Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

Tracking Real Time Layoffs with SEC Filings: A Preliminary Investigation

Leland D. Crane; Emily Green; Molly Harnish; Will McClennan; Paul E. Soto; Betsy Vrankovich; Jacob Williams

2024-020

Please cite this paper as:
Crane, Leland D., Emily Green, Molly Harnish, Will McClennan, Paul E. Soto, Betsy Vrankovich, and Jacob Williams (2024). "Tracking Real Time Layoffs with SEC Filings: A Preliminary Investigation," Finance and Economics Discussion Series 2024-020. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2024.020.

Tracking Real Time Layoffs with SEC Filings: A Preliminary Investigation*

Leland D. Crane Emily Green Molly Harnish Will McClennan
Paul E. Soto Betsy Vrankovich Jacob Williams†

April 11, 2024

Abstract

We explore a new source of data on layoffs: timely 8-K filings with the Securities and and Exchange Commission. We develop measures of both the number of reported layoff events and the number of affected workers. These series are highly correlated with the business cycle and other layoff indicators. Linking firm-level reported layoff events with WARN notices suggests that 8-K filings are sometimes available before WARN notices, and preliminary regression results suggest our layoff series are useful for forecasting. We also document the industry composition of the data and specific areas where the industry shares diverge.

*We thank Tomaz Cajner, Ryan Decker, Chris Kurz, Robert Kurtzman, Anderson Monken, and Nitish Sinha for valuable feedback. Opinions expressed herein are those of the authors alone and do not necessarily reflect the views of the Federal Reserve System or the Board of Governors.

†All authors are at the Federal Reserve Board of Governors

1 Introduction

Layoffs are among the most closely-watched macroeconomic indicators. Unemployment insurance (UI) claims and the Job Openings and Labor Turnover Survey (JOLTS) layoffs have long been considered early warnings for downturns. More recently, economists have begun tracking the Worker Adjustment and Retraining Notification (WARN) system which requires the public announcement of certain layoffs months in advance (Krolikowski and Lunsford, 2024). In this paper we explore an alternative timely layoff indicator based on 8-K filings.

Publicly-traded companies file 8-Ks with the Securities and Exchange Commission (SEC) when they need to disclose certain events to the public, and layoff plans are often recorded in 8-Ks. Some layoffs are recorded under a specific 8-K item number (discussed below) while others are recorded under more general items. The latter require natural language processing to identify; we experiment with sentence embeddings from a workhorse language model (BERT) as well as prompting a generative large language model (Llama 2) to identify layoffs. We also explore estimating the quantity of workers laid off by parsing the magnitudes discussed in the filings and linking to employment data from Compustat.

The resulting series are highly correlated with the business cycle and other layoff indicators. There are only two recessions in the sample period, but the 8-K series clearly captures the increase in layoffs for both. Interestingly, the 8-K layoff series have been elevated in 2023; this pattern is not present in other indicators. We leave it to future work to explore this pattern in more depth.

We also present preliminary evidence that the 8-K series are useful for forecasting important quantities such as the unemployment rate and initial unemployment insurance claims. These (in sample) results are robust to the inclusion of many controls including WARN notice data.

The sectoral composition of 8-K announced layoffs is also of interest. We compare the sector shares of 8-K announced layoffs with the composition of publicly-traded firms and the universe of (U.S.) firms. We find that—on a firm count basis—publicly-traded firms are highly concentrated in the manufacturing sector as compared to the universe of public and private firms. 8-K layoff events are even more highly concentrated in manufacturing than the distribution of publicly-traded firms would imply. A significant part of this divergence is due to the pharmaceutical industry: our evidence suggests that pharmaceutical firms are disproportionately likely to be public firms. Publicly-traded pharmaceutical firms also tend to be very small and thus subject to idiosyncratic shocks that would generate layoffs.

Finally, we capitalize on the firm-level nature of the data to link 8-K reported layoffs to the publicly-available WARN Notices, which also record layoffs. We find that—on the linked sample—neither data source clearly dominates in terms of timeliness. In many cases both data sources record the layoff event in the same week, but in a substantial fraction of cases the 8-Ks appear to record the layoff event four or more weeks in advance of the WARN notices. WARN notices also arrive weeks earlier than 8-Ks in many cases. We also examine whether the layoffs in announced in the (whole, unlinked) 8-K data would be subject to WARN notice requirements. Layoffs only need to be announced in a WARN notice if they meet certain thresholds based on plant size and the layoff size. Our evidence suggests that most of the layoffs in the 8-K data would not meet the WARN reporting requirements, though there are many caveats to this comparison.

The paper proceeds as follows: Section 2 describes the raw data and the construction of the series. Section 3 presents the series and comparisons to other labor market indicators. Section 4 explores the industry composition of the data, Section 5 highlights post-pandemic layoff dynamics, and Section 6 compares the data to WARN notices and develops forecasting results. Section 7 concludes.

2 Data

The SEC requires public companies to make certain filings and disclosures; for example, 10-K annual reports have been widely studied (Bodnaruk et al. (2015), Sufi (2009), Loughran and McDonald (2011)). Our focus is on form 8-K, or the "current reports." These forms are filed irregularly, whenever a reportable event occurs. This usually means within four days of a decision or action, though SEC guidance states that firms may wait to disclose layoffs until the employees are informed about the plan (SEC, 2023). The SEC specifies that different kinds of events be filed under different numbered item codes, which we use to narrow down the set of relevant filings.

Our analysis uses the item codes and the text of the 8-K filings. We begin with the universe of SEC filings, as available on EDGAR.[1] From each document we extract the form type, company identifier (CIK), SIC industry code, company name, filing date, period of report, and the text of the filing. We further clean the text by extracting individual items. To do this we modify and extend the EDGAR Crawler code of Loukas et al. (2021) (designed for reading 10-Ks) to parse the items in 8-Ks. The metadata and the itemized text from each filing are then loaded into a database.

Layoffs are often reported under item 2.05, "Costs Associated with Exit or Disposal Activities." This item—introduced in 2004—requires firms to report the disposal of any long-lived assets or the termination of employees if these actions lead to the firm incurring "material charges" (SEC, 2004).[2] The filing must include the planned course of action, estimated costs for various types of expenses related to the action, and an estimate of the resulting charge that will lead to future cash expenditures. Item 2.05 can be used to report many types of disposal costs—such as fees for terminating sales or marketing contracts—but in practice it appears that severance payments are nearly always among the reported charges.[3] As a result filing an 8-K with an item 2.05 is a very good indicator that layoffs will be occurring at the firm, and we include all items 2.05 as layoff events. Item 2.05 is only present in a small number of filings: we find that in 2022 0.38 percent of 8-Ks included an item 2.05

[1] We do not actually use EDGAR, instead we use the daily feed of all filings through the SEC's vendor Attain. The filings are in XBRL, an form of XML.

[2] The item requires reporting terminations and severance costs as described in FASB ASC paragraph 420-10-25-4.

[3] See the 2006-01-09 filing from Toys "R" Us for an example with several types of charges.

(a total of 258 filings.) Perhaps because of the relatively small number of filings and its focus on labor item 2.05 has received relatively little attention in the finance and accounting literatures, though Jung et al. (2016) and Laurion (2020) are exceptions.

2.1 Layoffs reported under other items

It turns out that layoffs are sometimes reported under 8-K items other than 2.05. In particular, we find layoff announcements under item 7.01 and item 8.01. Item 7.01 (“Regulation FD Disclosure”) relates to Regulation Fair Disclosure, the SEC rule requiring that firms giving material nonpublic information to certain market participants or shareholders also make a public disclosure of the information. Item 8.01 (“Other Events”) can be used to satisfy Regulation FD, or report other events “the registrant deems of importance to security holders” (SEC, 2004). Items 7.01 and 8.01 are much more common than item 2.05: in 2022 they were present in roughly 24 percent and 22 percent of 8-Ks, respectively. However, unlike item 2.05, items 7.01 and 8.01 are used to report many non-layoff events so layoffs are only a very small fraction of the items.

Filtering out the non-layoff-related reports is an empirical challenge. Our initial reading of the documents suggested that firms frequently use the same terms when announcing layoffs: the word “layoff” often appears, as does “reducing”/“reduce”/“reduced” followed by “workforce”. Starting with these phrases—and treating each sentence of each filing as a separate observation—we take an iterative approach:

1. We check that the existing stock of layoff-related phrases does not generate many false positives: false positives occur when sentences match one of the patterns but are not in fact about layoffs.

2. We flag all sentences matching the stock of phrases as true layoffs and fit a flexible model to predict whether or a not a sentence is about layoffs as a function of the sentence text. The model returns a probability that a given sentence is about layoffs.

3. we examine the set of sentences which are very likely about layoffs but not already flagged as such. This examination suggests new phrases that can mark layoffs. At this point the process repeats and we return to step 1 with an expanded stock of phrases.

The flexible model which predicts layoffs is based on BERT sentence embeddings (Devlin et al., 2019). The BERT model converts a sentence to a 768-dimensional vector, meant to capture the semantic meaning of the sentence. We feed the vector into a flexible classifier (XGBoost, Chen and Guestrin (2016)). XGBoost is extremely fast and stable, making it well-suited to our iterative approach.[4]

Among the phrases we identify are "reduction in force", "workforce reduction", "eliminate staff positions", and "reduce headcount". An advantage of this methodology is that—while the iterative ranking process is not particularly transparent—the final classification rule is just a list of patterns that sentences either match or don't match. All told, the layoffs from items 7.01 and 8.01 account for about 16 percent of the total, with item 2.05 comprising the rest. To avoid any double counting we treat all filings by a firm (as defined by the SEC's CIK identifier) on a given day as a single event.[5]

2.2 Further analysis with LLMs

One concern is that the procedure outlined above will only identify sentences with particular grammatical patterns as layoff-related. Sentences that use terminology or phrasing very different from the initial patterns we identify may never be ranked as likely layoffs. To address this issue we experimented with using a large language model to identify layoffs missed by the methodology. In particular, we took a subset of sentences that were scored as fairly likely layoffs (but not classified as layoffs) and prompted the Llama 2 model of

[4]Adding to the challenges of identifying layoffs, many filings include boilerplate disclaimers about forward-looking statements. These disclaimers often refer to hypothetical future layoffs that could materially affect the value of the company. We flag examples of these statements as a separate class when fitting the classifier, anticipating that it would be hard for a naive classifier to distinguish these remarks from statements that are actual, planned layoffs.

[5]There are only 8 cases where multiple items are flagged on a single day, these are all cases where a relevant item 7.01 and a relevant item 8.01 appear together.

Touvron et al. (2023) to classify the sentences as layoffs-related or not.[6] Llama 2 is a model similar in spirit to ChatGPT except that it can be run locally and is smaller in scale. See the Appendix for some additional details.

This procedure identifies several hundred additional layoffs beyond the roughly 7,000 previously identified. Interestingly, Llama 2 does appear capable of identifying layoffs in sentences that use a wide variety of linguistic devices. However, manual inspection shows a modest but non-negligible fraction of false positives. Given the complexities of handling these cases we do not include the Llama 2-identified layoffs in our series, though we see this as a promising avenue for future work. In any case, the experiment suggests that the iterative ranking procedure does capture the lion's share of layoffs, even if there is room for improvement.

3 Results

Figure 1 shows the monthly time series of layoffs as measured by 8-K filings. The solid black line includes all items 2.05, and the items 7.01 and 8.01 which match our string comparisons. While item 2.05 accounts for most of the relevant filings, the other items constitute about 16 percent of the total. The time series follows the business cycle, with both the Great Recession and the Covid shock clearly visible. The series also show a pronounced increase in early 2023, discussed more below. Interestingly, the other items (red dashed line) spike sharply during Covid.

Overall, while the 8-K series clearly show an increase during Covid the increase is not as large as other layoff indicators. Figure 2 compares the 8-K layoff announcements with initial UI claims and WARN notices. All series are indexed to average 100 in 2006 to make comparisons easier. In addition, we seasonally adjust the 8-K series using X-13 (there is a small degree of seasonality in the raw data, with more layoff announcements coming in

[6] We use the 7B chat-tuned version from https://huggingface.co/meta-llama/Llama-2-7b-chat-hf, run locally.

June and October than other months.) The first panel shows clearly that UI claims and WARN notices rose far more during Covid than 8-K layoff announcements. The SEC did allow for some leniency in filing early in the pandemic but a more important factor might be the unique nature of the layoffs. The pandemic saw a large wave of temporary layoffs, which likely include little or no severance payments. These kinds of layoffs may not be reported as often to the SEC, reducing the counts of 8-Ks. The second panel shows the same series, but suppresses the 2020 data for WARN and UI claims so the dynamics outside of the pandemic are easier to see. All three series move in lockstep at the onset of the Great Recession. Towards the end of the Great Recession, WARN and 8-Ks move down sharply, while initial UI claims take many years to fall. The dynamics around the Great Recession suggest that 8-K filings may have been a timely indicator of both the business cycle peak and the trough.

While counts of layoff events can be informative, it is also useful to track the number of workers affected. This is not trivial in the 8-K data as firms describe the magnitude of layoffs in different ways, though when there are specific numbers it is generally either a count of workers, or a percentage of the firm's workforce. The procedure searches for a quantity appearing in proximity to a quantifier ("reduc-", "decreas-", "lower-", "eliminat-", or "terminat-") and a word related to layoffs ("layoff", "headcount", or "workforce"). For example, if the 8-K mentioned either "We decreased our headcount by 10" or "The firm experienced a workforce reduction of 5%", our algorithm would return 10 and 5%, respectively. If a percentage is found, we multiply the percentage by the lagged value of total employee counts, obtained from Compustat. If a cardinal amount is mentioned, we use this value for the number of layoffs. If multiple quantities are captured within the 8-K, we sum the values to obtain the number of layoffs at the firm-filing level. We seasonally adjust this series as well.

Figure 3 shows the three month trailing moving average of 8-K workers counts (red line), as well as the other relevant series (which are not averaged.) Again, all are normalized by their 2006 averages to facilitate comparisons. We use a three month moving average because the worker count can be extremely volatile month to month. While the three month average still displays significant noise, the business cycles are also readily apparent. Interestingly, the 8-K worker count rose by far more in the Great Recession than any of the other series. During Covid it also increased by a factor comparable to initial UI claims, which (as discussed above) were inflated by temporary layoffs. While the 8-K worker count series may be noisy, it also appears to react strongly to the business cycle.

4 Industry Composition

In this section we examine how the layoff announcements are distributed across sectors and specific industries. SEC filings include the SIC industry code of the filing firm. As other agencies use NAICS codes instead, we convert the SIC codes to NAICS to facilitate comparison to other data. Where there is not a one-to-one correspondence we assign equal weight to each successor NAICS code. Our main comparisons are to firm count shares from the Census's Statistics of U.S. Businesses (SUSB) data, establishment count shares from the Business Dynamics Statistics (BDS), and firm count shares from Compustat. SUSB and BDS cover essentially the universe of employers in the U.S. (both private and public), while Compustat only covers firms that are publicly traded and thus regulated by the SEC. These comparisons complement the existing literature on differences between the publicly traded universe and comprehensive administrative data.[7]

Figure 4 shows the shares of firms, establishments, and layoff events across NAICS sectors in 2015. We focus on 2015 because this is the most recent year for which SUSB has statistics based on firm industry, rather than establishment industry, making it most comparable to the firm-level industry codes in Compustat and the SEC data.

Several things are worth noting. First, the SUSB firm shares and the BDS establishment shares (dark magenta and purple, respectively) are broadly similar. This suggests that

[7] For recent work see Tito (2019), Decker and Williams (2023), and Flynn and Ghent (2023).

the establishment distribution across industries is a reasonable proxy for the firm distribution. Second, some sectors are overrepresented in the publicly-traded universe. Information (NAICS 51), finance and insurance (NAICS 52), and especially manufacturing (NAICS 31-33) all have higher firm shares in Compustat as compared to SUSB. Third, layoff events from SEC filings are even more concentrated in manufacturing than the Compustat firm counts. Note that similar analyses for other years (not shown) suggest these broad patterns are generally stable over time.

The manufacturing sector shares in Figure 4 are particularly striking. Figure 5 focuses on the 3-digit subsectors within manufacturing, but still plots their share of total businesses or layoffs (i.e. not their share of manufacturing businesses/layoffs). Nearly every subsector is overrepresented in the publicly traded data, but chemicals (NAICS 325) and computers & electronics (NAICS 334) stick out as being particularly overrepresented in Compustat and the SEC filings. Note that while layoffs are very high in NAICS 334 in 2015, this pattern is not consistent across years. Interestingly, within the chemicals subsector, it is pharmaceutical manufacturing (NAICS 3254) that accounts for much of the gap between the SUSB/BDS and Compustat/8-Ks. To put it in perspective, according to SUSB about 4% of firms were in manufacturing in 2015, and about 0.03% were in pharmaceutical manufacturing. According to Compustat, about 37% of publicly-traded firms were in manufacturing, and about 11% of publicly-traded firms were in pharmaceutical manufacturing. The latter amounts to overrepresentation by a factor of more than 300. It seems that some factors specific to the pharmaceutical industry—perhaps the regulation of the products or the need to fund research—lead to more firms in that industry going public.

We focus on representativeness in terms of firm counts, since (unweighted) firm counts are the best point of comparison for layoff events. It is worth noting that most of the patterns described above also hold for employment shares too, though the overrepresentation issues are somewhat more moderate. Part of the reason is that while pharmaceutical firms appear to go public disproportionately, public pharmaceutical firms are also very small: median employment is only about 50 workers. Thus they tend to contribute less to employment-weighted statistics, though their small size may also make them more vulnerable to layoff-inducing shocks. The large number of very small firms manufacturing may help explain why the sector has so many layoffs reported in 8-Ks. Reweighting the layoff series is left for future work, though below we take a step in that direction by dropping some the industries that may be the most idiosyncratic.

5 Post-Pandemic Dynamics

Returning to Figure 1, the 8-K derived layoffs series has a notable peak in early 2023, a period when technology sector layoffs and recession fears were making headlines. The 8-K layoffs series also remained elevated throughout 2023, relative to its pre-pandemic level (Figure 3 shows similar patterns for the count of 8-K laid off workers.) This is interesting because other layoffs series did not change much over this period. Table 1 shows the percentage change in various series relative to their 2019 average. The first column covers 2022Q4 through 2023Q1—the period when the 8-K series spiked. The second row covers the remainder of 2023. 8-K layoff events rose 146 percent in 2022Q4-2023Q1 relative to their 2019 level, and remained 82 percent above their 2019 level for the rest of 2023. In contrast, initial UI claims, JOLTS-reported layoffs, and WARN notices (discussed more below) were all essentially flat.

We can investigate this pattern at the industry level, using the SIC industry codes that are part of the forms (here we use SIC codes instead of NAICS codes because we are not linking to other data sources, and we avoid the complications of bridging codes.) It turns out that much of the 2023 rise in 8-K layoffs can be accounted for by two industry groups: computer programming & data processing (SIC 737) and drugs (SIC 283). The former covers many high-tech companies, and the latter corresponds to the pharmaceutical industry discussed earlier. These industry groups generally account for a large number of layoffs, but also saw particularly large increases (in percentage terms) in late 2022 and early 2023. Taken together these industry groups accounted for more than 65 percent of the increase in 8-K layoffs. However, dropping these industries and redoing the calculations (row 2 of Table 1), we find that the purged series still rises considerably in late 2022 and 2023, a pattern at odds with the other indicators. We leave this for future investigation, but note that public companies are different from the universe of firms along many dimensions, including industry (as noted above), size and age.

6 Comparison with WARN Notices

The Worker Adjustment and Retraining Notification (WARN) Act requires that firms with 100 or more workers notify their employees—and the public—in advance of certain layoff events. Notification must be given if (1) a plant is closing that would displace 50 or more workers, (2) 500 or more workers are to be laid off from an establishment, or (3) 50-499 workers are to be laid off and the total is at least 33 percent of the location's workforce, for firms with more than 100 workers. The law requires that notification be given 60 days before the layoffs occur, though some states have additional notification requirements. While WARN notices cover many layoffs, some significant layoffs may not trigger a notice. For example, Okta's February 2023 layoff of 300 workers only amounted to five percent of their workforce.[8] As a result, WARN notices may be better at capturing plant closure than other layoffs.[9]

There are clear advantages and disadvantages to both WARN and the SEC filings. Importantly, WARN covers both private and public companies, and is only triggered by layoffs occurring the U.S. In contrast, layoffs only show up in the SEC filings if the firm is public and if the layoffs meet the bar for material disclosure. They also include staff in other countries. However, SEC filings can cover layoffs that fall below the employment threshold for WARN.

[8] See here.

[9] See the discussion in Krolikowski and Lunsford (2024) and GAO (2019).

We can compare the size of layoffs announced in 8-Ks to the WARN reporting thresholds to get a better sense of the likely overlap in filings. Firms often report the size of the layoff in the 8-K, either as a percent of employment or as the number of workers affected. The 8-K can be linked to information on firm employment in Compustat (which is ultimately derived from other SEC filings.) With the firm's employment level and the layoff as a percent of employment in hand, we can judge if the event satisfies the WARN act's thresholds.

There are many caveats to this exercise. In addition to the coverage issues mentioned above Compustat employment figures are only available annually and are thought to have some degree of measurement error. Also, WARN notices relate to employment and layoffs at a single establishment whereas SEC filings are for the firm as a whole. Finally the extraction of layoff quantities from the 8-Ks relies on NLP techniques that can introduce error. Nonetheless, we believe it is informative to examine the results.

Figure 6 shows the density of layoffs announced in 8-Ks, in grayscale. The distribution is centered around firms with 1,000-2,000 workers, with layoffs of less than 10 percent. The red region with dashed borders shows the area where WARN notices would be required.[10] It is evident that most 8-K announced layoffs would not be picked up by the WARN system. In this sample only about 23 percent of the layoffs meet the WARN filing threshold (assuming that no plant closings were involved.) This is consistent with the belief that WARN is better at capturing plant closings than other mass layoff events. Though no discontinuity is visible on the WARN filing threshold, it is an open question whether firms strategically choose the level of layoffs to avoid the mandated delay the WARN act requires.

6.1 Firm-level linking

WARN notices are useful in part because they can give signal well in advance of the actual layoffs: at least 60 days. The timing of 8-K announcements relative to the date of the actual

[10] There are four distinct regions in terms of the binding conditions for WARN notices. Between zero and 100 workers, no notices are required (except for plant closings). Between 100 and 150 workers, all layoffs of at least 50 workers must be reported. Between 150 and 1,500 workers, all layoffs above 33 percent of employment must be reported. Finally, above 1,500 workers all layoffs above 500 workers must be reported.

layoffs is less clear. The SEC's rules generally require that an 8-K be filed within four business days of the triggering event. Further, when 8-Ks are filed to satisfy Regulation FD the rules are even stricter, requiring simultaneous disclosure or at most a lag of one business day.[11] While the firm's commitment to a restructuring would generally be considered the triggering event for a filing, there is an exception, in that firms may wait to file an item 2.05 until the affected employees are notified.[12] This means that firms might be able to wait to file the 8-K until they tell the employees, which would presumably be at the same time as any WARN notice is filed. In practice, it is unclear how many firms wait this long. Anecdotally, some 8-Ks make it clear that the laid off employees are being notified simultaneously, while many others are more vague about the timing. Others explicitly state that the specific individuals affected have not yet been determined.[13] This suggests that layoffs are sometimes announced in 8-Ks before corresponding WARN notices are filed, or at least simultaneously.

We test this hypothesis by linking WARN notices to 8-K layoff announcements. Both datasets record the company's name, which we use to build the link. The matching is fuzzy, but our goal is only to get a subset of firms where we are confident of the match. We normalize the names (standardizing case and removing common suffixes like "Inc."), and require that the first several characters match exactly. In addition we require that the firm only have a single layoff event reported in the calendar year to avoid linking distinct layoffs.

The resulting dataset has 285 linked layoffs that appear in both the WARN notices and the 8-K filings. Figure 7 shows a histogram of the date of the 8-K filing, relative to the date of the WARN notice, in weeks. The mass around zero shows that the layoff is often announced in both datasets in the same week: roughly 22 percent of the linked events are reported in the same week in both datasets. Interestingly, there is significant variance as well. About 25 percent of linked layoffs appear in the 8-Ks four weeks or more in advance of the WARN notification. On the other end of the spectrum, about 18 percent of the WARN notices are

[11]See https://www.ecfr.gov/current/title-17/chapter-II/part-243
[12]See Question 109.01 in SEC (2023).
[13]See, for example, the 2023-01-03 filing from Pegasystems Inc

more than four weeks ahead of the 8-K filings. These results suggest that neither data source clearly dominates in terms of timeliness.

It is understandable that 8-Ks would sometimes be filed well in advance of WARN notices, for example, a firm may determine that layoffs are going to happen within six months, but may not have finalized the exact number and location of the layoffs. In this case the firm would file an 8-K immediately, but can wait to file a WARN notice, so long as they file 60 days before the layoffs. Consistent with this story, manually inspecting the 8-Ks filed well before the linked WARN notices shows that firms often expect the process to take several fiscal quarters, more than enough time to wait several weeks and then file WARN notices.

It is less clear what is happening when the WARN notices predate the 8-Ks significantly. After all, if the decision to do layoffs was a material event it should have been reported promptly. Examining the 8-Ks manually shows that a mix of factors are at play. Some of these cases are from the early pandemic, when the SEC allowed additional time to make filings. In other cases, the firm is involved in a complicated restructuring, so it appears that separate layoffs/plant closures are being reported in each data source.

6.2 Regressions

In this section we provide preliminary quantitative evidence on the predictive power of the 8-K data. Our exercise follows the VAR specification of Krolikowski and Lunsford (2024) (which in turn is based on Barnichon and Nekarda (2012)). Specifically, we regress various indicators on the 8-K series along with controls. The dependent variables of interest are the change in the log of the unemployment rate ($\Delta \ln ur_t$), the monthly log level of initial UI claims ($\ln uic_t$), and the log job separation rate $\ln s_t$, which is calculated as in Krolikowski and Lunsford (2024). The controls include two lags each of these variables, plus various combinations of the (change in the) 8-K series and the (change in the) WARN notice series. We run the regressions for the period 2006-2019.

The first four columns of Table 2 have the log change in unemployment as the dependent variable, $\Delta \ln ur_t$. Both the 8-K layoff event counts (Δ8Kevents$_{t-1}$, Δ8Kevents$_{t-2}$) and the laid off work counts (Δ8Kworkers$_{t-1}$, Δ8Kworkers$_{t-2}$) have statistically significant relationships with the dependent variable. These relationships appear weakened when we include the WARN series as well (columns three and four), as may be expected. However, the second lag of Δ8Kevents remains highly statistically significant event when the WARN data are included. Broadly similar patterns are apparent for the other dependent variable, though the association with 8-K seems weakest for the log separation rate $\ln s_t$. Both for $\Delta \ln ur_t$ and $\ln uic_t$ the second lag of Δ8Kevents seems to have the statistically strongest and most robust relationship.

We take these regressions as suggestive evidence that the 8-K data has forecasting signal for important labor market quantities, complementing the existing stock of indicators. Important future tasks include testing out-of-sample performance and performance over different sample periods.

7 Conclusion

SEC filings contain a wealth of information, much of it reported in a timely fashion. We have focused on a particular dimension, the announcement of layoffs in 8-K filings. These announcements track UI claims and WARN notices tightly during the Great Recession, though their behavior during and after the pandemic has deviated from other indicators. While neither WARN notices nor 8-Ks dominate with respect to timeliness, we find suggestive evidence that 8-Ks record a significant number of layoff events that would not be picked up via WARN notices. We also find suggestive evidences the the 8-K series are useful for forecasting, even when WARN data are included. Topics for future research include better understanding the reason 8-K layoff events did not rise more during the pandemic, exploring why the series has been elevated in 2023, and evaluating the utility of series for forecasting/nowcasting recessions as in Berge et al. (2016). In ongoing research we explore the role of severance payments in explaining whether or not layoffs are reported in 8-Ks.

The richness of the data allows for interesting comparisons as well as other topics for future research. Linking to WARN notices lets us better understand what kinds of layoffs are likely captured by each data sources, and the relative timing of the announcements. The industry patterns suggest some unusual dynamics in the pharmaceutical industry, which warrant further investigation. Finally, large language models have proven useful even in our limited experimentation, and may be useful for additional analysis.

References

Barnichon, Regis and Christopher Nekarda, "The Ins and Outs of Forecasting Unemployment: Using Labor Force Flows to Forecast the Labor Market," *Brookings Papers on Economic Activity*, 09 2012, *2012*, 83–131.

Berge, Travis J., Nitish R. Sinha, and Michael Smolyansky, "Which Market Indicators Best Forecast Recessions?," FEDS Notes 2016-08-02, Board of Governors of the Federal Reserve System (U.S.) August 2016.

Bodnaruk, Andriy, Tim Loughran, and Bill McDonald, "Using 10-K Text to Gauge Financial Constraints," *The Journal of Financial and Quantitative Analysis*, 2015, *50* (4), 623–646.

Chen, Tianqi and Carlos Guestrin, "XGBoost: A Scalable Tree Boosting System," in "Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining" KDD '16 Association for Computing Machinery New York, NY, USA 2016, pp. 785–794.

Compustat North America. S&P Global Market Intelligence, Capital IQ Platform, Wharton Research Data Services, http://wrds.wharton.upenn.edu/.

Decker, Ryan A. and Jacob Williams, "A Note on Industry Concentration Measurement," FEDS Notes 2023-02-03, Board of Governors of the Federal Reserve System (U.S.) February 2023.

Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," 2019.

Flynn, Sean J. and Andra Ghent, "Does Main Street Benefit from What Benefits Wall Street?," *Journal of Financial and Quantitative Analysis*, 2023, pp. 1–37.

GAO, "The Worker Adjustment And Retraining Notification Act: Revising the Act and Educational Materials Could Clarify Employer Responsibilities and Employee Rights," GAO-03-1003 2019.

Jung, Boochun, Byungki Kim, Woo-Jong Lee, and Choong-Yuel Yoo, "Corporate Layoffs and Accounting Conservatism," Working Paper 2016-001, KAIST College of Business January 2016.

Krolikowski, Pawel M. and Kurt G. Lunsford, "Advance Layoff Notices and Aggregate Job Loss," *Journal of Applied Econometrics*, 2024, *n/a* (n/a).

Laurion, Henry, "Implications of Non-GAAP Earnings for Real Activities and Accounting Choices," *Journal of Accounting and Economics*, 2020, *70* (1), 101333.

Loughran, Tim and Bill McDonald, "When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks," *The Journal of Finance*, 2011, *66* (1), 35–65.

Loukas, Lefteris, Manos Fergadiotis, Ion Androutsopoulos, and Prodromos Malakasiotis, "EDGAR-CORPUS: Billions of Tokens Make The World Go Round," in Udo Hahn, Veronique Hoste, and Amanda Stent, eds., *Proceedings of the Third Workshop on Economics and Natural Language Processing*, Association for Computational Linguistics Punta Cana, Dominican Republic November 2021, pp. 13–18.

SEC, "Additional Form 8-K Disclosure Requirements and Acceleration of Filing Date," 69 FR 15593 2004.

_, "Exchange Act Form 8-K," https://www.sec.gov/divisions/corpfin/guidance/8-kinterp.htm 2023.

Sufi, Amir, "Bank Lines of Credit in Corporate Finance: An Empirical Analysis," *The Review of Financial Studies*, 2009, *22* (3), 1057–1088.

Tito, Maria D., "Trade in Goods and Services: Measuring Domestic and Export Flows in Buyer-Supplier Data," FEDS Notes 2019-06-26, Board of Governors of the Federal Reserve System (U.S.) June 2019.

Touvron, Hugo, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom, "Llama 2: Open Foundation and Fine-Tuned Chat Models," 2023.

	Percent change relative to 2019	
	2022Q4-2023Q1	2023Q2-2023Q4
8-K Layoffs Events	145.56	81.72
8-Ks, ex. SIC codes 283 and 737	68.78	41.14
Initial UI Claims	-2.21	3.77
JOLTS Layoffs	-9.48	-10.97
WARN Notices	-10.74	7.45

Note: All entries are the percent difference between the average for the specificed period and the average for 2019.
Source: BLS, Department of Labor, SEC, Krolikowski and Lunsford (2024), authors' calculations

Table 1: Layoff Indicators, Change Relative to 2019

	$\Delta \ln ur_t$				$\ln s_t$				$\ln uic_t$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
Δ8Kevents$_{t-1}$	0.038**			0.025		0.132**			0.083		0.045			0.019	
	(0.017)			(0.020)		(0.052)			(0.063)		(0.032)			(0.034)	
Δ8Kevents$_{t-2}$	0.065***			0.056***		0.081*			0.047		0.100***			0.082***	
	(0.015)			(0.018)		(0.046)			(0.052)		(0.029)			(0.031)	
Δ8Kworkers$_{t-1}$		0.063***			0.034		0.064			−0.043		0.123**			0.087
		(0.020)			(0.024)		(0.084)			(0.083)		(0.058)			(0.060)
Δ8Kworkers$_{t-2}$		0.054**			0.033		0.140*			0.061		0.129**			0.102*
		(0.021)			(0.023)		(0.071)			(0.076)		(0.053)			(0.054)
$\Delta\widehat{WARN}_{t-1}$			0.167***	0.120**	0.141***			0.537***	0.481**	0.551***			0.238**	0.197	0.170
			(0.049)	(0.052)	(0.053)			(0.166)	(0.185)	(0.179)			(0.116)	(0.121)	(0.108)
$\Delta\widehat{WARN}_{t-2}$			0.091*	0.032	0.069			0.227	0.087	0.221			0.160	0.084	0.096
			(0.051)	(0.054)	(0.052)			(0.163)	(0.196)	(0.173)			(0.102)	(0.100)	(0.094)
R^2	0.475	0.439	0.456	0.489	0.462	0.807	0.800	0.811	0.816	0.813	0.987	0.986	0.986	0.987	0.987
Baseline model R^2	0.415	0.415	0.415	0.415	0.415	0.796	0.796	0.796	0.796	0.796	0.985	0.985	0.985	0.985	0.985

Note: *** $p<0.01$, ** $p<0.05$, * $p<0.10$. Robust standard errors in parenthesis. Sample period: 2006-2019. All dependent variables are taken from Krolikowski and Lunsford (2024) and use their transformations. All specifications include two lags of $\Delta \ln ur_t$, $\ln s_t$, and $\ln uic_t$ as controls. "Baseline R^2" gives the R^2 of the regression using only these controls.

Table 2: Forecasting Regressions: 8-K Series and WARN Notices

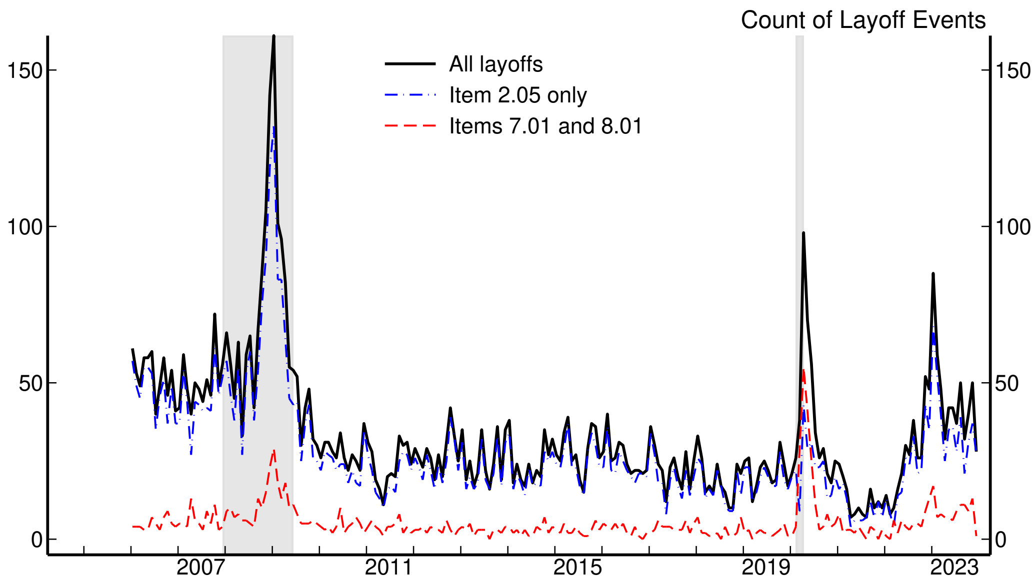

Note: Series are not seasonally adjusted
Source: SEC, Krolikowski and Lunsford (2024), authors' calculations

Figure 1: 8-K Layoff Events

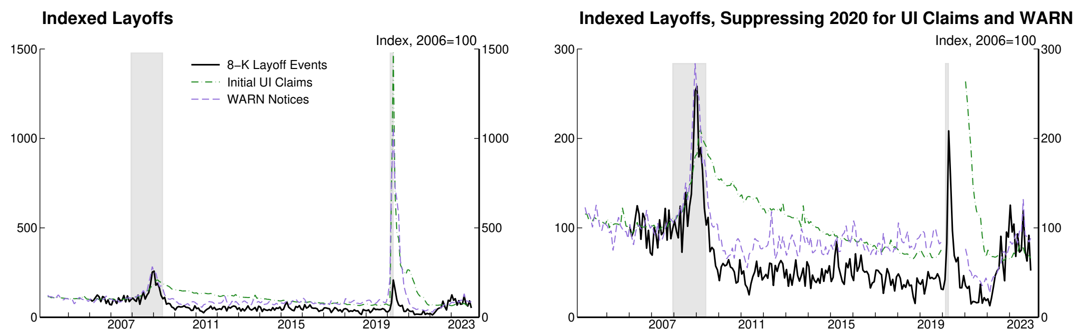

Note: All series are seasonally adjusted. All series are divided through by their 2006 average and multiplied by 100. "WARN Notices" is the measure of layoffs implied by the "WARN factor" from Krolikowski and Lunsford (2024).
Source: Department of Labor, SEC, Krolikowski and Lunsford (2024), authors' calculations

Figure 2: Comparison of 8-K Layoffs, UI Claims, and WARN Notices

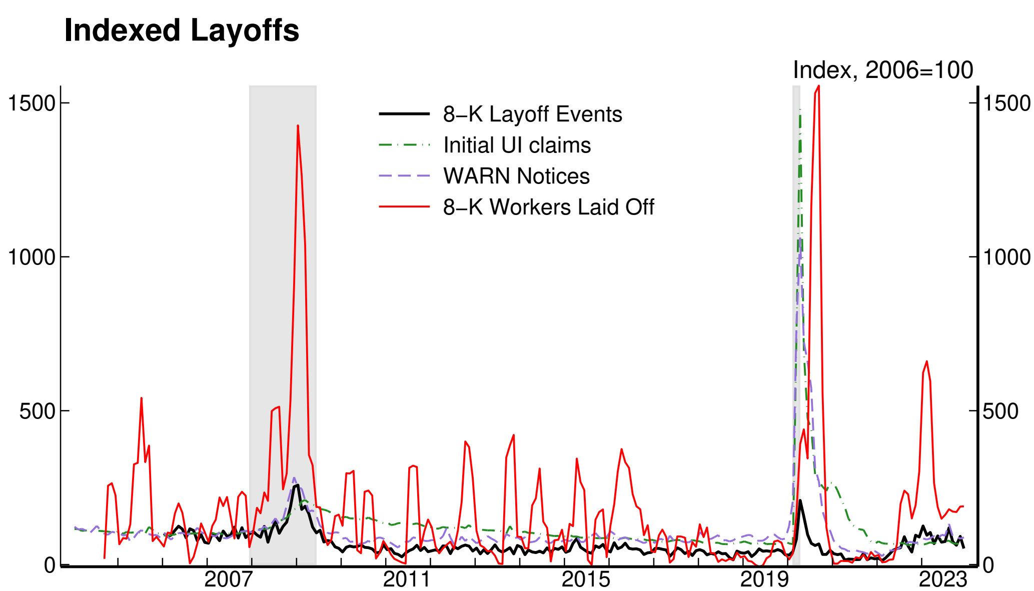

Note: All series are seasonally adjusted. All series are divided through by their 2006 average and multiplied by 100. 8-K Workers Laid Off is a trailing three month moving average.
Source: Department of Labor, S&P, SEC, Krolikowski and Lunsford (2024), authors' calculations

Figure 3: 8-K Workers Laid Off

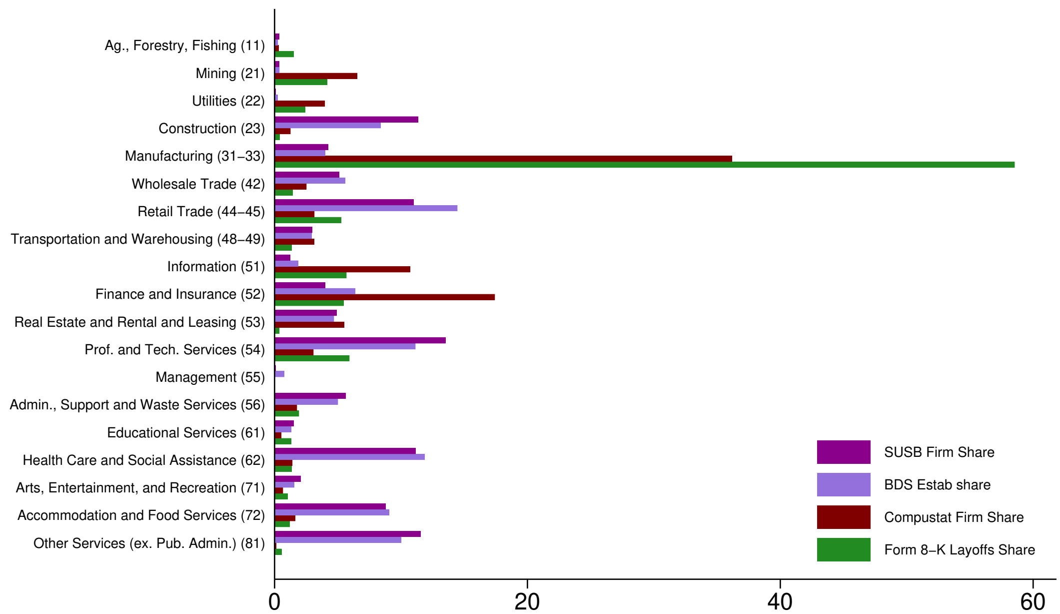

Note: Bars show the percent of the total relevant quantity (business count or layoff event count) in that NAICS sector. SUSB firm-level data is classified by the NAICS of the firm as a whole.
Source: Census Bureau, S&P, SEC, authors' calculations.

Figure 4: Shares of Businesses or Layoff Events Across Sectors, 2015

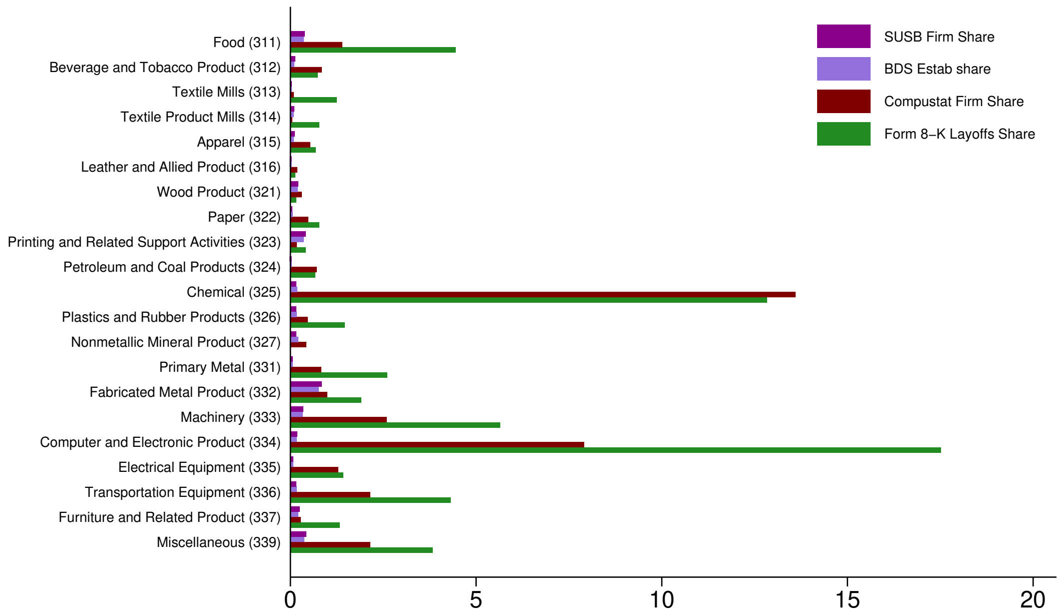

Note: Bars show the percent of the total economy-wide (not manufacturing-specific) relevant quantity (business count or layoff event count) in that three-digit industry. The SUSB firm-level data is classified by the NAICS of the firm as a whole.
Source: Census Bureau, S&P, SEC, authors' calculations.

Figure 5: Manufacturing Industry Shares of Total Businesses or Layoffs, 2015

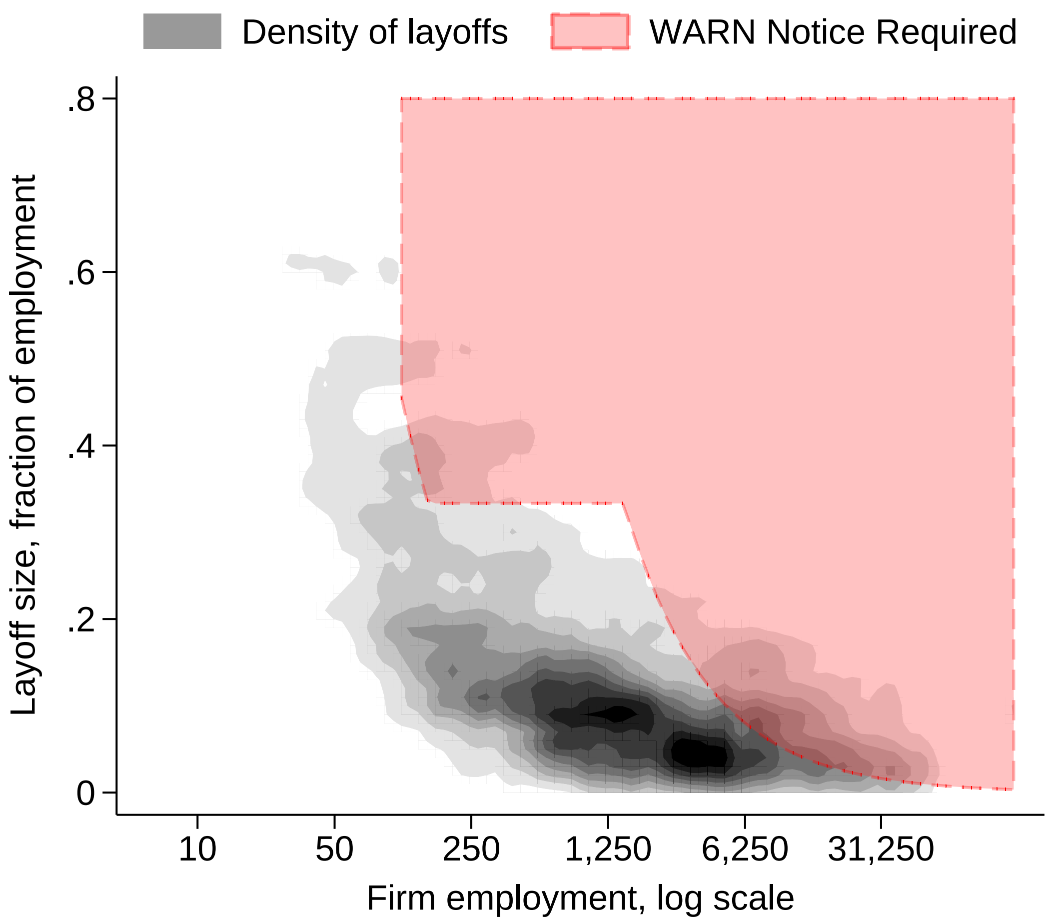

Note: X axis is the count of workers at the firm on a log scale. Y axis is the magnitude of the layoff as a faction of base employment (i.e., the value on the x axis).
Source: SEC, S&P, authors' calculations.

Figure 6: Density of 8-K Reported Layoffs

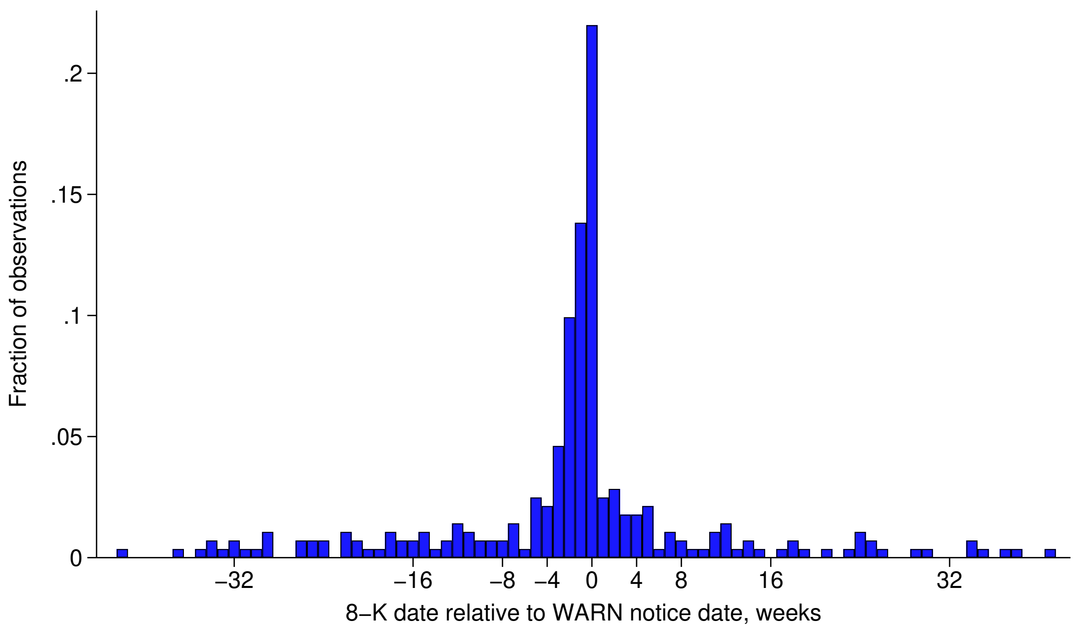

Note: Histogram plots the distribution of the 8-K filing date minus the WARN notice filing date. Each bar is one week.
Source: SEC, https://layoffdata.com/, authors' calculations.

Figure 7: Histogram of 8-K Date Relative to WARN Date

A Llama 2 LLM Details

We using the 7 billion parameter version of Llama 2, specifically the model `meta-llama/Llama-2-7b-chat-hf` from the Huggingface Hub. We use the following prompt:

```
<s>[INST] <<SYS>>
You are a world-class PhD economist, knowledgable on business topics.
<</SYS>>

We are interested whether a firm is announcing layoffs. Layoffs include any
    firings or involuntary elimination of existing employee positions. We do not
     consider temporary furloughs, hiring freezes, or the departure of
    executives to be layoffs, unless they are accompanied by additional firings.

Consider the following excerpt a firm's SEC filings:

    Text: "{x}"

Based on the text, is the firm definitely announcing layoffs or otherwise firing
     existing employees, according to the definition above? Briefly explain your
     reasoning step by step. Towards the end of your answer include either the
    exact phrase "Therefore, yes, the firm is laying people off." or "Therefore,
     no, the passage does not confirm the firm is laying people off."
[/INST]
```

where `x` is a sentence drawn from an 8-K filing. Then we parse the responses for the desired phrases. The hope is that by asking the LLM to reason first and deliver the final answer last it will benefit from chain of thought reasoning.

We leave it for future work to use more sophisticated methods, such as assessing the token probabilities directly instead of relying on the LLM to produce exact strings for matching.