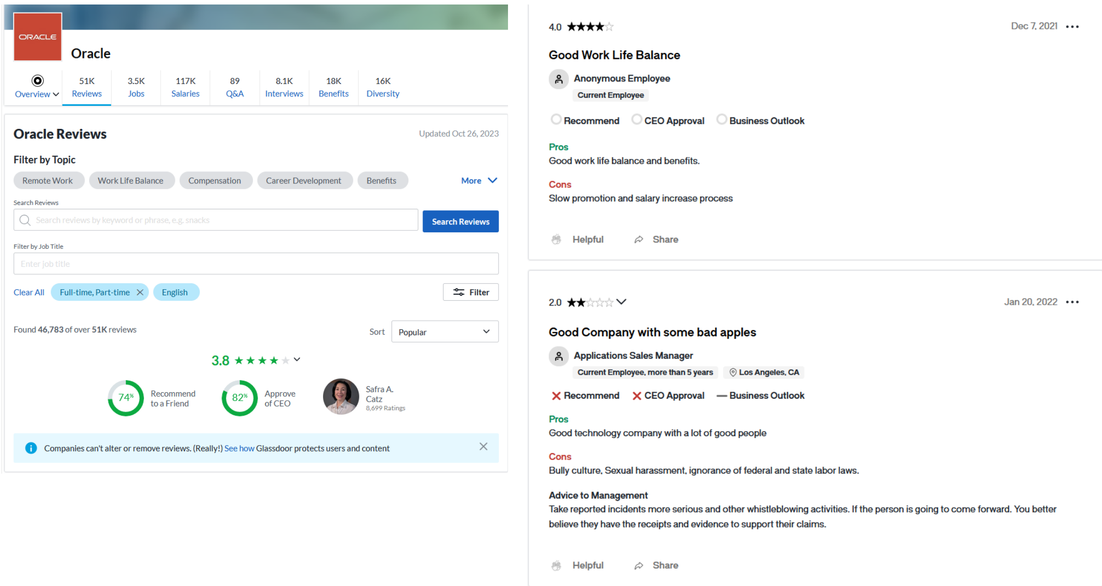

Figure A3: An Example of Company Reviews on Glassdoor

State name	Treatment start year	Effective year-month	Source
Arizona	2019	201804	Link
California	2019	201901	Link
Illinois	2020	202001	Link
Maryland	2019	201810	Link
Nevada	2020	201907	Link
New Jersey	2020	201903	Link
New Mexico	2021	202005	Link
New York	2019	201807	Link
Oregon	2021	202010	Link
Tennessee	2019	201805	Link
Vermont	2019	201807	Link
Virginia	2020	201907	Link
Washington	2019	201806	Link

Table A12: Legislation Limiting Non-Disclosure Agreements (L-NDAs). This table provides information on when and where regulations were enacted to prohibit the use of non-disclosure agreements in pre-employment contracts to silence employees about workplace harassment problems. "Treatment start year" refers to the year when the $Post_{s,t}$ dummy in the Equation 3 starts taking a value of one.

B.1 Preprocessing Employee Review Text Corpus

To ensure that the employee review texts are in a consistent and structured format before fitting the word2vec model, I perform pre-processing steps on the review texts. The following steps are performed:

- **Sentence Identification**: To train the word2vec model, I divide the review texts into sentences. Before identifying sentences, I remove all multiple whitespaces and empty new line elements from the review text corpus. The sentence tokenization function of the NLTK package is then used to identify individual sentences within the text.
- **Text Cleaning**: I take the following steps to clean the sentences and prepare them for further processing:
  - *Removal of Non-Word Characters*: All non-word characters, such as special symbols or emoticons, are removed from each sentence.
  - *Stop-word Removal*: Common stop words such as "the," "and," or "is" are removed from sentences because they do not provide meaningful information.

- *Digit and Punctuation Removal*: Words containing digits and punctuation are removed from sentences.
- *Lemmatization*: Words are lemmatized, meaning they are converted to their base or dictionary form, to ensure consistency in how words appear.

- **Identifying Phrases: *Unigrams and Bigrams***: Mikolov et al. (2013a) show that having phrases (i.e., 2 or more word combinations) improve the model performance in better learning the context than simply using single words (unigrams). To be able to identify the phrases (2 word combinations or bigrams), I use Phrases model.

B.2 Word2vec Model Parameters

I follow Mikolov et al. (2013b,a) and Li et al. (2021) in choosing the parameters for fitting the word2vec model, taking into account the specific characteristics of employee reviews. The parameter choices are as follows:

- *Model*: I use the Skip-Gram model, which is effective at identifying unique words used in similar contexts, making it suitable for capturing harassment-related terms, including those with common misspellings (e.g., "harass" is often spelled "harrass" in reviews). Importantly, this model takes word order into account, giving less weight to words further away from the target word, unlike the alternative model, Continuous Bag of Words (CBOW), where word order does not matter.
- *Algorithm*: I choose the Hierarchical Softmax algorithm instead of the Negative Sampling algorithm. This decision is based on the fact that my study focuses on harassment-specific texts, which are less frequent than, for example, salary-related reviews. As suggested by Mikolov et al. (2013a), the Hierarchical Softmax algorithm is more suitable for modeling less frequent words.
- *Window size*: The window size determines how many words to the left or right of the target word the algorithm should consider when understanding its context. While increasing the window size can improve the accuracy of the model, words further away

from the target word are less likely to be directly related to it. Following the literature, I set it to 5.

- *Minimum appearance*: This parameter helps to eliminate words that appear only a few times in the entire review corpus and do not contribute significantly to the model's performance. Following Mikolov et al. (2013a), I set the minimum appearance threshold to 5.

- *Word vector size*: The size of the word vectors (i.e., word embedding) captures the meaning of the words, and its optimal value depends on the type of text corpus used. While Mikolov et al. (2013a) used a vector size of 300 in their training on the large Google News database, I chose a vector size of 100 for my study, given the limited range of topics covered in employee reviews.

B.3 Identifying Skilled Employees on LinkedIn

To construct proxies for skill, I rely on employees' educational attainment and accumulated experience. LinkedIn public profile data from the BrightData Initiative allows me to identify employees with different levels of human capital. In particular, I can observe in this data where and when employees have worked or studied. LinkedIn has standardized IDs for universities and degrees. This allows me to exploit both dimensions and construct 2 variables based on educational attainment.

(i) *Workers with masters degree or above*: To identify an employee who has an advanced degree, I check whether any of the following degree titles appear on the degrees obtained using regular experession methods: (Note that I remove punctuation and lowercase text information for a better match):

'master', 'ms', msc', 'msce', 'ma', 'meng', 'med','mba', 'graduate', 'phd', 'mphil', 'md', 'doctor', 'dr', 'pharmd', 'msis'

If educational information does not contain any degree information about then I group them into non-masters or below.

(ii) *Workers with Top 100 Forbes College degree*: This is motivated by the recent work of Niessen-Ruenzi and Zimmerer (2021). Since LinkedIn has already standardized the names of the universities, I use regex methods and search for the matches of the 100 best schools in the educational attainment section. The Table A13 below shows the list of universities.

Rank	College	Rank	College
1	Massachusetts Institute of Technology	51	Wesleyan University
2	Stanford University	52	Texas A&M University, College Station
3	University of California, Berkeley	53	Hamilton College
4	Princeton University	54	Boston University
5	Columbia University	55	Middlebury College
6	University of California, Los Angeles	56	Santa Clara University
7	Williams College	57	Brigham Young University
8	Yale University	58	Purdue University
9	Duke University	59	Washington and Lee University
10	University of Pennsylvania	60	New York University
11	Northwestern University	61	George Washington University
12	Rice University	62	Trinity College (CT)
13	Vanderbilt University	63	San Diego State University
14	Dartmouth College	64	University of Georgia
15	Harvard University	65	Binghamton University, SUNY
16	Cornell University	66	CUNY, Baruch College
17	University of California, San Diego	67	Florida State University
18	Johns Hopkins University	68	University of Minnesota, Twin Cities
19	Brown University	69	Davidson College
20	University of Chicago	70	North Carolina State University, Raleigh
21	University of Southern California	71	Barnard College
22	Georgetown University	72	University of Richmond
23	University of California, Davis	73	Vassar College
24	Amherst College	74	University of Connecticut
25	University of Michigan, Ann Arbor	75	New Jersey Institute of Technology
26	University of Florida	76	California State University, Fullerton
27	Washington University in St. Louis	77	Lafayette College
28	University of North Carolina, Chapel Hill	78	University of Miami (FL)
29	University of Virginia	79	Northeastern University
30	University of California, Irvine	80	California State University, Long Beach
31	Emory University	81	Michigan State University
32	Tufts University	82	Virginia Tech
33	University of Washington, Seattle	83	Southern Methodist University
34	University of Illinois, Urbana-Champaign	84	University of California, Riverside
35	Georgia Institute of Technology	85	Pomona College
36	University of Notre Dame	86	Stony Brook University, SUNY
37	Wellesley College	87	University of California, Santa Cruz
38	Swarthmore College	88	University of Utah
39	University of California, Santa Barbara	89	Bucknell University
40	University of Maryland, College Park	90	University at Buffalo
41	William & Mary	91	Indiana University, Bloomington
42	Boston College	92	Grinnell College
43	University of Texas, Austin	93	Rutgers University
44	Colgate University	94	Villanova University
45	California Institute of Technology	95	Bryn Mawr College
46	Carnegie Mellon University	96	Colorado College
47	Claremont McKenna College	97	University of Rochester
48	Bowdoin College	98	California Polytechnic State University, San Luis Obispo
49	University of Wisconsin, Madison	99	University of Illinois at Chicago
50	Wake Forest University	100	Loyola Marymount University

Table A13: Forbes America's Top 100 Colleges. This table shows the list of 100 best colleges that are included in Forbes America's Top Colleges 2023. The ranking is obtained in April, 2023, and is available in the following webpage: `https://www.forbes.com/top-colleges/`.