

Externalities from Silence: Non-Disclosure Agreements Distort Firm Reputation*

Jason Sockin
University of Pennsylvania

Aaron Sojourner
University of Minnesota

Evan Starr
University of Maryland

April 25, 2021

Abstract

Workers face information problems when choosing a potential employer, negative information on employers is undersupplied. We examine one potential driver of this shortage: that workers are prohibited from sharing negative information about their employer because they are bound by over-broad non-disclosure agreements (NDAs) that extend beyond trade secrets. Three states recently changed laws to narrow the kinds of information NDAs could suppress. We study how these changes affected the supply of information from workers about firm employment practices on Glassdoor. The flow of negative information about employment practices did increase after the legal changes, especially in industries a greater use of NDAs. These was driven by an increase in 1- and 2-star ratings, and a 5% increase in reviews related to harassment. Employees were also less likely to conceal aspects of their identity, consistent with both reduced concern about retaliation risks and increased value of the information supplied to others. Lastly, in markets where NDAs are more common, the legal changes created more dispersion in firms' ratings, suggesting that broad NDAs facilitate equilibria where firms with worse employment practices seek to 'pool' their public reputations with firms with better practices. These results highlight that, by preventing outsiders from learning about undesirable firm employment practices, over-broad NDAs impose potential negative externalities on job seekers and competitor firms.

JEL: M54, K31, J58

Keywords: Imperfect Information, Non-Disclosure Agreements, Firm Reputation.

*Preliminary.

1 Introduction

When choosing a potential employer, workers face an information problem: How can they ascertain accurate, relevant information about what it will be like to work for each potential employer? Platforms that collect reviews of employers (e.g., Glassdoor) mitigate such information problems, but only insofar as the information supplied is accurate and covers the attributes that are most relevant to workers. Prior research indicates that job seekers especially value negative information about potential employers, but that negative information is undersupplied (Marinescu et al., 2018; Sockin and Sojourner, 2020). In this study, we examine a potential explanation for why negative information about firms is undersupplied: those who possess such negative information are restrained from sharing it by broad non-disclosure agreements (NDAs).

NDAs cover 57% of workers in the US (Balasubramanian et al., 2021), and while they may have legitimate purposes in protecting trade secrets, they can also be written and interpreted very broadly to prohibit the diffusion of general information about an employer.¹ In theory, by threatening workers who speak out with potential costs, NDAs may discourage workers from sharing negative information about the firm. The #MeToo movement has spotlighted prominent examples of how NDAs help perpetuate wrongdoing by silencing victims of sexual harassment (Facchinei, 2020). The employer and employee who sign an NDA are unlikely to internalize these costs fully because (a) NDAs are often signed as a condition of employment before any information is realized, (b) all relevant parties, especially future job seekers, are not involved in the contracting process, and (c) because outside parties cannot price information of which they are not aware.

In so doing, NDAs impose negative externalities on others who value such information, especially future job seekers who would otherwise redirect their labor market search and on employers who prefer high-road talent management strategies but struggle to credibly differ-

¹A close cousin of NDAs are non-disparagement agreements, which prohibit workers from disparaging the firm, though such restrictions are often commingled within NDAs. Accordingly, we use the broad term ‘NDA’ to refer to contracts that prohibit works from sharing information.

entiate themselves from those pursuing low-road strategies. [Benson et al. \(2020\)](#) describe a model where partially-informed workers seek to recognize firm quality using public employer reputations, where low-information equilibria exist such that low-road firms can survive by pooling themselves with high-road firms, and higher-information equilibria exist such that a smaller share of employers can survive using low-road strategies and it describes experimental evidence from an online labor market consistent with the model. Harnessing data from Glassdoor, [Sockin and Sojourner \(2020\)](#) offer evidence that jobseekers in the broader labor market strongly value negative information about a firm’s employment practices and that workers are reluctant to supply such information out of fear of employer retaliation.

To study whether and how NDAs affect the supply of (negative) information in labor markets, we examine how individual reviews of employers on Glassdoor change between jobs which vary in their propensity for being bound by an NDA and in response to the passage of laws which prohibit firms from using NDAs to conceal unlawful activity (laws we refer to as ‘narrowing NDAs’). Our main results suggest that NDAs suppress the provision of negative information about employers. After the passage of laws that weaken NDAs, the supply of reviews becomes more negative, especially in industries with higher NDA prevalence. These effects disproportionately come from an increase in 1- and 2-star reviews (on a 5-star scale), occur across all job characteristics, and also show a 5% increase in negative review text related to harassment.

We investigate several dimensions of heterogeneity to shed light on the mechanisms that underlie why NDAs inhibit sharing negative information. In particular, we suggest that employees who have more negative information to share and who are most likely to fear the consequences of violating an NDA should be more willing to share negative reviews following the passage of policies that weaken NDAs. In particular, current (vs. former) employees may be particularly unwilling to share negative information about their firm and risk violating an NDA, out of fear of reprisals ([Sockin and Sojourner, 2020](#)). Among current employees, we also expect that since late-tenure employees have likely had better experiences at the

firm (since they elected to stay), early-tenure employees will be more likely to share negative reviews of their employer when states narrow NDAs. We find evidence of both of these patterns. We also examine gender differentials, and find that men are somewhat more likely to write negative reviews in response to these policies, though gender is missing for many reviews. To avoid these state law changes, firms operating in multiple states can invoke another state’s laws in their employment contracts. In line with this argument, we find that the increase in negative reviews is larger in firms that operate in a single state.

The policies that narrow NDAs reduce workers’ risk of employer retaliation for speaking out. Accordingly, workers leaving negative reviews may also be less likely to conceal aspects of their identity (e.g., by omitting a job title in a review). Consistent with this idea, we find that workers in industries where NDAs are more prevalent are 1.7% less likely to conceal their job title following the passage of these laws, relative to industries where NDAs are less prevalent (and relative to the same difference in control states). To the outside jobseeker, reviewer identity concealment degrades the value of the information supplied, likely because it makes it more difficult for the jobseeker to judge the information’s relevance for their own decision ([Sockin and Sojourner, 2020](#)). Policies that narrow NDAs create new value for jobseekers by both increasing the supply of negative information and its value by reducing identity concealment.

Finally, we study whether policies that weaken NDAs help workers distinguish between ‘high-road’ from ‘low-road’ prospective employers—where the reputation of low-road firms may have been otherwise propped up by the suppressive effects of NDAs. The negative reviews that arise in response to the passage of these NDA policies in high-NDA-use industries are disproportionately coming from firms that previously had high ratings. More directly, the variance of reviews across firms within a market rises after NDAs are weakened, suggesting that NDAs were compressing the distribution of reviews, preventing ‘high-road’ employers from separating themselves from ‘low-road’ employers, consistent with predictions of employer reputation theory and evidence from [Benson et al. \(2020\)](#). As jobseekers grow

more informed, it becomes harder for low-road employers to pool with high-road ones.

These results are, for the most part, robust to different measures of NDA exposure, different sets of control states, continuous measures of the ‘pros’ and ‘cons’ of the reviews, alternative weighting schemes, stacked designs, and specifications with firm-state fixed effects.

Since the passage of policies that narrowed NDAs may discourage wrongdoing, our interpretation of these results is that they reflect an uptick in the reporting of prior negative experiences rather than an increase in actual negative experiences per se. Indeed, we interpret these results as short-term effects. Over time, policies that improve the flow of information about employer practices may improve the distribution of actual employer quality. ‘Low-road’ employers may substantively address the issues raised in negative reviews or may find it more difficult to hire.

Taken together, our work integrates across and contributes to the body of literature on labor market sorting as well as to the contracting literature around restrictive covenants. In the latter literature, legal scholars have been particularly concerned about “hushing contracts” because of the externalities they impose on workers and firms, and frequently appeal to public policy arguments as a remedy (Bast, 1999; Bagchi, 2015; Hoffman and Lampmann, 2019; Note, 2006; Lobel, 2020). Legal scholars and policymakers have also pointed to NDAs as barriers to workers and enforcement agencies recognizing racial and gender discrimination within firms (Yang and Liu, 2021). However, due to data limitations, only recently have scholars documented the ubiquity of NDAs (Starr et al., 2019; Balasubramanian et al., 2021), and no prior empirical work has examined the potential externalities arising from such contracts, likely because it is historically rare to observe a shift in how courts treat NDAs.²

By documenting that NDAs allow firms to suppress the flow of negative information, our work joins the body of literature in economics concerned with labor market sorting,

²Hoffman and Lampmann (2019) specifically note that “We do not have a counterfactual firmly in hand. That is, to know what hush contracts do . . . , the gold standard test would be to find a legal regime that switched from enforcement to nonenforcement of hush contracts To date, such a natural experiment has been unavailable.” (p. 174-175)

firm reputation, and the value of non-wage amenities (Sorkin, 2018; Sullivan and To, 2014; Maestas et al., 2018). Our results suggest that policies that weaken firms’ ability to use NDAs to conceal wrongdoing and that prohibit retaliation against those supplying such information help mitigate negative externalities experienced by prospective job applicants, because (a) negative information is more likely to be supplied and such information is likely to have higher value due to reduced identity concealment by volunteers. Such policies also mitigate negative externalities borne by some employers because they allow high-road employers to credibly differentiate themselves in the labor market, likely increasing labor supply to those firms Benson et al. (2020); Sorkin and Sojourner (2020).

2 Institutional Background

NDAs are one in a class of employment provisions known as restrictive covenants, which seek to restrict what workers can do during and following the end of an employment relationship (Lobel 2020). Other common restrictions include noncompetition agreements, which prohibit departing workers from joining or starting competitors, and nonsolicitation agreements, which prohibit departing workers from soliciting former clients or coworkers. These restrictive covenants can be agreed to at the outset of employment, while employed, or in severance arrangements or settlement agreements. NDAs typically prohibit the use or disclosure of trade secrets, but they can also “purport to protect information that is otherwise public, discoverable, or would not otherwise seem to be particularly confidential” (Flanagan and Gerstein 2019, p.532). For example, in *Schwans Home Serv., Inc., 364 N.L.R.B. No. 20 2016* (page 6), the NDA stipulated that as a condition of employment:

“Employee shall neither directly nor indirectly (i) disclose to any person not in the employ of Employer any Confidential or Proprietary Information, or (ii) *use any such information to the Employee’s benefit, the benefit of any third party or [e]mployer, or to the detriment of Employer . . .*” (emphasis added)

Although legal scholars have long cautioned that NDAs may beget negative externalities

(Bast, 1999), this work has been entirely theoretical, likely because data on NDAs was scarce until recently. In the first detailed analysis of NDAs, Balasubramanian et al. (2021) study 33,000 workers and 1,800 US firms and find that NDAs are the most common restrictive covenant, covering 57% of the workforce and in use by 88% of firms for at least some workers. Moreover, they also show that NDAs are the baseline restrictive covenant: if a worker has (or the firm uses) a noncompete or nonsolicitation agreement, there is at least a 95% chance that the firm also uses an NDA. Finally, they document substantial variation in NDA use by industry and occupation. With regards to industries, accommodation and food services have the lowest incidence (44.1%) and professional, scientific, and technical services having the highest (69.2%).

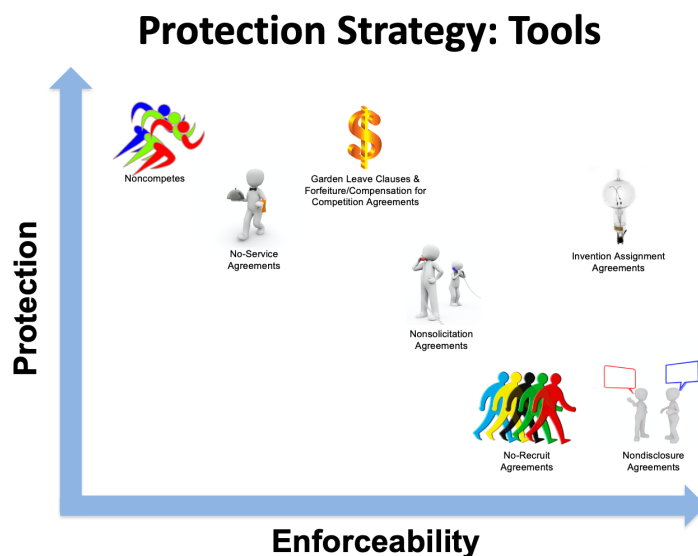
Historically, courts have been quite willing to enforce NDAs. Indeed, among the other similar restrictive covenants, courts have been most willing to enforce NDAs (See Figure 1). This willingness to enforce NDAs derives both from the presumption that there is value for the two parties who have privately agreed not to share such information, and that NDAs impose comparatively weaker restrictions on workers.³

The NDA enforcement status quo has remained largely stable, until recently, when the #MeToo movement put a national spotlight on how men used NDAs to silence women that they had sexually harassed (Silver-Greenberg and Kitroeff, 2020). More importantly, because these victims were prohibited from speaking out, NDAs helped perpetuate such wrongdoing (Carlson, 2019). The #MeToo movement—and the Catholic church sex abuse scandal before that (Philip, 2002)—provide stark examples of negative externalities that arise in the context of NDAs: by preventing the diffusion of valuable (negative) information, NDAs create costs borne by individuals not party to the agreement who may be harmed by “bad actors” in the future, and by “good actors” who are unable to differentiate themselves.

As a result of the #MeToo movement and the essential role NDAs play in covering up and perpetuating wrongdoing, many states reconsidered how to regulate NDAs (Harris,

³In contrast, noncompetes prohibit actual moves to competitors, while non-solicitation agreements prohibit the solicitation of co-workers, clients, or vendors.

Figure 1: Enforceability of NDAs and related restrictions from Beck (2020)



Notes: Figure available in Beck (2020).

2019). While several states passed measures to protect workers from sexual harassment in the workplace or regarding post-harassment settlement agreements, only three states made it unlawful for firms to use NDAs (or other contracts) in the employment context to conceal any unlawful activity (Johnson et al., 2019). We focus on these narrowing changes that narrowed the scope of NDAs, which all passed around the same time.

Beginning on January 1, 2019, California SB 1300 made it an “unlawful employment practice, in exchange for a raise or bonus, or as a condition of employment or continued employment” to require an employee to sign any “document that purports to deny the employee the right to disclose information about unlawful acts in the workplace, including, but not limited to, sexual harassment.”⁴ The bill also prohibits retaliation.

Similarly, beginning on March 18, 2019, New Jersey Senate Bill 121 provides that a “provision in any employment contract” that would prohibit current or former employees from revealing “the details relating to a claim of discrimination, retaliation, or harassment”

⁴See https://leginfo.ca.gov/faces/billNavClient.xhtml?bill_id=201720180SB1300 for more details.

are “against public policy and unenforceable” (Hughes and Nacchio 2019).⁵ The bill also protects workers against any retaliation and requires the party seeking to attempt to enforce a contract against public policy to be liable for attorneys fees.

Lastly, Illinois Senate Bill 75, effective January 1, 2020 notes that “Any agreement, clause, covenant, or waiver that is a unilateral condition of employment or continued employment and has the purpose or effect of preventing an employee or prospective employee from making truthful statements or disclosures about alleged unlawful employment practices is against public policy, void to the extent it prevents such statements or disclosures, and severable from an otherwise valid and enforceable contract under this Act.”⁶ The bill also protects workers against retaliation.

In our empirical work below, we exploit both variation in the likelihood that a worker would be bound by an NDA (using data on industries’ NDA prevalence from Balasubramanian et al. 2021), alongside these three shocks, which narrowed the scope of NDAs.

3 Data

To analyze how NDAs influence the information workers share about their employer, we leverage two unique datasets. The core dataset we leverage is from Glassdoor. It contains individual reviews that current or former employees wrote about their employer. Workers are incentivized to leave reviews on the website through a “give-to-get” policy, whereby contributors to the website gain access to the nexus of information. Reviews contain the worker’s overall rating of their firm on a Likert scale of 1–5 stars, with more stars signaling more satisfaction. Each review also contains free-response text from which we can gain deeper insight into the content of the review.⁷

We also exploit variation in the incidence of NDAs in jobs to examine additional heterogeneity. To do this, we aggregate the NDA data in Balasubramanian et al. (2021) to

⁵See https://www.njleg.state.nj.us/2018/Bills/S0500/121_R2.PDF for the full bill.

⁶See <https://www.ilga.gov/legislation/publicacts/101/101-0221.htm> for the full bill.

⁷Appendix Figure 5 provides a sample, blank review form.

calculate the use of NDAs by industry,⁸ and merge it in to the Glassdoor data. We use the industry level as to opposed to, e.g., the occupation or industry-occupation level, because, as we discuss later, occupations may be left blank in the data as a form of concealment, while industry is fixed by the firm. Combining the Glassdoor and Payscale industry classifications results in fifteen industries that we use throughout the paper.⁹

4 Cross-Sectional Analyses

In this section we examine how variation in the industry-level incidence of NDAs relates to the types of reviews individuals give of their employers. If NDAs suppress the amount of negative information shared across platforms, then we should see that where NDAs are more common that NDAs are likely to be associated with higher overall average ratings (because the negative reviews are missing).

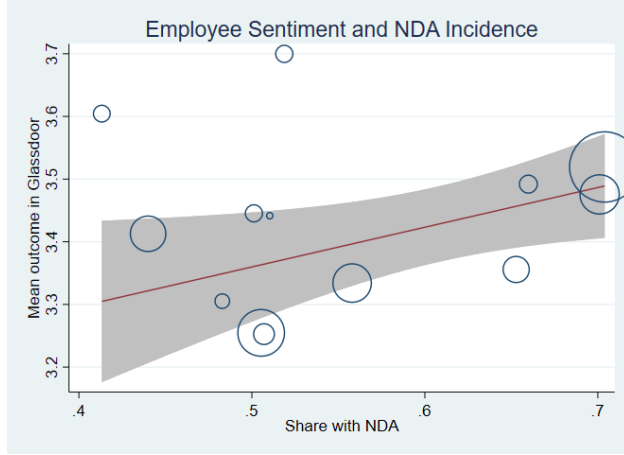
Figure 2 belows convey this relation via a scatterplot. For each industry, we calculate the average overall rating and plot it against the share of workers with an NDA from the Payscale survey. The result shows a strong positive correlation between the intensity of NDA usage and the average reported job satisfaction from workers in that portion of the labor market.

To examine what variation is driving this relationship, we can regress the continuous measures of NDA intensity on overall rating, iteratively removing the variance across time, locations, job titles, and workers. The results are displayed in Table 1 below. Column (1) indicates that jobs with a 10 percentage point increase in the likelihood of signing an NDA give their employers a 0.086 higher rating on Glassdoor.com. This positive effect moderates

⁸The data in [Balasubramanian et al. \(2021\)](#) derive from individual intake data collected by Payscale.com. In particular, individuals who visit Payscale.com can fill out information about themselves to receive an estimate of their earnings potential. In 2017, Payscale.com added a question on NDAs to their intake survey. Individuals were incentivized to provide accurate information because the validity of their earnings prediction depended on it. The individual responses cover 33k workers.

⁹These include Accommodation and Food Services, Agriculture, Arts and Entertainment, Construction, Finance and Insurance, Health Care and Social Assistance, Information, Manufacturing, Mining, Other Services, Professional, Scientific and Technical Services, Real Estate, Retail Trade, Transportation and Warehousing, and Utilities.

Figure 2: NDA Incidence and Average Employer Rating by Industry



Notes: Each dot represents a unique industry. Observations are weighted by the number of Glassdoor reviews within each pairing. NDA intensity by industry obtained from Payscale data.

somewhat when controlling for location fixed effects (column 2), and falls by approximately half when holding constant worker job titles. Finally, column (4) exploits the fact that individuals can leave multiple reviews for different employers with different likelihoods of using NDAs, allowing for the inclusion of individual fixed effects. Even in this stringent specification, workers in industries more likely to be bound by NDAs give higher reviews.

Table 1: NDA Usage and Review Outcomes

	Overall rating			
(Industry) NDA intensity	0.855*** (0.072)	0.769*** (0.070)	0.466*** (0.056)	0.315*** (0.070)
Year-week FE	✓	✓	✓	✓
Location FE		✓	✓	✓
Job title FE			✓	✓
Worker FE				✓
N	3179411	3108815	2904699	245443
Adjusted R ²	0.01	0.02	0.07	0.39

Notes: The dependent variable in each regression specification is detailed above each column. Standard errors are two-way clustered by firm and state. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively.

The core challenge with these cross-sectional analyses, however, is that there are many reasons why NDAs might drive higher ratings—and only one of them is that NDAs are suppressing negative information. For example, a likely theory is that NDAs actually cause increases in average satisfaction. This is because NDAs likely increase profitability, if they are effective in preventing valuable trade secrets and information from leaking out. If the firm uses these profits to invest in the amenities for their workers, then workers may indeed be more satisfied in their jobs. Alternatively, it’s also possible that NDAs are correlated with many other characteristics that reflect firm quality. Accordingly, this analysis can tell us little about whether NDAs actually improve firm quality, or whether the observed associations are a mirage driven by the suppression of negative information. To disentangle those stories, we need a shock which exogenously allows the suppressed negative information to be released, while retaining the trade secret protection aspects of NDAs.

5 Empirical Approach

To address whether NDAs actually cause workers to withhold negative reviews of their employers, we exploit the policy shocks in California, Illinois, and New Jersey described above, which curtailed the ability of NDAs to conceal information related to unlawful wrongdoing, including but not related to discrimination, harrassment, and retaliation. We will broadly refer to these policies as ‘narrowing NDAs.’

The passage of these three policies that narrowed NDAs makes a difference-in-difference style analysis a natural fit, since there are myriad control states that did not pass these provisions. The key challenge with this approach, however, is that the bills that narrowed NDAs also came with a wide swath of other new laws geared towards addressing similar issues raised by the #MeToo movement, such as new training regimens. As a result, a comparison of the states that narrowed NDAs to the states that did not is unlikely to pick-up just the effects of the NDA. Accordingly, to better isolate the NDA-specific aspect of these policies,

we exploit industry-level heterogeneity in the likelihood that an individual is bound by an NDA to perform within-state analyses that net out any state-wide effects that these policies had. In addition, since the within-state analyses may be biased by industry-specific trends, we leverage a triple-difference analyses to net out these differences using control states. We revisit alternative specifications in Appendix A.

Accordingly, our main specification is of the form:

$$Y_{ikst} = \beta \times nda_{\iota(k)} \times Post_{st} + \lambda_{st} + \lambda_{\iota(k)s} + \lambda_{\iota(k)t} + \lambda_k + \epsilon_{ikst} \quad (1)$$

where Y_{ikst} is a measure of worker satisfaction for worker i employed at firm k in state s in calendar year-week t , $nda_{\iota(k)}$ is the intensity with which NDAs are used in industry $\iota(k)$, $Post_{st}$ is a dummy for legislation was passed in state s after year-week t , λ_{st} are state-year-week fixed effects, $\lambda_{\iota(k)s}$ industry-state fixed effects, $\lambda_{\iota(k)t}$ industry-year-week fixed effects, and λ_k firm fixed effects. We iterate between models leveraging variation within the treated states between high and low NDA incidence jobs, and a triple difference model that incorporates other control states. In terms of control states, we would ideally like to include states whose trends reflect the counterfactual trends had our treated states not narrowed NDAs. As this information is unobservable, we consider models that include all US states and also just the regional neighbors defined as bordering states, though our results are largely consistent between these sets of control groups.

In the most saturated triple difference specification, our coefficient of interest β is identified by comparing how worker reviews changed for the same firm before and after the policies that narrowed NDAs passed (first difference), in high NDA-use-industries versus low-NDA-use-industries within the same state (second difference), and relative to the same differences in states that did not pass legislation (third difference). If NDAs have a suppressing effect on firm reputation—by restricting the flow of negative content—then β should be negative.

6 Results

6.1 Review Ratings

Our main set of analyses examines how reviews of employers change following the adoption of policies that narrowed NDAs, for workers more versus less likely to be bound by one. Table 2 examines how individual reviews changed on average following the narrowing of NDAs between high and low NDA-use jobs in the treated states (columns 1 and 2) and in triple difference specifications that include neighboring states and all states (columns 3-6). In even columns we also add firm-state fixed effects such that the identifying variation is coming from reviews of the same firm in the same state. Overall, the main point estimates are quite similar in magnitude, suggesting that, after laws narrowed NDAs, average ratings supplied fell about 0.3 stars (10% of sample mean) in industries where everyone has NDAs, or by 0.03 stars (1% of sample mean) more in industries with 10 percentage point higher NDA prevalence than comparison industries.

Table 2: Three Law Changes and Review Rating, Industry Variation

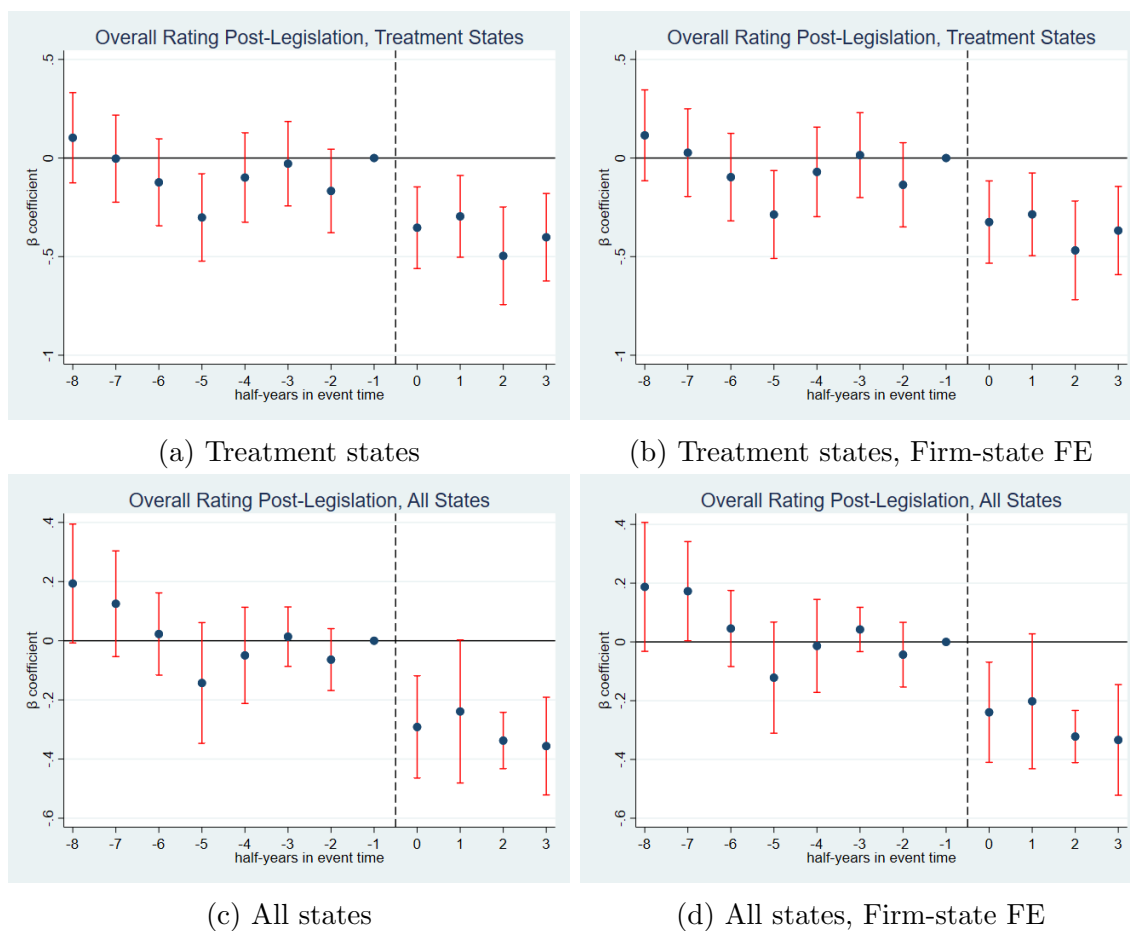
	Star rating					
	CA-IL-NJ	CA-IL-NJ	+ Neighbors	+ Neighbors	All	All
(Industry) NDA intensity x post-legislation	-0.308*** (0.063)	-0.307*** (0.064)	-0.321*** (0.065)	-0.304*** (0.069)	-0.326*** (0.072)	-0.314*** (0.074)
States in sample	CA-IL-NJ	CA-IL-NJ	+ Neighbors	+ Neighbors	All	All
Uses firm-state fixed effects		✓		✓		✓
Outcome sample average	3.49	3.49	3.44	3.44	3.44	3.44
N	681184	673757	1374556	1334445	3114773	2975530
Adjusted R ²	0.16	0.17	0.16	0.17	0.15	0.17

Notes: The dependent variable in each regression is employee overall star rating. Regressions include firm, industry–state, and state-year-week fixed effects. Columns 3–6 additionally include industry-year-week fixed effects. Standard errors are clustered by firm for the first two specifications and two-way clustered by industry and state for the latter four. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively.

The key identifying assumption in these models is that the comparison group reflects the unobserved counterfactual of the treated group in the post-period. To assess whether the control groups moved in parallel fashion the pre-period, Figure 3 reports the dynamic

responses for the within-treated-state analyses and the specification with all states. In each case we observe parallel pre-trends before the policies take effect, with drops in the average ratings in the post period.

Figure 3: Three Law Changes and Employee Ratings, Dynamic Responses 2013–2020, Industry Variation



Notes: The dependent variable in each regression is employee overall star rating. The sample period is 2013–2020 and point estimates are relative to the calendar half-year before the legislation goes into effect. Standard errors are clustered by firm for the upper panels and two-way clustered by industry and state for the lower panels. Red vertical bars indicate 95% confidence intervals around each point estimate.

While the analyses above examine the average effect on reviews, it is also illuminating to consider what effects narrowing NDAs have on the distribution of reviews. To do this, we exploit the fact that Glassdoor reviews offer 5 discrete choices, with 5 being the highest level of satisfaction and 1 being the lowest. Accordingly we run linear probability models with indicators for each specific satisfaction level as the dependent variable, and examine where

in the distribution of reviews the NDA policies are having the most bite.

Consistent with the idea that NDAs are preventing workers from sharing negative information, we find that following the policies that weakened NDAs, industries in which NDAs were more common were more likely to provide 1 and 2 star reviews, and substantially less likely to provide 5 star reviews. To provide some sense of the size of these estimates, a 10 percentage point increase in the use of NDAs increases the probability of seeing a one star review by 0.6 percentage points (column 2) after NDAs are weakened, which is a 3.6% increase relative to the sample average. Similarly, the likelihood of a 5 star review falls by approximately 3.3% for the same comparison.

Table 3: Narrowing NDAs and the Distribution of Review Ratings

	1{1 star}		1{2 stars}		1{3 stars}		1{4 stars}		1{5 stars}	
(Industry) NDA intensity x post-legislation	0.055*** (0.013)	0.053*** (0.013)	0.035*** (0.010)	0.023** (0.008)	0.009 (0.014)	0.027*** (0.003)	-0.033* (0.017)	-0.011 (0.012)	-0.065*** (0.022)	-0.092*** (0.022)
States in sample	CA-IL-NJ	All	CA-IL-NJ	All	CA-IL-NJ	All	CA-IL-NJ	All	CA-IL-NJ	All
Uses firm-state fixed effects										
Outcome sample average	0.152	0.154	0.111	0.116	0.177	0.184	0.221	0.225	0.340	0.321
N	681184	3114773	681184	3114773	681184	3114773	681184	3114773	681184	3114773
Adjusted R ²	0.10	0.08	0.02	0.02	0.04	0.04	0.04	0.03	0.19	0.18

Notes: The dependent variable in each regression is a dummy variable for the specific star rating. Regressions include firm, industry–state, and state-year-week fixed effects. Columns 3–6 additionally include industry-year-week fixed effects. Standard errors are clustered by firm for the first two specifications and two-way clustered by industry and state for the latter four. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively.

We also examine the content areas in which these negative ratings might arise. Glassdoor allows individuals to rate several aspects of their job, including career opportunities, compensation and benefits, culture and values, senior management and work life balance. In Table 4 we examine how the narrowing of NDAs change how individuals rate each of these areas of the firm. In each subarea, after the NDA policies are passed workers with a high likelihood of signing an NDA rate each of these areas lower, regardless of the particular control group. Interestingly, on the basis of the size of the point estimate in the triple-difference models, the most negative effects are driven by views about senior management (across all control groups).

Table 4: Three Law Changes and Review Sub-Ratings

	Glassdoor star rating					
<i>Panel A: Career opportunities</i>						
(Industry) NDA intensity x post-legislation	-0.234*** (0.070)	-0.227*** (0.070)	-0.272*** (0.070)	-0.241*** (0.074)	-0.285*** (0.082)	-0.263*** (0.079)
Outcome sample average	3.26	3.26	3.22	3.22	3.22	3.22
N	745542	738234	1494606	1454820	3370612	3232635
Adjusted R ²	0.17	0.17	0.16	0.17	0.15	0.16
<i>Panel B: Compensation & benefits</i>						
(Industry) NDA intensity x post-legislation	-0.369*** (0.060)	-0.356*** (0.061)	-0.240*** (0.064)	-0.229*** (0.067)	-0.235*** (0.056)	-0.245*** (0.050)
Outcome sample average	3.35	3.35	3.31	3.31	3.32	3.31
N	745132	737830	1493602	1453828	3369834	3231924
Adjusted R ²	0.20	0.20	0.19	0.20	0.18	0.20
<i>Panel C: Culture & values</i>						
(Industry) NDA intensity x post-legislation	-0.081 (0.076)	-0.076 (0.077)	-0.248*** (0.074)	-0.223*** (0.078)	-0.251*** (0.051)	-0.243*** (0.044)
Outcome sample average	3.42	3.42	3.38	3.38	3.38	3.38
N	743234	735949	1489196	1449558	3357875	3220388
Adjusted R ²	0.17	0.18	0.17	0.18	0.16	0.17
<i>Panel D: Senior management</i>						
(Industry) NDA intensity x post-legislation	-0.323*** (0.075)	-0.321*** (0.075)	-0.338*** (0.076)	-0.319*** (0.081)	-0.348*** (0.075)	-0.347*** (0.064)
Outcome sample average	3.12	3.12	3.07	3.07	3.07	3.07
N	738256	730963	1479583	1440045	3336422	3199332
Adjusted R ²	0.18	0.18	0.17	0.18	0.16	0.17
<i>Panel E: Work/life balance</i>						
(Industry) NDA intensity x post-legislation	-0.194*** (0.061)	-0.190*** (0.061)	-0.234*** (0.065)	-0.208*** (0.067)	-0.225*** (0.038)	-0.188*** (0.030)
Outcome sample average	3.40	3.40	3.35	3.34	3.34	3.33
N	746842	739517	1496865	1457034	3376255	3238169
Adjusted R ²	0.16	0.16	0.15	0.16	0.15	0.16
States in sample	CA-IL-NJ	CA-IL-NJ	+ Neighbors	+ Neighbors	All	All
Uses firm-state fixed effects		✓		✓		✓

Notes: The dependent variable in each regression is a dummy variable for the specific star rating. Regressions include firm, industry–state, and state-year-week fixed effects. Columns 3–6 additionally include industry-year-week fixed effects. Standard errors are clustered by firm for the first two specifications and two-way clustered by firm and state for the latter four. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively.

The reviews on Glassdoor also contain written responses regarding the worker’s subjective views on the pros and cons of the workplace. We use these written reviews to bolster two parts of our analysis. First, rather than using a discrete measure of employee satisfaction, we can examine measures that capture how much effort employees expend sharing about

the ‘cons’ relative to the ‘pros’ of the workplace. Table 5 reports this analysis, using as a dependent variable the share of the review text that is attributable to the ‘pros’ field, a continuous measure of worker satisfaction. As with our prior results, individuals are not only giving firms lower ratings, they are also spending more effort in elaborating on the downsides of working at the firm.

Table 5: Post-Legislation Effect on Continuous Measure of Job Satisfaction

	Pros section share of review text					
(Industry) NDA intensity x post-legislation	-0.038*** (0.008)	-0.039*** (0.008)	-0.024*** (0.009)	-0.021** (0.010)	-0.024** (0.009)	-0.022** (0.008)
States in sample	CA-IL-NJ	CA-IL-NJ	+ Neighbors	+ Neighbors	All	All
Uses firm-state fixed effects		✓		✓		✓
Outcome sample average	0.49	0.49	0.48	0.48	0.48	0.48
N	681184	673757	1374556	1334445	3114773	2975530
Adjusted R ²	0.13	0.13	0.13	0.13	0.12	0.13

Notes: The dependent variable is the share of the review’s text that is attributable to the pros section. Regressions include firm, industry–state, and state-year-week fixed effects. Columns 3 and 4 additionally include industry-year-week fixed effects. Standard errors are clustered by firm for the first two specifications and two-way clustered by industry and state for the latter four. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively.

Lastly, because these laws were passed in the wake of the #MeToo movement, and all of them deal with harassment in some way, in Table 6 we examine whether individuals are more likely to use language related to harassment. Specifically, we create an indicator $1\{harassment\}$ if the review reveals some dimension of ‘abuse’ in the “Cons” field of the review. This indicator is equal to 1 if any of the following terms are mentioned in the Cons field: harass, stalk, abuse, abusive, bully, mobb, intimidat, victim, violen, humiliat, threaten, hostile, assault, sexual, innuendo. We then implement the same difference in differences to see if employee abuse is more frequently discussed after these laws are passed. Indeed, in all specifications we see that reviews indicating harassment are more common. To interpret the size of the point estimate, a 10 percentage point increase in the likelihood of signing an NDA increases the likelihood of writing a review about harassment by 5.5-9%.

Table 6: Three Law Changes and Harassment, Industry Variation

	1(Worker mentions employee harassment in review text)					
(Industry) NDA intensity x post-legislation	0.017*** (0.004)	0.017*** (0.004)	0.012** (0.005)	0.011** (0.006)	0.010*** (0.003)	0.008** (0.003)
States in sample	CA-IL-NJ	CA-IL-NJ	+ Neighbors	+ Neighbors	All	All
Uses firm-state fixed effects		✓		✓		✓
Outcome sample average	0.021	0.021	0.021	0.021	0.021	0.021
N	681170	673743	1374526	1334417	3114708	2975475
Adjusted R ²	0.03	0.03	0.03	0.03	0.02	0.02

Notes: The dependent variable in each regression is a dummy variable for the ‘cons’ section of the review text mentions a word pertaining to employee harassment. Regressions include firm, industry–state, and state-year-week fixed effects. Columns 3–6 additionally include industry-year-week fixed effects. Standard errors are clustered by firm for the first two specifications and two-way clustered by industry and state for the latter four. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively.

6.2 Heterogeneity

The prior section establishes that following regulation prohibiting firms from using NDAs to conceal unlawful information, workers in jobs that have a high propensity to sign NDAs are more likely to provide negative reviews of their employer (relative to both workers less likely to sign NDAs in their same state and relative to the same difference workers in states whose policies did not change). In this section we examine several dimensions of heterogeneity to explore which types of workers in which types of firms NDAs are particularly likely to silence.

Our analyses are driven by both theoretical and practical concerns. From a theoretical perspective, NDAs are likely to conceal negative information when (a) an individual actually has negative information to share, and (b) when the additional costs to sharing information induced by the NDA are larger. To address the latter, we consider reviews by workers who are current versus former employees. Current employees may face additional retaliation by the firm for violating an NDA (as opposed to only expected legal costs), and so are more likely to be silenced by an NDA than former workers.

Table 7 shows the results of our main specifications, splitting the sample by whether the individual is a current or former employee. Consistent with our expectations, we find that in both samples narrowing NDAs reduced the average reviews, but the effects are 2-4 times larger for current as opposed to former employees, depending on the control group.

In addition, we also expect that among workers who are currently employed that workers who are early in their tenure would have more negative information to share. This is because workers in later tenures have selected to stay, suggesting that they have had relatively good experiences. Table 8 reports the results of a this analysis, where we split our sample by workers who are early in their tenure (two years or less) at the firm and workers who are later in their tenure (at least three years). Consistent with the idea that NDA-bound short-tenure workers would be less likely to share negative information, we find that when NDAs are weakened such workers provide more negative reviews.

Another natural cut of the data based on the nature of the environment under which

Table 7: Three Law Changes and Review Rating, by Current/Former Status

	Current employees			Former employees		
(Industry) NDA intensity x post-legislation	-0.234*** (0.074)	-0.342*** (0.073)	-0.316*** (0.072)	-0.149** (0.072)	-0.080 (0.070)	-0.188*** (0.055)
States in sample	CA-IL-NJ	+ Neighbors	All	CA-IL-NJ	+ Neighbors	All
Uses firm-state fixed effects						
Outcome sample average	3.906	3.86	3.853	3.010	2.98	2.977
N	354843	713498	1627821	311337	634430	1438411
Adjusted R ²	0.20	0.20	0.19	0.19	0.19	0.17

Notes: The dependent variable in each regression is a dummy variable for the specific star rating. Regressions include firm, industry–state, and state-year-week fixed effects. Columns 3–6 additionally include industry-year-week fixed effects. Standard errors are clustered by firm for the first two specifications and two-way clustered by industry and state for the latter four. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively.

Table 8: Three Law Changes and Review Rating, by Short/Long Tenure

	Short tenure employees			Long tenure employees		
(Industry) NDA intensity x post-legislation	-0.398*** (0.073)	-0.360*** (0.073)	-0.375*** (0.053)	0.026 (0.092)	-0.220* (0.120)	-0.198** (0.086)
States in sample	CA-IL-NJ	+ Neighbors	All	CA-IL-NJ	+ Neighbors	All
Uses firm-state fixed effects						
Outcome sample average	3.519	3.47	3.455	3.457	3.42	3.442
N	339599	679330	1548061	209897	431117	988400
Adjusted R ²	0.18	0.18	0.17	0.17	0.17	0.16

Notes: The dependent variable in each regression is a dummy variable for the specific star rating. Regressions include firm, industry–state, and state-year-week fixed effects. Columns 3–6 additionally include industry-year-week fixed effects. Standard errors are clustered by firm for the first two specifications and two-way clustered by industry and state for the latter four. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively.

these laws were passed is whether women bound by NDAs become more likely to share negative information about their employer. To do this, we re-run our main analyses and split the sample by the gender of the individual reviewer. Table 9 reports the results from our main specifications. We find that, while women do indeed provide more negative reviews, the point estimates are somewhat larger for men. It’s also worth pointing out that many respondents do not include their gender (indeed the sample with gender unknown is larger than the sample of males or females). We find similarly negative point estimates for the

gender-unknown group, but cannot distill which of these are female or male reviews.

Table 9: Three Law Changes and Review Rating, by Gender

	Male employees			Female employees			Gender unknown		
(Industry) NDA intensity x post-legislation	-0.409*** (0.097)	-0.390** (0.158)	-0.293** (0.124)	-0.247** (0.109)	-0.288* (0.162)	-0.255* (0.140)	-0.121 (0.082)	-0.240*** (0.074)	-0.287*** (0.052)
States in sample	CA-IL-NJ	+ Neighbors	All	CA-IL-NJ	+ Neighbors	All	CA-IL-NJ	+ Neighbors	All
Uses firm-state fixed effects									
Outcome sample average	3.483	3.43	3.423	3.350	3.31	3.313	3.561	3.53	3.525
N	203093	397520	894651	154282	325751	757308	295608	601541	1371747
Adjusted R ²	0.17	0.17	0.16	0.16	0.16	0.15	0.17	0.17	0.16

Notes: The dependent variable in each regression is employee overall star rating. Regressions include firm, industry–state, and state-year-week fixed effects. Columns 3–6 additionally include industry-year-week fixed effects. Standard errors are clustered by firm for the first two specifications and two-way clustered by industry and state for the latter two. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively.

Finally, one way that firms can potentially avoid the policies that weakened NDAs is by using choice of law provisions that stipulate that in the event of contract breach a different state’s law be applied (Sanga, 2014). While we do not know what choice of law provision each firm has in its employment contract, we design our empirical specification based on the assumption that single-state employers have chosen the laws of the state in which they operate, while multi-state employees could choose myriad state laws.

To divide our sample into single-state and multi-state firms, we group firms based on whether all of their reviews come from a single state, or whether they come from multiple states. Table 10 reports the results from our main analyses, splitting the sample by whether the firm is a single-state or multi-state employer. Consistent with our expectations Table 10 shows that the negative reviews stemming from narrowing NDAs in high NDA-use jobs are driven more by firms that operate in a single state, and are comparatively weaker when compared to firms that operate in multiple states.

6.3 Concealment

Up to this point, we have shown that by narrowing NDAs individuals more likely to be bound by them are more likely to share negative information, especially if those individuals are

Table 10: Three Law Changes and Review Rating, by Firm Multi-State Status

	Operates in One State			Operates in Many States		
(Industry) NDA intensity x post-legislation	-0.896*** (0.147)	-0.436** (0.203)	-0.490** (0.201)	-0.232*** (0.068)	-0.314*** (0.068)	-0.306*** (0.063)
States in sample	CA-IL-NJ	+ Neighbors	All	CA-IL-NJ	+ Neighbors	All
Uses firm-state fixed effects						
Outcome sample average	3.61	3.57	3.58	3.46	3.43	3.42
N	102388	184328	378771	578795	1189776	2734031
Adjusted R ²	0.20	0.21	0.21	0.15	0.15	0.14

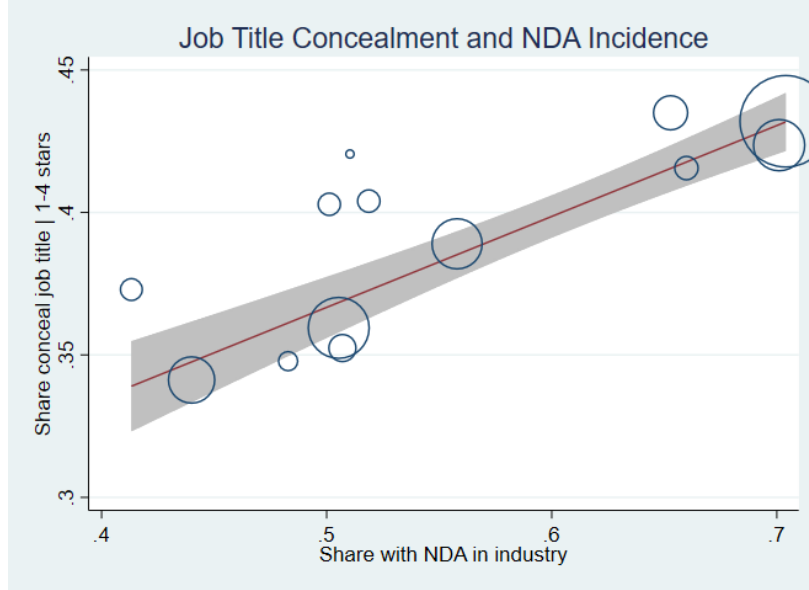
Notes: The dependent variable in each regression is employee overall star rating. Short tenure is defined as two or fewer years with the firm, long tenure as three or more years with the firm. Regressions include firm, industry–state, and state-year-week fixed effects. Columns 3–6 additionally include industry-year-week fixed effects. Standard errors are clustered by firm for the first two specifications and two-way clustered by industry and state for the latter two. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively. Restricted to occupation–industry pairings with at least 10 observations in Payscale.

current employees, early in their tenure, or if their firm operates in a single state. However, an important element of negative reviews is not only the content, but also their credibility. Prior research indicates that when workers leave negative reviews, they tend to conceal aspects of their identity, such as their job title, likely as a way to protect themselves from any retaliation (Socin and Sojourner, 2020). The lack of identifiable information in negative reviews casts doubt on the reliability of such information.

To examine this hypothesis, we examine as a dependent variable whether the respondent includes their job title in their review. Figure 4 examines the cross-sectional relationship between NDA incidence and job-title concealment, conditional on the reviews not being a 5-star review. The figure shows a strong positive correlation such that the likelihood of concealment rises with the use of NDAs. This relationship may be driven, however, by other characteristics. For example, if industries in which NDAs are more likely to be deployed come with greater risk of otherwise sharing a negative review (for whatever reason), then it may not be NDAs that are responsible for this relationship.

To isolate the causal effect of NDAs on the likelihood of concealment, we bolster this cross-sectional analysis by examining how these patterns change when states weakened NDAs using our main specification, but with a new dependent variable indicating job title concealment.

Figure 4: NDA Usage and Rates of Job Title Concealment for 1–4 Star Reviews



Notes: Observations are weighted by the number of Glassdoor reviews within each pairing. NDA intensity obtained from Payscale data. For more detail on the mapping between Glassdoor and Payscale, see Appendix.

Table 11 shows that, consistent with the cross-sectional results, after NDAs are weakened workers in industries are less likely to conceal their job title in their reviews relative to workers in jobs where NDAs are less common (and relative to the same difference in states where NDA policies didn't change).

Table 11: Three Law Changes and Job Title Concealment, Industry Variation

	1(Worker conceals job title 1–4 stars rating)					
(Industry) NDA intensity x post-legislation	-0.265*** (0.020)	-0.269*** (0.020)	-0.052*** (0.017)	-0.057*** (0.018)	-0.059*** (0.015)	-0.060*** (0.012)
States in sample	CA-IL-NJ	CA-IL-NJ	+ Neighbors	+ Neighbors	All	All
Uses firm-state fixed effects		✓		✓		✓
Outcome sample average	0.19	0.19	0.18	0.18	0.17	0.17
N	443347	437825	918556	888810	2095786	1992552
Adjusted R ²	0.12	0.12	0.11	0.12	0.11	0.11

Notes: The dependent variable in each regression is a dummy for the worker conceals the job title attached to the review. Regressions include firm, industry–state, and state-year-week fixed effects. Columns 3–6 additionally include industry-year-week fixed effects. Standard errors are clustered by firm for the first two specifications and two-way clustered by industry and state for the latter two. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively.

One way to assess quantitatively how much the potential for concealment matters for

our initial analysis of reviews is to consider models that use NDA variation that drops those who conceal their job title. To do this we use an occupation-industry measure of NDA incidence (derived from Balasubramanian et al. 2021), which relies on job titles to categorize workers. Table 21 reports these estimates, using our main specification. Relative to our results that leverage just industry NDA incidence, these coefficients are 30%-50% smaller. This analysis thus suggests that limiting the sample to only those who do not conceal information understates the relationship between NDA restrictions and firm reviews.

6.4 Implications for the Distribution of Firm Reviews

Finally, we examine potential NDA-induced negative externalities borne by firms. In particular, because NDAs may prop-up the reputation of ‘low-road’ employers, it may be more difficult for ‘high-road’ employers to effectively separate themselves in the labor market. We address this question with two analyses: The first examines if the policies that weakened NDAs caused more negative reviews for employers with previously high ratings or lower ratings. If NDAs allow ‘low-road’ employers to effectively pool with ‘high-road’ employers, then we should see that firms with previously higher ratings were more likely to receive negative reviews following the passage of laws that weakened NDAs.

Table 12 performs this analysis by splitting the sample into ‘above-average’ firms and ‘below-average’ firms in terms of the firm’s average overall rating between 2015 and 2018 in an industry-state. The results suggest that in high NDA use industries weaker NDAs result in more negative reviews for both low-rating firms and high-rating firms, but that the negative reviews are more than twice as large for the above average firms.

Second, if indeed these policies that weakened NDAs allow high-road employers to stand out, it should be the case that dispersion in reviews within the market increases—as some firms receive more negative reviews while others receive positive reviews. This would have implications for both firms and workers, as workers may be more able to avoid low-quality firms if the distribution of firm reviews is less condensed. If, in contrast, all firms receive

Table 12: Review Rating Split by Firm Quality

	Below average firm			Above average firm		
(Industry) NDA intensity x post-legislation	-0.204** (0.100)	-0.196** (0.100)	-0.194* (0.101)	-0.545*** (0.109)	-0.547*** (0.108)	-0.547*** (0.107)
States in sample	CA-IL-NJ	+ Neighbors	All	CA-IL-NJ	+ Neighbors	All
Uses firm-state fixed effects	✓	✓	✓	✓	✓	✓
Outcome sample average	3.07	3.02	3.02	3.89	3.85	3.85
N	332226	659875	1475774	341531	674609	1499201
Adjusted R ²	0.09	0.09	0.09	0.11	0.11	0.11

Notes: The dependent variable in each regression is employee overall star rating. To determine firm quality, we calculate the average overall rating by firm–state from 2015–2018. A firm–state is considered below (above) average if it falls below (above) the mean firm quality within an industry–state. Regressions include firm–state, industry–state, and state-year-week fixed effects. Columns 3–6 additionally include industry-year-week fixed effects. Standard errors are clustered by firm for the first two specifications and two-way clustered by industry and state for the latter two. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively.

more negative reviews in equal measure, then dispersion should be constant. We investigate this possibility by calculating the average rating for each firm in a given calendar quarter, $\bar{R}_{k_{st}}$, calculating the variance across firms within an industry for each calendar quarter, and again estimating a difference-in-difference specification according to Equation 2,

$$Y_{nst} = \beta \times nda_n \times postLaw_{st} + \lambda_{st} + \lambda_{ns} + \lambda_{n\tau} + \epsilon_{nst} \quad (2)$$

where Y_{nst} is the standard deviation of the firms’ average ratings for industry n in state s in calendar quarter τ . The results, summarized in Table 13, show that while the estimates without the triple differencing are imprecise, the dispersion in firm ratings increases following the passage of policies that weakened NDAs in high-NDA use industries.

7 Sensitivity Analyses

While our results above are largely robust to different potential control groups, in this section we consider several additional sensitivity analyses to further probe the robustness our results.

Table 13: Dispersion in Firm Quality within Labor Markets

	Standard deviation of firm rating					
	CA-IL-NJ	+ Neighbors	All	CA-IL-NJ	+ Neighbors	All
(Industry) NDA intensity x post-legislation	0.149 (0.106)	0.143* (0.075)	0.139** (0.049)	0.128** (0.042)	0.090** (0.037)	0.083*** (0.013)
States in sample	CA-IL-NJ	+ Neighbors	All	CA-IL-NJ	+ Neighbors	All
Industry-state weighted by review count				✓	✓	✓
Outcome sample average	1.34	1.34	1.33	1.34	1.34	1.33
N	864	3072	8520	864	3072	8520
Adjusted R ²	0.62	0.57	0.53	0.80	0.75	0.73

Notes: The dependent variable in each regression is the standard deviation of the (mean) firm rating among firms within an industry–state–quarter. Regressions include industry–state, industry–quarter and state–quarter fixed effects. Standard errors are clustered by industry for the first and fourth specifications, and two-way clustered by industry and state for the other four. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively. Restricted to state–industry pairings with at least 30 firms on average for a given quarter.

First, we consider how states and industries disproportionately contribute to our coefficients of interest. Because the arrival rate of new reviews is not the same across states and industries, and given that each review is given equal weight in our benchmark specifications, naturally states and industries that constitute larger shares of the sample will more heavily weigh on our coefficient of interest. For instance, California represents 18% of the reviews sample, where as Illinois and New Jersey reflect 5.0% and 2.5% of the sample, respectively. To redistribute weight towards smaller states and industries, we consider an alternative specification, where industry-state pairings are given equal weight, meaning each review has weight $1/\sum 1_{i(k)s}$. The results, displayed in Table 14, differ little from our benchmark results.

Second, a recent literature highlights several concerns with staggered adoption two-way fixed effects models. Our results are unlikely to be affected by these issues because the policies we examine were all adopted within one year of each other and because we exploit within-state across-industry variation in the incidence of NDAs. Nevertheless, we employ a stacked regression approach to avoid the fact that some of our treated states serve as control states (Cengiz et al. 2017). To do this, we create a dataset with just one treated state and all the ‘clean’ control states. We then append to that dataset another dataset with a different

Table 14: Review Rating using Equal Weights Controlling for Distribution of Reviews

	Star rating					
(Industry) NDA intensity x post-legislation	-0.135 (0.123)	-0.145 (0.125)	-0.404*** (0.129)	-0.436*** (0.138)	-0.347*** (0.017)	-0.377*** (0.102)
States in sample	CA-IL-NJ	CA-IL-NJ	+ Neighbors	+ Neighbors	All	All
Weight state-industries equally		✓		✓		✓
Outcome sample average	3.487	3.486	3.44	3.44	3.443	3.438
N	681001	673592	1372206	1332438	3100703	2964245
Adjusted R ²	0.18	0.19	0.22	0.24	0.20	0.23

Notes: The dependent variable in each regression is employee overall star rating. Each review is weighted by $1/N_{ikg(j)st}$ such that each industry–state–year receives equal weight. Sample is restricted to industry–state pairings that receive no fewer than 20 reviews in any year. Regressions include firm, industry–state, and state-year-week fixed effects. Columns 3–6 additionally include industry-year-week fixed effects. Standard errors are clustered by firm for the first two specifications and two-way clustered by industry and state for the latter two. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively.

treated state and all the ‘clean’ control states. We do this again for the third treated state, such that the data are stacked but that within each dataset there is no variation coming from other treated states. We then run our main specification with fixed effects for each dataset. The results, presented in Table 15 confirm our baseline results.

Table 15: Robustness: Stacked Regression

	Star rating			
(Industry) NDA intensity x post-legislation	-0.327*** (0.070)	-0.317*** (0.075)	-0.349*** (0.000)	-0.376*** (0.099)
States in sample	All	All	All	All
Uses firm-state fixed effects		✓		✓
Weight state-industries equally			✓	✓
Outcome sample average	3.44	3.44	3.44	3.44
N	8108035	8089163	8064411	8045632
Adjusted R ²	0.19	0.25	0.25	0.34

Notes: Stacked for each treatment state with control sample. Regressions include firm, industry–state, state-year-week fixed effects, industry-year-week, and treatment-state sample fixed effects. Standard errors are clustered by firm for the first two specifications and two-way clustered by industry and state for the latter two. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively.

Next, we consider an alternative specification where in lieu of a continuous measure of

NDA intensity across industries, we partition industries into low- and high-intensity. An industry is considered low (high) intensity if the share of workers covered by an NDA is below or above the median across industries. In Equation 1, we then substitute a dummy variable for high NDA intensity, $highNDA_{\iota(k)}$, for the continuous intensity measure $nda_{\iota(k)}$. The results for each of our main outcomes of interest are detailed in Appendix B. Following the passage of these laws, high-NDA industries experienced on average 0.06 stars decline in overall ratings, 1.1 percentage points less incidence of job title concealment, and a roughly 5% increase in employees’ detailing harassment at work. We next construct an alternative measure of NDA intensity from the Payscale data at the industry–occupation level, where occupations are *onet50* occupation categories. To obtain occupations for Glassdoor reviews, we use a mapping from job title to occupation that was constructed based on a textual analysis machine learning algorithm. We then re-estimate Equation 1 using a continuous measure of NDA intensity at the industry-occupation level, $nda_{\iota(k)o(j)}$ in lieu of the industry-specific measure $nda_{\iota(k)}$. The results for each of our main outcomes of interest are detailed in Appendix C. Again, the main findings are similar.

Last, our analysis thus far has been restricted to reviews for which a location is available in order to assign reviews to our treatment or control states. However, leaving the location of the review blank is not uncommon, representing just over 42% of reviews. To try and incorporate these reviews into our analysis, we implement an imputation procedure by which reviews are assigned to their highest likelihood state. Although the location for these reviews is missing, the firm is not. The intuition behind our imputation process is to use all of the firms’ reviews for which the location is not missing to estimate a latent distribution of the firms’ reviews across states. We then assign every review for firm k with the location blank to the state s with the highest probability $p_{k,s} = \sum 1_{\iota(k)s} / \sum 1_{\iota(k)}$. If concealing location is not random and instead a strategic decision when revealing more-negative information, then failing to incorporate these reviews may bias our results. At the same time though, incorporating these no-location reviews injects measurement error that will bias our estimates

toward zero. We re-estimate our baseline specification, iteratively lowering the threshold for inclusion into the sample (Table 16). While incorporating these reviews attenuates the magnitude of the effect, our estimates remain robustly negative.

Table 16: Review Rating with Imputation of Missing Location

	Star rating				
(Industry) NDA intensity x post-legislation	-0.326*** (0.072)	-0.287*** (0.066)	-0.245*** (0.056)	-0.236*** (0.055)	-0.232*** (0.057)
States in sample	All	All	All	All	All
Threshold for including missing location	none	100%	75%	50%	25%
Outcome sample average	3.443	3.437	3.436	3.435	3.430
N	3114773	3476701	3903145	4295117	4811065
Adjusted R ²	0.15	0.15	0.15	0.15	0.15

Notes: The dependent variable in each regression is employee overall star rating. Regressions include firm, industry–state, and industry-year-week fixed effects, as well as controls for employment status and whether current employee. Regressions also include fixed effects for year-month x a dummy for location left blank. Threshold for inclusion reflects the minimum number of annual reviews an industry–state can receive from 2015–2020. Standard errors are two-way clustered by industry and state. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively.

8 Conclusion

This study is motivated both by the longstanding concern in the legal literature that overbroad NDAs can create negative externalities (Bast, 1999; Hoffman and Lampmann, 2019), and that prior research has not empirically examined whether such externalities manifest. To this end, we leverage employee reviews on Glassdoor, variation in the incidence of NDAs, and state policy shocks that prohibited firms from using NDAs to conceal unlawful activity to document three primary results: First, NDAs suppress negative information about the firm—especially from workers who face high costs from violating an NDA and who likely have more negative information to share. Second, when employees do share negative reviews of their employer, NDAs cause them to conceal their identity, likely reducing the helpfulness of the review. Third, NDAs compress the distribution of firm reviews, making it harder for

high-road employers to differentiate themselves from low-road employers whose reputations are propped up by NDAs. In the remainder of this section we discuss our interpretation of these findings, and their contributions, implications, and limitations.

A key limitation in interpreting these results is that we cannot observe actual wrongdoing. That is, the fact that workers in high-NDA-use jobs are more likely to provide negative reviews of their employer after the passage of laws that weaken NDAs may suggest that these laws encourage more wrongdoing. While we cannot completely rule out this interpretation, we think it is unlikely for three reasons. First, negative information on employers is durable and can be shared at any point in time, even if it occurred years ago. Second, if firms know they cannot prohibit workers from speaking out by using NDAs, then they will likely be less inclined towards perpetrating wrongdoing in the first place. Third, our results are relatively short-run, giving firms limited time to react to these negative reviews. Indeed, over the long run it is conceivable that firms respond to the issues raised in these negative reviews such that firm ratings improve over time.¹⁰ Accordingly, our preferred interpretation of the results is that the policies that weaken NDAs encouraged workers to share their negative experiences with firms—not that they created more negative experiences. We hope that as time elapses future research examines the potential dynamic implications for firm behavior.

Our results suggests value from policies to ensure workers can freely disparage their employers. Our empirical shock to NDAs—while basically the only one in the literature—is quite weak in that they only free employees to disclose a narrow class of unlawful behavior. This might explain the small magnitude of the effects. However, where non-disclosure clauses that protect trade secrets have some economic benefit, this benefit is absent in the case of non-disparagement clauses. They are basically a way for a low-road employer to cut an insider employee in on the rent the worker’s silence creates through externalities on others. The parties to non-disparagement clauses do not internalize the costs of silence on others and facilitate the survival of low-road employment practices.

¹⁰As in [Johnson \(2020\)](#), revealed wrongdoing by one firm may cause other firms to improve their behavior.

A natural implication of our results relates to the literature on labor market sorting. Prior research has found that non-wage amenities account for a substantial portion of value for workers (Sorkin, 2018; Maestas et al., 2018; Sullivan and To, 2014). However, as we document, firms can use NDAs to conceal negative information on the amenities they offer. Our findings suggest that policies that prohibit firms from using NDAs to conceal wrongdoing increases negative information flows, which should in theory impact labor market sorting since firms with negative reviews are likely to be less attractive to workers. While prior research has documented this fact (Sorkin and Sojourner, 2020), data limitations prevent us from identifying precisely how NDAs affect labor market search behavior. Thus while we think it is likely that the policies that weakened NDAs hurt the ability of firms that received negative reviews to recruit, we cannot document this fact directly. We hope that future research can examine this next step.

Finally, while our focus in this study is on the externalities that NDAs create, policies that prohibit firms from using NDAs to conceal wrongdoing may come with important tradeoffs for the individuals who experienced the wrongdoing in the first place. Some individuals may not wish to share their negative experiences, and may prefer to receive a compensating differential in exchange for their silence. Correlational evidence from Balasubramanian et al. (2021) suggests that the wage premium for signing an NDA in the employment context is on average 6-7% (though the effect is -5.3% when NDAs are coupled with other postemployment restrictive covenants). If state policies prohibit firms from using NDAs to keep workers from speaking out, then the size of the compensating differential workers receive for signing an NDA will likely shrink—especially in severance or settlement arrangements. Given that these payments are typically private and endogenous to the expected or actual wrongdoing, they are necessarily difficult to study. Nevertheless, we hope that future research will engage not only with the potential externalities that NDAs create, but also with how much workers value their freedom to speak out in the first place—and how much they are compensated for giving up that freedom.

References

- A. Bagchi. Other people’s contracts. Yale J. on Reg., 32:211, 2015. [Cited on page 5.]
- N. Balasubramanian, E. Starr, and S. Yamaguchi. Bundling postemployment restrictive covenants: When, why, and how it matters. Working paper, 2021. [Cited on pages 2, 5, 7, 10, and 32.]
- C. M. Bast. At what price silence: are confidentiality agreements enforceable. Wm. Mitchell L. Rev., 25:627, 1999. [Cited on pages 5, 7, and 30.]
- R. Beck. President biden’s proposed ban of (most) noncompetes: Protection strategies and steps to take now. <https://www.faircompetitionlaw.com/2020/12/02/president-bidens-proposed-ban-of-most-noncompetes-protection-strategy-and-steps-to-take-now> 2020. [Cited on page 8.]
- A. Benson, A. Sojourner, and A. Umyarov. Can reputation discipline the gig economy?: Experimental evidence from an online labor market. Management Science, 66:1802–1825, 2020. [Cited on pages 3, 4, and 6.]
- G. Carlson. Gretchen carlson: Fox news, i want my voice back. <https://www.nytimes.com/2019/12/12/opinion/gretchen-carlson-bombshell-movie.html>, 2019. [Cited on page 7.]
- B. Facchinei. Gretchen carlson urges 2020 candidates to end ndas. <https://www.newsyc.com/stories/gretchen-carlson-it-can-be-suffocating-to-be-silenced/>, 2020. [Cited on page 2.]
- E. Harris. Despite #metoo glare, efforts to ban secret settlements stop short. <https://www.nytimes.com/2019/06/14/arts/metoo-movement-nda.html>, 2019. [Cited on page 7.]
- D. A. Hoffman and E. Lampmann. Hushing contracts. Wash. UL Rev., 97:165, 2019. [Cited on pages 5 and 30.]
- A. Johnson, K. Menefee, and R. Sekaran. Progress in Advancing Me Too Workplace Reforms in 20statesby2020. National Women’s Law Center, 2019. [Cited on page 8.]
- M. S. Johnson. Regulation by shaming: Deterrence effects of publicizing violations of workplace safety and health laws. American economic review, 110(6):1866–1904, 2020. [Cited on page 31.]
- O. Lobel. The contract thicket: Addition and supra addition in private law theory. Working paper, 2020. [Cited on page 5.]
- N. Maestas, K. J. Mullen, D. Powell, T. von Wachter, and J. B. Wenger. The value of working conditions in the united states and implications for the structure of wages. Technical Report 25204, National Bureau of Economic Research, October 2018. [Cited on pages 6 and 32.]
- I. Marinescu, N. Klein, A. Chamberlain, and M. Smart. Incentives can reduce bias in online employer reviews. Academy of Management Proceedings, 2018(1):11189, 2018. [Cited on page 2.]
- A. Note. Law and economics look at contracts against public policy, 119 harv. L. Rev., 1445:1460, 2006. [Cited on page 5.]

- R. M. Philip. Silence at our expense: Balancing safety and secrecy in non-disclosure agreements. Seton Hall. L. Rev., 33:845, 2002. [Cited on page 7.]
- S. Sanga. Choice of law: an empirical analysis. Journal of Empirical Legal Studies, 11(4):894–928, 2014. [Cited on page 22.]
- J. Silver-Greenberg and N. Kitroeff. How bloomberg buys the silence of unhappy employees. www.nytimes.com/2020/03/02/business/michael-bloomberg-nda.html, 2020. [Cited on page 7.]
- J. Sockin and A. Sojourner. What’s the inside scoop? addressing workers’ information problems with firm reputation. Working paper, 2020. [Cited on pages 2, 3, 4, 6, 23, and 32.]
- I. Sorkin. Ranking firms using revealed preference. The Quarterly Journal of Economics, 133(3): 1331–1393, 2018. [Cited on pages 6 and 32.]
- E. Starr, J. J. Prescott, and N. Bishara. Noncompetes in the U.S. Labor Force. U of Michigan Law & Econ Research Paper, (18-013), 2019. [Cited on page 5.]
- P. Sullivan and T. To. Search and nonwage job characteristics. Journal of Human Resources, 49 (2):472–507, 2014. [Cited on pages 6 and 32.]
- J. Yang and J. Liu. Strengthening accountability for discrimination. Technical report, Economic Policy Institute, Washington, D.C., Jan. 2021. URL <https://files.epi.org/pdf/218473.pdf>. [Cited on page 5.]

A Appendix

Figure 5: Example blank firm review form from Glassdoor

It only takes a minute! And your anonymous review will help other job seekers.



Company

University of Minnesota

Overall Rating*



Are you a current or former employee?

Current

Former

Employment Status*

Select

Your Job Title at University of Minnesota*

Review Headline*

Pros*

5 word minimum

Share some of the best reasons to work at University of Minnesota

Cons*

5 word minimum

Share some of the downsides of working at University of Minnesota

Thank you for contributing to the community. Your opinion will help others make decisions about jobs and companies.

Please stick to the [Community Guidelines](#) and do not post:

- Aggressive or discriminatory language
- Profanities
- Trade secrets/confidential information

Thank you for doing your part to keep Glassdoor the most trusted place to find a job and company you love. See the [Community Guidelines](#) for more details.

B Appendix: Split into Low and High NDA Industries

In this section we show that our results are robust to defining NDA intensity using a dichotomous measure reflecting above average NDA use.

Table 17: Dispersion in Firm Quality within Labor Markets, Controlling for Firm Composition

	Standard deviation of firm rating					
	CA-IL-NJ	+ Neighbors	All	CA-IL-NJ	+ Neighbors	All
(Industry) NDA intensity x post-legislation	0.253** (0.072)	0.126 (0.082)	0.126** (0.043)	0.212* (0.087)	0.208** (0.062)	0.167*** (0.016)
States in sample	CA-IL-NJ	+ Neighbors	All	CA-IL-NJ	+ Neighbors	All
Industry-state weighted by review count				✓	✓	✓
Outcome sample average	1.03	1.03	1.03	1.03	1.03	1.03
N	408	1176	3000	408	1176	3000
Adjusted R ²	0.61	0.58	0.58	0.76	0.72	0.69

Notes: The dependent variable in each regression is the standard deviation of the (mean) firm rating among firms within an industry–state–quarter. Regressions include industry–state, industry–quarter and state–quarter fixed effects. Standard errors are clustered by industry for the first two specifications and two-way clustered by industry and state for the latter two. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively. Restricted to state–industry pairings with at least 30 firms on average for a given quarter. Sample of firms over which standard deviation of averages is calculated restricted to those that are present in at least 16 of the 24 calendar quarters for the given industry–state pairing.

Table 18: Three Law Changes and Review Rating

	Star rating					
	CA-IL-NJ	CA-IL-NJ	+ Neighbors	+ Neighbors	All	All
High NDA intensity industry x post-legislation	-0.067*** (0.012)	-0.067*** (0.012)	-0.061*** (0.013)	-0.058*** (0.014)	-0.060*** (0.013)	-0.059*** (0.012)
States in sample	CA-IL-NJ	CA-IL-NJ	+ Neighbors	+ Neighbors	All	All
Uses firm-state fixed effects		✓		✓		✓
Outcome sample average	3.49	3.49	3.44	3.44	3.44	3.44
N	681184	673757	1374556	1334445	3114773	2975530
Adjusted R ²	0.16	0.17	0.16	0.17	0.15	0.17

Notes: The dependent variable in each regression is employee overall star rating. Regressions include firm, industry–state, and industry-year-week fixed effects, as well as controls for employment status and whether current employee. Standard errors are clustered by firm for the first specification and two-way clustered by firm and state for the latter two. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively.

Table 19: Three Law Changes and Job Title Concealment

	1(Worker conceals job title 1–4 stars rating)					
High NDA intensity industry x post-legislation	-0.041*** (0.004)	-0.042*** (0.004)	-0.010*** (0.003)	-0.011*** (0.004)	-0.011** (0.004)	-0.012*** (0.004)
States in sample	CA-IL-NJ	CA-IL-NJ	+ Neighbors	+ Neighbors	All	All
Uses firm-state fixed effects		✓		✓		✓
Outcome sample average	0.19	0.19	0.18	0.18	0.17	0.17
N	443347	437825	918556	888810	2095786	1992552
Adjusted R ²	0.12	0.12	0.11	0.12	0.11	0.11

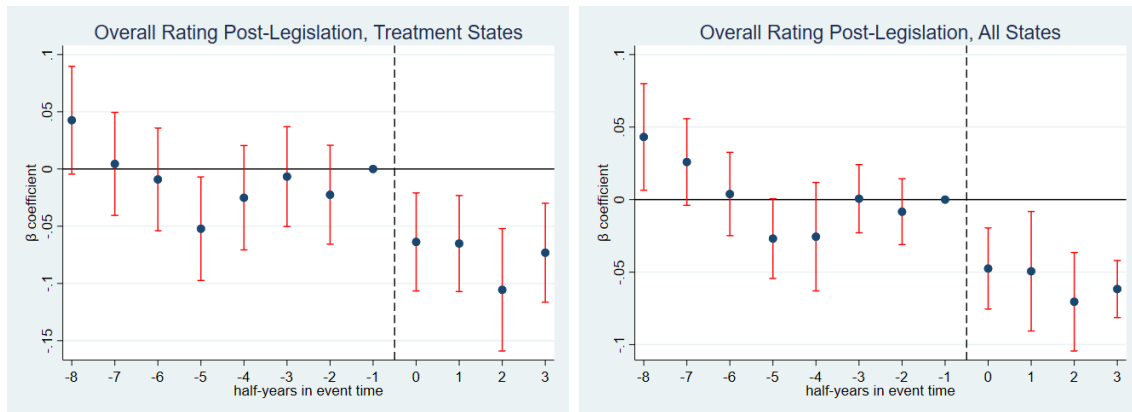
Notes: The dependent variable in each regression is a dummy variable for job title is concealed. Regressions include firm, industry–state, and state-year-week fixed effects. Columns 3–6 additionally include industry-year-week fixed effects. Standard errors are clustered by firm for the first specification and two-way clustered by industry and state for the latter two. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively.

Table 20: Three Law Changes and Mentioning of Harassment

	1(Worker mentions employee harassment in review text)					
High NDA intensity industry x post-legislation	0.003*** (0.001)	0.003*** (0.001)	0.002 (0.001)	0.002 (0.001)	0.001* (0.001)	0.001* (0.001)
States in sample	CA-IL-NJ	CA-IL-NJ	+ Neighbors	+ Neighbors	All	All
Uses firm-state fixed effects		✓		✓		✓
Outcome sample average	0.021	0.021	0.021	0.021	0.021	0.021
N	681170	673743	1374526	1334417	3114708	2975475
Adjusted R ²	0.03	0.03	0.03	0.03	0.02	0.02

Notes: The dependent variable in each regression is a dummy variable for review text mentions harassment. Regressions include firm, industry–state, and state-year-week fixed effects. Columns 3–6 additionally include industry-year-week fixed effects. Standard errors are clustered by firm for the first specification and two-way clustered by firm and state for the latter two. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively.

Figure 6: Three Law Changes and Employee Ratings, Dynamic Responses 2013–2020, Industry Partition by Low- and High-NDA Intensity



(a) Treatment states

(b) All states

Notes: The dependent variable in each regression is employee overall star rating. The sample period is 2013–2020 and point estimates are relative to the calendar half-year before the legislation goes into effect. Standard errors are clustered by firm for the upper panels and two-way clustered by industry and state for the lower panels. Red vertical bars indicate 95% confidence intervals around each point estimate.

C Appendix: NDA Intensity by Industry–Occupation

In this section we show that our results are robust to using NDA variation from Payscale at the more granular level of industry–occupation rather than our benchmark of industry. To reduce measurement error in NDA intensity, we restrict our attention to industry-occupation pairings for which there were at least ten unique observations in the Payscale employee survey.

Table 21: Three Law Changes and Review Rating by Industry–Occupation

	Dependent variable					
	(Ind–occ) NDA intensity x post-legislation	(Ind–occ) NDA intensity x post-legislation	(Ind–occ) NDA intensity x post-legislation	(Ind–occ) NDA intensity x post-legislation	(Ind–occ) NDA intensity x post-legislation	(Ind–occ) NDA intensity x post-legislation
(Ind–occ) NDA intensity x post-legislation	-0.242*** (0.048)	-0.242*** (0.049)	-0.159*** (0.047)	-0.148*** (0.049)	-0.164*** (0.055)	-0.166*** (0.060)
States in sample	CA-IL-NJ	CA-IL-NJ	+ Neighbors	+ Neighbors	All	All
Uses firm-state fixed effects		✓		✓		✓
Outcome sample average	3.53	3.53	3.48	3.48	3.48	3.47
N	389199	383837	783548	754568	1794635	1694630
Adjusted R ²	0.17	0.17	0.16	0.17	0.16	0.17


Notes: The dependent variable in each regression is employee overall star rating. Regressions include firm, occupation-state, industry–state, and state-year-week fixed effects. Columns 3 and 4 additionally include occupation-year-week and industry-year-week fixed effects. Standard errors are clustered by firm for the first two specifications and two-way clustered by firm and state for the latter two. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively. Restricted to occupation–industry pairings with at least 10 observations in Payscale.

Table 22: Three Law Changes and Discussion of Harassment

	Dependent variable					
	(Ind–occ) NDA intensity x post-legislation	(Ind–occ) NDA intensity x post-legislation	(Ind–occ) NDA intensity x post-legislation	(Ind–occ) NDA intensity x post-legislation	(Ind–occ) NDA intensity x post-legislation	(Ind–occ) NDA intensity x post-legislation
(Ind–occ) NDA intensity x post-legislation	0.012*** (0.004)	0.012*** (0.004)	0.008* (0.004)	0.007* (0.004)	0.005** (0.003)	0.006** (0.003)
States in sample	CA-IL-NJ	CA-IL-NJ	+ Neighbors	+ Neighbors	All	All
Uses firm-state fixed effects		✓		✓		✓
Outcome sample average	0.02	0.02	0.02	0.02	0.02	0.02
N	389194	383832	783536	754558	1794600	1694603
Adjusted R ²	0.03	0.03	0.03	0.04	0.03	0.03

Notes: The dependent variable in each regression is a dummy variable for Cons field of review mentions harassment. Regressions include firm, industry-state, and state-year-week fixed effects. Columns 3 and 4 additionally include industry-year-week fixed effects. Standard errors are clustered by firm for the first two specifications and two-way clustered by firm and state for the latter two. One, two, and three stars denote significance at the 10, 5, and 1 percent levels, respectively. Restricted to industry–state pairings with at least 10 observations in Payscale.

Figure 7: Three Law Changes and Employee Ratings, Dynamic Responses 2013–2020, Industry–Occupation NDA Intensity



(a) Treatment states

(b) All states

Notes: The dependent variable in each regression is employee overall star rating. The sample period is 2013–2020 and point estimates are relative to the calendar half-year before the legislation goes into effect. Standard errors are clustered by firm for the upper panels and two-way clustered by firm and state for the lower panels. Red vertical bars indicate 95% confidence intervals around each point estimate.