

Bassanini, Andrea; Caroli, Eve; Chaves Ferreira, Bruno; Reberiou, Antoine

Working Paper

Don't Downsize This! Social Reactions to Mass Dismissals on Twitter

GLO Discussion Paper, No. 1254

Provided in Cooperation with:
Global Labor Organization (GLO)

Suggested Citation: Bassanini, Andrea; Caroli, Eve; Chaves Ferreira, Bruno; Reberiou, Antoine (2023) : Don't Downsize This! Social Reactions to Mass Dismissals on Twitter, GLO Discussion Paper, No. 1254, Global Labor Organization (GLO), Essen

This Version is available at:
<http://hdl.handle.net/10419/270427>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Don't Downsize This!

Social Reactions to Mass Dismissals on Twitter¹

Andrea Bassanini
(OECD, GLO and IZA)

Eve Caroli
(LEDa, Université Paris Dauphine, PSL and IZA)

Bruno Chaves Ferreira
(Université Paris Dauphine, PSL)

Antoine Reberioux
(Université Paris Cité, LADYSS)

March 2023

Abstract

We study social reactions to job destructions on Twitter. We use information on large-scale restructuring events announced in the United Kingdom over the period 2013-2018. We match it with data collected on Twitter regarding the number and sentiments of the tweets posted around the time of the announcement and involving the company name. We show that job-destruction announcements immediately elicit numerous and strongly negative reactions. These reactions are almost five times larger than the positive reactions to job-creation announcements. We also show that the negative reactions triggered by job-destruction announcements are associated with significant losses in the market value of the downsizing firms. This suggests that the damage to the firms' image triggered by the negative social reactions following job destructions entail a financial cost for firms that adds to the standard adjustment costs of dismissals.

Keywords: job destructions, firm value, social media, sentiment analysis, mass dismissals.

JEL codes: J63, L82, M21, M51

¹ We are grateful to the European Foundation for the Improvement of Living and Working Conditions for making the Restructuring Events database available to us. We gratefully acknowledge highly valuable assistance with data scraping and sentiment analysis from Yifei Fan, Svitlana Galeshchuk and Abir Jaza from Governance Analytics (PSL IRIS) and research assistance from Anna Farrugia. We are also indebted to Omar Bamieh, Eric Brousseau, Filippo Maria D'Arcangelo, Kevin Geay, Mathilde Godard, Sophie Hatte, Matthias Heinz, Julien Jourdan, Alessandro Saia and participants at seminars at the OECD (Paris) and LEST (Aix-en-Provence) and at the EALE 2021 conference for useful comments and suggestions. All remaining errors are ours. Andrea Bassanini, Eve Caroli and Antoine Reberioux acknowledge financial support from the French National Research Agency for financial support (Grant ANR-19-CE26-0017).

1. Introduction

In this paper, we study social reactions to dismissals on social networks and, more specifically, on Twitter. To date, the literature has essentially focused on the reactions of peer company executives and business analysts. In their seminal paper, Flanagan and O'Shaughnessy (2005) show that having a layoff in the course of the year lowers the rating attributed to a company by the senior executives, outside directors and financial analysts answering *Fortune's* America's Most Admired Corporations (AMAC) survey. This result has been subsequently confirmed by Love and Kraatz (2009) and Schulz and Johann (2018). In contrast, evidence of social reactions to dismissals from outside the business world is only anecdotal (see e.g. Michael Moore's film, *Roger and Me*, released in 1989, which features the strong opposition to the mass dismissals carried out by General Motors in Flint, Michigan). However, reactions in the community and beyond may be strongly negative since large-scale dismissals are known to impose large costs on communities and are often perceived as unfair (Charness and Levine, 2000; Hallock, 2009, Gathmann et al., 2020). When expressed online these reactions may rapidly become viral, since social networks are now a ubiquitous medium of information diffusion (Brady et al., 2017) and the main vehicle of public evaluation.

In the present paper, we show that firms' announcements of mass dismissals generate strong negative reactions involving the company names on Twitter. We focus on Twitter since it is a leading social-media platform and, as such, one of the most important "echo chambers" of public perceptions. Moreover, on Twitter, reactions to information are almost instantaneous, which permits clear identification of the impact of mass-dismissal announcements on social buzz. We show that the negative social reactions triggered by job-destruction announcements are more than four times larger than the positive reactions entailed by job creations'. Moreover, we provide evidence indicating that these negative social reactions have sizeable consequences for firms: we show that even if the job-destruction announcements *per se* have, on average, no effect on firm market value, the negative buzz that they generate on Twitter reduces stock prices, controlling for the characteristics of the events.

We rely on the Restructuring Events database, which is part of the European Restructuring Monitor (ERM) and provides information on announcements of large-scale job destructions (i.e. mass dismissals) and job creations reported on the press for a large number of European Union (EU) companies since the early 2000s. We consider announcements made by companies in the United Kingdom (UK) in 2013-2018. We match each announcement with information collected on Twitter regarding the number and content of the tweets involving the company

name posted during a time period ranging from 45 days before to 10 days after the announcement. We first show that the number of tweets mentioning the company name upsurges after a job-destruction announcement. To make sure that this effect is not the same for any human-resource management decision made by firms, we compare it with the effect of job-creation announcements. We find that the latter also trigger an increase in the number of tweets mentioning the company name, but this increase is significantly smaller than that following job-destruction announcements. As a second step, we focus on the content of the tweets. We carry out a sentiment analysis (see Gentzkow et al., 2019 for a review) based on the VADER (Valence Aware Dictionary and Sentiment Reasoner) lexicon, which is particularly designed to assess sentiment on social media (Hutto and Gilbert, 2014 – see Shapiro and Wilson, 2022, and Shapiro et al., 2021, for other applications in economics). This lexicon attributes a positive or negative score to approximately 7,500 words according to the sentiment they express. We use it to compute the share of negative (resp. positive) words in each tweet. Our analysis shows that the average negativity of the tweets significantly increases following job-destruction announcements. We also provide evidence that the incidence of negative words in the tweets following job destructions increases by a much larger amount than the incidence of positive words in the tweets following job creations. These findings suggest that large-scale job destructions generate a large negative buzz involving the company's name.

Finally, we examine whether, despite the fact that reactions on Twitter are short-lived, these negative social reactions may have noticeable consequences for firms. To do so, we investigate the relationship between these reactions to job-destruction announcements and firm market value. We compute cumulative abnormal returns (*CAR*) for those firms in our database that are listed on the London Stock Exchange. We show that an increase in the overall volume of negative sentiments expressed on Twitter in reaction to job-destruction announcements triggers a reduction in *CAR*, not only in the short run, but also in the medium run.

Our paper relates to three strands of literature. First, it contributes to the literature on the costs borne by firms when downsizing. Economic theory suggests that employment destructions generate adjustment costs (Nickell, 1986; Bertola, 1992). The empirical literature has shown that legal and contractual provisions are key determinants of these costs (Hamermesh, 1995; Kramarz and Michaud, 2010; Boeri and Van Ours, 2013). Remaining workers have also been shown to react to dismissals of their colleagues by reducing their effort and organisational engagement (Datta et al., 2010; Drzensky and Heinz, 2016; Sucher and Gupta, 2018; Heinz et al., 2020), thereby raising unit labour costs. Finally, some studies show that firms refrain from

cutting jobs close to their headquarters and hypothesise that this strategy could aim at avoiding the reputational cost that dismissals of its members may induce with the local community (Landier et al., 2009; Abraham et al., 2014; Bassanini et al., 2017). In the present paper, we provide direct evidence that job-destruction announcements trigger negative reactions on social media which are, in turn, associated with a loss in the firm's market value. This suggests that negative social reactions may generate costs in terms of image for the firm, that add to the more traditional adjustment costs.

Second, our research contributes to the literature on social media and business companies. Social media play a crucial role in product markets. On the one hand, they provide an important source of information for consumers, and are therefore often directly integrated in firms' marketing strategies (Chen and Xie, 2008, Majumdar and Bose, 2019). Rui et al. (2013) as well as Asur and Huberman (2010) show, for example, that the sentiments expressed by Twitter users significantly influence movie sales in the first weeks after movies are released. On the other hand, social media are increasingly used by consumers to get companies to do what they think would be fair (Hendel et al., 2017) and there is evidence in the literature that company bashing on Twitter deteriorates firms' reputation in the long run (Hansen et al., 2018). We speak to this literature by showing that human resource management decisions also entail a significant buzz involving companies' name, which may affect the image of the firm.

Finally, our paper also speaks to the literature on media slant (e.g. Gentzkow and Shapiro, 2010; Durante and Knight, 2012; Puglisi and Snyder, 2015). In particular, our research relates to Heinz and Swinnen (2015) and Friebe and Heinz (2014) who document media slant against dismissals in the written press in Germany. We provide evidence that dismissals also trigger many more reactions than job creations on social media.

The rest of the paper is structured as follows. Section 2 describes the data and presents summary statistics. Section 3 lays out our empirical strategy. Section 4 presents the main empirical results concerning reactions on Twitter to job-destruction and job-creation announcements. Section 5 investigates the relationship between negative reactions to job destructions and firm value. Section 6 concludes.

2. Data

The first dataset that we use is the ERM Restructuring Events database.² It contains factsheets with data on large-scale restructuring events – i.e. job-creation and job-destruction announcements – reported in the principal national newspapers and on TV websites in each EU member state since 2002. We consider restructuring events reported in the UK over 2013-2018.³ The UK is indeed one of the EU countries where Twitter started expanding first: 12 million Britons were already using Twitter in 2013 as compared to 5.6 million French,⁴ for example. 2013 is the first year in which Twitter was massively used in the UK: the number of users increased by 34% with respect to 2012, while annual user growth rates decreased to less than 15% in each subsequent year. All 1,264 restructuring events contained in the database entail either job creations and/or job destructions. For each event, we know the date at which it was officially announced by the firm, as reported in the national press. We also have information on the number of planned job destructions and/or job creations. We drop events for which planned job creations and job destructions are simultaneously positive.

For each event, we scrape from Twitter all tweets the text of which includes the name of the company that announced this event. These tweets are scraped over a time period ranging from 45 days before the announcement to 10 days afterwards.⁵ We drop events corresponding to companies which name can be confused with famous people (e.g. McCain which can also refer to the late US senator John McCain who died during our sample period), with geographical locations (e.g. Oakland International) or with generic expressions (e.g. Call Connection or New Look). We also drop events corresponding to companies that attracted more than ten thousand tweets in several days during the pre-event period (such as Amazon, Google or McDonald's) since Twitter shuts down access when very large numbers of tweets are scraped for a given company over several days in a row. Our database eventually contains 1,046 events, corresponding to 766 companies since some of them announced several events in our time window – see Appendix Table A.1. 51% of these events involve job destructions while 49% involve job creations. As evidenced in Appendix Table A.2, the mean size of job-destruction

² Available at <http://www.eurofound.europa.eu/observatories/emcc/erm/factsheets>.

³ Mass dismissals reported in the ERM Restructuring Events database account for 15% of all dismissals in the UK (including dismissals for personal reasons and individual redundancies); therefore this dataset likely covers most large-scale dismissal events since mass dismissals are only a small proportion of all dismissals – see OECD (2018).

⁴ See <https://www.emarketer.com/Article/More-than-One-Fifth-of-UK-Consumers-Use-Twitter/1010623> for data on the UK and <https://www.emarketer.com/Articles/Print.aspx?R=1009851> for data on France.

⁵ We scrape tweets only up to ten days after the event since reactions on Twitter are short-lived – as will be confirmed by our results. In contrast, we scrape back to 45 days before the event because we want to check that pre-trends are horizontal over a long period of time before the announcement takes place.

events is slightly larger than that of job creations with, on average 469 jobs destroyed in job-destruction events as compared to 402 jobs created in job-creation events.

For each tweet, we have information on the username, the exact date of the post, the number of first-level replies⁶ and the content of the text, including the number of words.⁷ We drop the tweets in which the name of the company appears in the username. Those tweets are indeed likely to have been posted by the companies themselves (Majumdar and Bose, 2019), while we are interested in social reactions to job destructions rather than in information disclosed by companies. Overall, our database contains 11,949,136 tweets.

For each tweet, we compute the number of positive and negative words using the VADER lexicon.⁸ This has been shown to be particularly suited to sentiments expressed in social media (see Hutto and Gilbert, 2014). This lexicon contains a list of about 7,500 words that have been allocated a score on a continuous scale ranging from -4 to +4. This score reflects the intensity of the negative/positive sentiment expressed by the word, with -4 capturing the most negative and +4 the most positive sentiment. We consider as negative words those attracting a strictly negative score and positive words those attracting a strictly positive score. Following Tetlock et al. (2008), for each tweet, we compute the ratio of negative to total words (*RatioNeg*), the ratio of positive to total words (*RatioPos*) and the ratio of the difference between the number of positive and negative words to the sum of positive and negative words (*RatioDiff*). The latter indicator captures the dominant sentiment of the tweet: from complete negativity (-1) to complete positivity (+1).

We aggregate all tweet-level information at the day-by-announcement level. By doing so, we obtain a database containing, for each event, the daily number of tweets and the daily number of first-level replies. For each event, we also obtain the average values of *RatioNeg*, *RatioPos* and *RatioDiff* for each day. As evidenced in Appendix Table A.3, the number of tweets mentioning the name of a company that announced job destructions is larger after the announcement than before. The daily average of the number of tweets indeed amounts to 235 between $t = 0$ and $t = +10$ – where t indicates the time distance in days from the date of the announcement –, as compared to 192 between $t = -45$ and $t = -1$. A much smaller difference is

⁶ A first-level reply to a tweet is a direct reply to that tweet. On each tweet, Twitter provides a counter of these direct replies which are thereby dated at the date of the posting of that tweet. The counter does not include indirect replies, i.e. replies to replies to a tweet.

⁷ To transform hashtags into words, we rely on the Ekphrasis text-processing tool developed by Baziotis et al. (2017) which performs tokenisation, word normalisation, word segmentation and spell correction, using word statistics from Wikipedia along with 330 million tweets.

⁸ https://github.com/cjhutto/vaderSentiment/blob/master/vaderSentiment/vader_lexicon.txt.

observed for job creations with a daily average number of tweets of 211 after the announcement as compared to 207 before – see Appendix Table A.4. The average number of first-level replies to those tweets increases following both job destructions and job creations. As could be expected, the ratio of negative to total words (*RatioNeg*) increases following job-destruction announcements (from 3% to 4.7%) while the ratio of positive to total words (*RatioPos*) decreases (from 7.2% to 5.8%). *RatioDiff*, which captures the overall positivity of the sentiments expressed, also goes down from 0.422 to 0.130 – see Appendix Table A.3. As regards job creations, the ratios of negative and positive words to total words remain almost stable after the announcement has taken place – see Appendix Table A.4. In contrast, *RatioDiff* increases – although by a much smaller amount than its decrease following job destructions – from 0.509 to 0.572. These descriptive statistics suggest that the (negative) buzz following job-destruction events is of a larger magnitude than the (positive) buzz induced by job-creation announcements. In the next section, we lay out our methodology to investigate this relation.

3. Empirical Model

Our objective is to estimate the impact of job-destruction announcements on the number of tweets and first-level replies, on the one hand, and on the sentiments expressed by those tweets, on the other hand, and compare it with the effect of job-creation announcements. We then investigate how these reactions affect firms' financial performance in the short and medium run.

3.1 Twitter reactions to job-destruction and job-creation announcements

We first consider the change in the number of tweets and in sentiments following the announcement of a job-destruction (resp. job-creation) event. To do so, we estimate the following model:

$$Y_{j,c} = \sum_{t=-45}^{-4} \beta_t D_{j,c}^t + \sum_{t=-2}^{+10} \beta_t D_{j,c}^t + \mu_j + \gamma_c + \varepsilon_{j,c} \quad (1)$$

where $Y_{j,c}$ is the outcome variable for event j at calendar day c – i.e., alternatively, the number of daily tweets, the number of first-level replies to daily tweets, the daily average ratio of negative to total words (*RatioNeg*), the daily average ratio of positive to total words (*RatioPos*) and the daily average ratio of the difference between the number of positive and negative words to the sum of positive and negative words (*RatioDiff*).

$D_{j,c}^t = \mathbf{1}\{c - c_{0j} = t\}$ is a distance-to-event indicator taking value 1 t days away from the date c_{0j} of the announcement of event j . We use $t = -3$ (i.e. 3 days before the announcement) instead of $t = -1$ as the reference point to allow for the fact that the date of the event may be somewhat imprecise. The ERM Restructuring Events database indeed indicates the date of the official communication of job destruction or job creation by the firm, as reported in the press. However, in a number of cases, the news leaked in the press in advance, without official communication from company executives. Our data contain several examples of leakages one or two days before a company spokesman confirmed the announcement to the press. For example, the downsizing announced by Marks & Spencer on September 5th 2016 leaked on Skynews (and other newspapers, such as the Herald Scotland)⁹ two days before. Similarly, the downsizing of 150 persons in Brighton by Legal and General in September 2017 was officially confirmed by a company spokesperson two days after the staff received an email and the information leaked to the press.¹⁰ In such cases, taking $t = -1$ as a reference would wrongly underestimate the true effect of the announcement.

μ_j is an event fixed effect capturing time-invariant unobservable heterogeneity across firms and events. γ_c are calendar-day fixed effects accounting for the fact that Twitter activity is unevenly distributed over the days of the week¹¹ and that other events may take place on the same days as the job destruction or job creation announcements we consider. Standard errors $\varepsilon_{j,c}$ are clustered at the event level.

The model presented in equation (1) is a standard dynamic two-way fixed effect model with a single, binary treatment in a staggered design (see de Chaisemartin and D'Haultfœuille, 2023, for a simple discussion). Therefore, coefficients obtained by estimating equation (1) by OLS would suffer from the contamination problem evidenced by Sun and Abraham (2021) and Callaway and Sant'Anna (2021). Hence, they could not be interpreted as unbiased estimates of the dynamic causal effects of the announcement.¹² In order to address this issue we use the DID_1 estimator proposed by de Chaisemartin and D'Haultfœuille (2021, 2023). In a model

⁹ <https://www.heraldscotland.com/news/14721772.marks-spencer-to-cut-around-500-jobs-at-its-head-office/>

¹⁰ <https://www.theargus.co.uk/news/15525648.shock-as-legal-and-general-moves-150-jobs-to-midlands/>

¹¹ The number of tweets is indeed lower on Fridays and even more so on weekends, while job-destruction and job-creation announcements are rare during weekends.

¹² Sun and Abraham (2021) show that, if the true treatment effects vary across each post-treatment period t , the standard two-way fixed effect OLS estimator for β_t is a weighted sum of two terms: one is a weighted average of the true treatment effects across all groups at t (therefore the true treatment effect if the latter is homogeneous across groups); and the other is a weighted sum of the treatment effects at all other periods. Consequently, the expected value of the standard two-way fixed effect OLS estimator would not be the true treatment effect.

with a single binary treatment and a staggered design, it is indeed identical to the estimator of Callaway and Sant'Anna (2021) when not-yet-treated observations are used as a control group.

Our key parameter of interest is β_0 , which yields the magnitude of the change in the outcome variable from $t = -3$ to $t = 0$. The identifying assumption is that the level of the outcome variable would have remained constant in the absence of job-destruction or job-creation announcement. This is plausible to the extent that this level does not show any significant daily variation over the month and a half preceding the event: we indeed test that none of the $\beta_{t=-45,\dots,-4}$ in equation (1) are significantly different from 0. In the case of job destructions, we expect $\hat{\beta}_0$ to be positive and significantly different from zero when the dependent variables are the number of daily tweets, the number of first-level replies or the ratio of negative to total words. In contrast, we expect it to be significantly negative for *RatioPos* and *RatioDiff* to the extent that these variables take higher values when the sentiments expressed in the tweets get more positive. In the case of job creations, we expect $\hat{\beta}_0$ to be positive (even though of potentially smaller magnitude than for job destructions) when the dependent variables are the number of daily tweets, the number of first-level replies, the ratio of positive to total words, as well as for *RatioDiff*. By contrast, we expect it to be negative for *RatioNeg*. We also check that $\hat{\beta}_0$ is significantly different from $\hat{\beta}_{-1}$ and $\hat{\beta}_{-2}$, to make sure that the change we observe at the announcement date is larger than any potential change taking place one or two days before.

3.2 Comparing reactions to job-creation and job-destruction announcements

As a second step, we want to gauge the significance of any differential effect of job-destruction vs job-creation announcements on the buzz involving the company name on Twitter.

When considering quantitative outcomes such as the number of tweets and the number of first-level replies, we estimate the following equation:

$$Y_{j,c} = \sum_{t=-4}^{-4} \beta_t D_{j,c}^t + \sum_{t=-2}^{+10} \beta_t D_{j,c}^t + \sum_{t=-45}^{-4} \beta'_t D_{j,c}^t * JD_j + \sum_{t=-2}^{+10} \beta'_t D_{j,c}^t * JD_j + \mu_j + \gamma_c + \varepsilon_{jc} \quad (2)$$

where JD_j is a dummy variable equal to 1 when event j is a job-destruction announcement and 0 when event j is a job-creation announcement. Our main parameter of interest is β'_0 which

captures the relative effect of a job-destruction event with respect to a job-creation event at time $t = 0$ as compared to $t = -3$. If job destructions generate more buzz on Twitter than job creations, $\hat{\beta}'_0$ will be positive and significant. To make sure that our comparison is meaningful, we check that pre-event trends are not significantly different between job destructions and job creations, i.e. $\hat{\beta}'_t$ for $t \leq -4$ are not significantly different from 0.

When considering the sentiments expressed by the tweets, we expect negative sentiments to increase following job-destruction but decrease following job-creation announcements, and positive sentiments to vary the other way round. So, estimating equation (2) on our standard negativity and positivity indicators would yield trivial results: $\hat{\beta}'_0$ would be positive for the ratio of negative to total words and it would be negative for the ratio of positive to total words and for *RatioDiff*. To obtain a meaningful comparison, we consider new, modified dependent variables. The first one is defined as follows:

$$Y_{j,c}^{(1)} = JD_j * RatioNeg_{j,c} + JC_j * RatioPos_{j,c}$$

where JC_j is a dummy variable equal to 1 when event j is a job-creation announcement and 0 when event j is a job-destruction announcement. $Y_{j,c}^{(1)}$ is therefore equal to the ratio of negative to total words in case of job destructions and to the ratio of positive to total words in case of job creations. When estimating equation (2) for $Y_{j,c}^{(1)}$, $\hat{\beta}'_0$ will be positive and significant if the negative reactions triggered by job-destruction announcements are more important than the positive reactions triggered by job-creation announcements.

Similarly, we define $Y_{j,c}^{(2)}$ as the ratio of positive to total words in case of job destructions and to the ratio of negative to total words in case of job creations. When estimating equation (2) for this outcome, $\hat{\beta}'_0$ will be negative and significant if job destructions reduce positive sentiments more than job creations reduce negative sentiments.

The third and last variable we consider is:

$$Y_{j,c}^{(3)} = -JD_j * RatioDiff_{j,c} + JC_j * RatioDiff_{j,c}$$

where, as defined in Section 2:

$$RatioDiff_{j,c} = \frac{Positive\ Words_{j,c} - Negative\ Words_{j,c}}{Positive\ Words_{j,c} + Negative\ Words_{j,c}}$$

$Y_{j,c}^{(3)}$ is therefore equal to the excess number of positive words (standardized by the sum of positive and negative words) in the case of job creations and to the excess number of negative words (standardized by the sum of positive and negative words) in the case of job destructions. When estimating equation (2) on $Y_{j,c}^{(3)}$, $\hat{\beta}'_0$ will be positive and significant if negative reactions in case of job destructions are of greater magnitude than positive reactions in case of job creations.

4. Social Reactions on Twitter

4.1 Number of Tweets and First-Level Replies

We first estimate equation (1) for the number of tweets. The regression coefficients and standard errors are plotted against the time distance to the announcement – see Figure 1. $t = -3$ is taken as a reference, hence the reported coefficient is equal to 0. The time window we represent on the graphs is restricted to $[-10; +10]$ since none of the coefficients estimated before $t = -10$ are significant at conventional levels.

As evidenced in Panel A, the estimated daily number of tweets mentioning a company's name is not significantly different between $t = -3$ and any other date preceding a job-destruction announcement. In contrast, the estimated number of tweets almost triples on the day of the announcement: +303, significant at the 1% level, as compared to an average of 178 at $t = -3$. It is still significantly higher at $t = +1$ and $t = +2$, although by a smaller amount (+131 and +71, respectively). It finally goes back to its baseline value three days after the job destructions were announced. Reactions on Twitter are much more limited in case of job creations – see Panel B of Figure 1. The increase in the number of tweets is indeed almost five times smaller than for job destructions: +69 at $t = 0$, as compared to an average of 192 at $t = -3$. It is also more short-lived since the number of tweets goes back to its initial level the day after the announcement.

Although the difference in reactions to job-destruction and job-creation events is quite stark from a graphical point of view, to make sure that it is statistically significant, we estimate equation (2) with the number of tweets as the dependent variable. The results are presented in Table 1. Trends appear to be parallel in the pre-event period: at all pre-event dates ($t < 0$), the estimated difference in the number of tweets across job destructions and job creations is never significantly different from that at $t = -3$. In contrast, at the time of the announcement ($t =$

0), the number of tweets increases much more in case of job destructions than in case of job creations with a difference of 234 tweets, significant at the 1% level. The gap between the number of tweets posted in reaction to job-destruction and job-creation announcements remains significantly larger at $t = +1$ (+109, significant at the 5% level) and at $t = +2$ (+66, significant at the 10% level) than at the reference date, while it becomes insignificant at later dates. As a consequence, we only report the estimated coefficients until $t = +5$.¹³ Overall, these results indicate that job-destruction announcements trigger many more reactions on Twitter than job-creations'.

We then re-estimate equation (1) for the number of first-level replies to the tweets posted during our time window. These replies are an additional indicator of buzz on social media, as they are the most direct way for users to engage in a conversation. The number of first-level replies is provided by Twitter on the initial tweet. As a consequence, replies are mechanically dated the day when the initial tweet was posted even if they have been posted later on. This is why the observed dynamics of the number of replies is particularly short-termed. As can be seen on Panel A of Figure 2, the estimated number of first-level replies increases significantly on the day of the announcement of a job destruction: +86 as compared to an average of 50 replies at $t = -3$. The number of replies is also higher at $t = +1$, but the difference with $t = -3$ is not significant at conventional levels. This suggests that only the tweets posted on the day of the event systematically trigger a particularly large number of replies. As regards job-creation events, interestingly, they do not trigger any significant increase in the number of replies.¹⁴

When comparing reactions to job destructions and job creations, column (2) of Table 1 shows that our identifying assumptions are satisfied when considering the number of replies. Pre-event trends are parallel, since for all dates before the announcement ($t < 0$), the estimated difference in the number of replies across job destructions and job creations is never significantly different from that at $t = -3$. In contrast, at the time of the announcement, the difference in the number of replies triggered by job destructions and job creations significantly increases with respect to

¹³ Although we take $t = -3$ as the reference date, we also check that the difference in the number of tweets at $t = 0$ is significantly higher than at $t = -1$ and $t = -2$. This is actually the case at the 1% level for both dates, as indicated at the bottom of Table 1.

¹⁴ As shown on Figure 2 – Panel B, the estimated number of replies increases by 48 at $t = +6$ (with respect to $t = -3$). Yet, this change is not significant at conventional levels since it is driven by one single outlier unrelated to the event, i.e. a tweet posted by a famous singer 6 days after company [NAME] announced some job creations (the company name has been suppressed for confidentiality reasons). This tweet, which stated: "I'll do more later! I'm just going into [NAME] to get somethin for dinner!", indeed attracted 23,648 replies... If we remove this tweet from our sample, the number of replies at $t = +6$ becomes very similar to its value at $t = -3$ – see Panel B of Appendix Figure B.1.

$t = -3$ (+78). This gap fades away the day after the event, consistent with the time dynamics evidenced on Figure 2. These findings confirm that job creations trigger much fewer reactions on Twitter than job destructions do.

4.2 Sentiment Analysis

Our first indicator of sentiments is the average daily ratio of negative to total words (*RatioNeg*). As can be seen on Figure 3 – Panel A, following a job-destruction announcement, *RatioNeg* doubles: +3 percentage points (significant at the 1% level), as compared to a baseline level of 3.2% at $t = -3$. The ratio of negative to total words remains significantly higher than its reference value up to 4 days after the event: at $t = +4$, it is still 0.7 percentage points higher than at $t = -3$ (significant at the 5% level). This first piece of evidence suggests that the official announcement of a job-destruction episode entails a strongly negative buzz on Twitter. Consistently, the estimated ratio of positive to total words (*RatioPos*) decreases following a job-destruction announcement. As evidenced in Figure 4 – Panel A, it goes down by 2.3 percentage points at $t = 0$ (significant at the 1% level), as compared to a baseline level of 7% at $t = -3$ and the deviation from the reference value lasts for several days. This pattern is confirmed by our third indicator, *RatioDiff* – see Sections 2 and 3. As evidenced in Figure 5 – Panel A, the estimated excess number of positive over negative words (standardized by the sum of positive and negative words) decreases following a job-destruction announcement: -53 percentage points (significant at the 1% level) as compared to a baseline level of 41.5% at $t = -3$. This reduction is long lasting by Twitter standards since it only fades away 6 days after the event. The analysis carried out in Section 4.1 suggested that job-destruction announcements trigger a substantial buzz on Twitter. This new series of findings shows that this buzz is strongly negative and that the increase in negativity (and the reduction in positivity) lasts longer than the increase in the number of tweets or replies. This means that even when the number of tweets and replies comes back to its pre-event level, for several days their content remains significantly more negative (less positive) than it used to be.

Our analysis also suggests that job creations trigger positive social reactions, although more limited in size and time than the reactions triggered by job destructions. Panel B of Figure 3 shows that the ratio of negative to total words decreases following a job-creation announcement (-0.50 percentage points at $t = 0$ as compared to 2.5% at $t = -3$, significant at the 1% level) and that this reduction disappears on the following day. Consistent with this decrease in negativity, the ratio of positive to total words significantly increases at $t = 0$ as compared to

$t = -3$ (+0.9 percentage points at $t = 0$ as compared to 8.1% at $t = -3$, significant at the 1% level), and so does the excess number of positive words, *RatioDiff* (+8.7 percentage points at $t = 0$, as compared to a baseline value of 52.4% at $t = -3$) – see Panels B of Figures 4 and 5.

Visual inspection of Figures 3, 4 and 5 suggests that the intensity of negative sentiments expressed in reaction to job destructions is much stronger than the intensity of positive sentiments expressed in reaction to job creations. To make sure that these differences are statistically significant, we estimate equation (2) for our three sentiment-based outcome variables – see Table 2. Column (1) provides the results for $Y^{(1)}$ which captures the ratio of negative to total words in case of job destructions and the ratio of positive to total words in case of job creations. By the time of the announcement, job destructions trigger a much larger increase in the ratio of negative to total words than the increase in the ratio of positive to total words triggered by job creations – with the difference being significant at the 1% level. Similarly, the reduction in the share of positive words following a job-destruction announcement is significantly larger than the reduction in negative words following a job-creation announcement – see column (2) of Table 2. Finally, as shown in Table 2 – column (3), the increase in the standardised excess number of negative over positive words is significantly larger upon announcement of a job-destruction event than the increase in the excess number of positive over negative words triggered by a job-creation announcement. Overall, our results support the idea that there exists an asymmetry between job-destruction and job-creation announcements: the negative buzz generated by the former is significantly stronger – about five times larger – than the positive buzz generated by the latter.

These findings suggest that job destructions entail a strong negative buzz involving the company's name on a largely-used social network, namely Twitter. They also provide evidence that job destructions and job creations are not symmetric in this respect, since the former generate a much fiercer firm bashing than any positive comments triggered by the latter.

4.3 Robustness checks

A first concern regarding the reactions to job destruction announcements is that they could be overestimated if firms disclose additional information likely to generate reactions on Twitter at the time they announce job destructions. If this were the case, the negative buzz we capture could be partly triggered by reactions to this additional information rather than by job destruction announcements only. To make sure that this is not the case, we manually inspected our 532 job destruction announcements and identified those for which the company

contemporaneously disclosed information such as brand sale, the fact that the firm went into administration, the backing of a restructuring plan by creditors, etc. 135 such events were identified and dropped from our database, representing 25% of the initial events. We then re-estimated equation (1) for job destructions. Our results are unchanged,¹⁵ thus suggesting that even when additional information is disclosed, job destructions are the most salient information to which Twitter users react.

Regarding our sentiment analysis, a concern is that measuring sentiments based on a simple word count is, of course, crude. As a robustness check, we use a score constructed using the contextual VADER algorithm of sentiment analysis.¹⁶ This algorithm aggregates the scores assigned in the lexicon taking into account punctuation, negation, capital letters, the use of intensifiers – such as e.g. "extremely", "much", "really" – and the three preceding words, so that "not so great" is coded as negative whereas "great" is coded as positive. For each tweet, the VADER score generated by the algorithm is standardised so that values range from -1 (extremely negative sentiment) to +1 (extremely positive sentiment). As for the word-count measures, we aggregate it at the day-by-announcement level. As evidenced in Appendix Figure B2 – Panel A, the positive sentiments expressed by the tweets collapse following a job-destruction announcement: the estimated VADER score decreases by 0.166 at $t = 0$ as compared to a baseline value of 0.148 at $t = -3$ (with the change being significant at the 1% level). This reduction is particularly long lasting since it is still significant at the 5% level at $t = +7$. As can be seen on the chart, some anticipation takes place at $t = -1$ and $t = -2$, but tests for the difference across the regression coefficients are significant at the 1% level, showing that the reduction taking place at $t = 0$ is significantly larger than at prior dates. In contrast, job creation announcements do not trigger any significant change in the VADER score – see Panel B of Figure B2.

Another worry could be that the reactions we observe in case of job destructions could be due to a small number of Twitter users, namely those who have lost their job, or their relatives. To tackle this issue, we estimate equation (1) separately for two different outcome variables computed at the day-by-event level: first, the total number of users – i.e. the number of distinct user identifiers who have tweeted the company name – and second, the number of multiple

¹⁵ At $t = 0$, the number of tweets increases by +266, the number of first-level replies by +69, the ratio of negative to total words by +3.1 percentage points, the ratio of positive to total words decreases by -2.4 percentage points and the excess number of positive over negative words (standardized by the sum of positive and negative words) decreases by -51 percentage points, all significant at the 1% level.

¹⁶ <https://github.com/cjhutto/vaderSentiment>. This score has been used in the literature in economics by Shapiro and Wilson (2022) and Shapiro et al (2021), among others.

users defined as users who have tweeted mentioning the company name and for whom we have tweets mentioning at least another company for which we have an event in the same quarter in our dataset. The latter are individuals who have tweeted multiple events in a short period of time and are hence unlikely to have been personally affected by all of them. As evidenced in Appendix Figure B.3 – Panel A, the estimated total number of users more than doubles following job-destruction announcements: +227, significant at the 1% level. Interestingly, out of those 227 additional users, 54% are multiple users. As can be seen on Panel B of Figure B.3, the estimated number of multiple users indeed increases by +122 following job destruction announcements, thus suggesting that a large proportion of those individuals who reacted to a job destruction announcement also tweeted in reaction to other events during the same quarter and are hence unlikely to do so only because they have been personally impacted.

Finally, the average size of job-destruction events is a little larger than that of job creations (469 vs 402) – see Appendix Table A.2. One could therefore be concerned that social reactions to job destructions could be more massive and intense just because these events are of slightly larger scale. In fact, the increase in the number of tweets on the day of the event appears to be strongly correlated with the size of this event.¹⁷ To make sure that this is not driving our results, we re-estimate equation (2) for our five main outcome variables on the sample of events whose size is above the median size of all events – see Appendix Table A.5. We do so because Panel C of Appendix Table A.2 shows that, within the group of events with size above the median, the average size of job creations is slightly larger than that of job destructions (783 vs 743). So, any differential impact of job-destruction announcements on social reactions in this group would not be due to job destructions being of larger scale than job creations. As evidenced in Appendix Table A.5, above-median job destructions trigger more reactions than above-median job creations: they attract altogether more tweets, more replies and more acute sentiments – see columns (1) to (5). This suggests that the difference we find across job-destruction and job-creation announcements is not driven by differences in size across the corresponding events.

5. Negative Buzz and Firm Value

In the previous section, we have shown that job-destruction announcements induce a negative buzz on Twitter, and that this buzz is of a larger magnitude than the positive buzz induced by job-creation events. Our findings also indicate that the increase in the number of tweets and

¹⁷ The correlation is 0.32, significant at the 1% level.

first-level replies following job-destruction announcements is short lived – one to three days – and that the increase in negativity (and the decrease in positivity) of the sentiments expressed by these tweets does not last for more than six days. One may wonder why being subject to negative social reactions for such a short period of time should be of concern to firms.

In what follows, we investigate the relationship between the negative social reactions to job destructions and firm market value. To do so, we compute daily stock returns for those companies in our dataset that are listed on the London Stock Exchange. Daily percentage market returns (including dividends) and market indexes were obtained from the Thomson Reuters Eikon database. When dividends were missing, we collected them manually using the Investing.com website. Eventually, information on market returns was available for 54 firms with job-destruction events, corresponding to 99 events.

We estimate the change in firm market value that can be attributed to the negative buzz following job destructions, using the cross-sectional event-study approach that is standard in the literature in finance (MacKinlay, 1997; Kothari and Warner, 2007; Ahern and Dittmar, 2012). We first estimate normal returns on a time window preceding the event. More specifically we model daily stock returns as:

$$R_{i,j,t} = \alpha_{i,j} + \beta_{i,j} * MR_t + \varepsilon_{i,j,t} \quad (3)$$

where $R_{i,j,t}$ is the return on the stock market value of firm i , before event j , at time distance t from the day of the event. MR_t denotes the benchmark market return, as computed using the FTSE 250 index.¹⁸ Following Farber and Hallock (2009), we use a time window ranging from 60 to 31 calendar days prior to the event (denoted $[-60;-31]$) to estimate equation (3). We then compute abnormal returns as:

$$AR_{i,j,t} = R_{i,j,t} - \hat{R}_{i,j,t}$$

where $\hat{R}_{i,j,t}$ are the normal returns predicted using equation (3). To distinguish short-term versus longer-term financial performance, we cumulate abnormal returns over a one-week period $[-2;+4]$, and alternatively, over a 3-month period $[-2;+90]$ around the event. Starting two days before the event allows us taking into account the existence of potential leakages of information before the event is officially announced (Ahern and Dittmar, 2012). After trimming

¹⁸ The FTSE 250 index is the capitalization-weighted index consisting of the 101st to the 350th largest companies listed on the London Stock Exchange (LSE). As compared to the FTSE 100 which focuses on the 100 largest companies, this index is usually considered a better indicator of the financial performance of UK firms since it is composed of less internationally focused companies (Alkhatib and Harasheh, 2018; Law et al., 2020; Rosini and Shenai, 2020).

the top and bottom 1% cumulative abnormal returns (*CAR*) to avoid that our results be driven by outliers, both regression samples contain 97 events.

As evidenced in Appendix Table A.6 below, the mean value of 1-week *CAR* following a job destruction announcement is slightly positive. This is consistent with results in the recent literature¹⁹ and suggests that job-destruction events do not systematically reduce firm market value.

In order to capture the negative buzz, we take into account both the volume of activity on Twitter and the negativity of the sentiments expressed in each tweet. To do so, we define the net negative buzz as the difference between the daily number of negative and positive words in all tweets mentioning the company name. To account for potential leakages of information before the event and consistent with what we do for *CAR*, we compute the average of this indicator over the [-2;0] time window and compare it with its pre-leakage value at $t = -3$. We denote this difference as $\Delta NegBuzz$.

We then estimate the relationship between $\Delta NegBuzz$ and *CAR*. Thus doing, we compare the gap in cumulative returns between the market and firms that have been subject to strong bashing after job destruction on the one hand, and the gap in cumulative returns between the market and firms that have been subject to fewer negative reactions, on the other hand.²⁰ To make sure that our estimates do not capture the impact of the size of job destructions, we control for the number of jobs destroyed in each of the events we consider. Our regressions also include industry, year, month and weekday dummies. Results are shown in Table 3. As evidenced in column (1), an increase in the net negative buzz following job destructions is associated with a decrease in the one-week *CAR*, significant at the 1% level: when $\Delta NegBuzz$ increases by 10% of one standard deviation (+44), the *CAR* decrease by 0.24 percentage points, i.e. 4.4% of one standard deviation. This negative effect of firm bashing on cumulative abnormal returns does not only affect firms in the short run. It actually reinforces in the medium run, since an increase in

¹⁹ Although the oldest literature on the impact of job destructions on firm stock value suggested that *CAR* were on average negative at times of job-destruction announcements (e.g. Abowd et al., 1990; Hallock, 1998), Farber and Hallock (2009) show that since the 1990s, the average *CAR* tend to be insignificantly different from zero at times of job-destruction announcements.

²⁰ The implicit identifying assumption is that there is no simultaneous upsurge in bashing on average in the market against firms that have not announced any job destruction. This appears to be a reasonable assumption since we find that at the dates at which a firm announces a job destruction in our sample, on average $\Delta NegBuzz$ is not significantly different from 0 for firms that have not made an announcement: $\Delta NegBuzz$ is equal to +74 (significant at conventional levels) for firms which made an announcement and to an insignificant -4 for firms that did not.

$\Delta NegBuzz$ by 10% of a standard deviation generates a reduction in the 3-month CAR by 0.89 percentage points, i.e. 16.2% of one standard deviation – see column (2) of Table 3.

One may worry that these effects could be driven by the fact that low CAR trigger negative buzz rather than the opposite. This could occur, for example, if investors complain on Twitter when financial performance is bad. However, if this were the case, the association between lower CAR and negative buzz should be observed at any time, i.e. also before the job-destruction event takes place. To check for this possibility, we run a placebo experiment. For all "true" events, we consider fictive events taking place at all possible dates between 6 weeks and 2 weeks before the "true" event. For each fictive event, we redefine $t = 0$ as the date of this event and compute the corresponding CAR using the same methodology as above. The distribution of these CAR is similar to the one presented in Appendix Table A.6 with a mean of 0.48, a first quartile of -2.16 and a third quartile equal to 2.22 for 1-week CAR and a mean of 3.52, a first quartile of -13.09 and a third quartile equal to 16.70 for 3-month CAR . We regress these CAR on the corresponding $\Delta NegBuzz$ computed for the fictive event, controlling for industry, year, month and weekday dummies. As shown in Appendix Table A.7, none of these estimates are significant at conventional levels. This suggests that, although CAR are of the same order of magnitude in "normal times" – i.e. outside job-destruction announcements – as around the date of the "true" event, they are unrelated to a potential negative buzz in periods in which no job destruction is announced. This allows us ruling out the fact that low CAR trigger negative buzz.

6. Conclusion

We have shown in this paper that job-destruction announcements trigger numerous and strongly negative reactions on one of the most important social media, i.e. Twitter. On the day of the announcement, the number of tweets and first-level replies sharply increases (it almost triples in both cases) as does the negativity of the posted tweets: the ratio of negative to total words doubles while the ratio of positive to total words and the excess number of positive over negative words significantly decreases. The reactions to job-destruction events are systematically larger than the reactions to job-creation events. The latter trigger fewer tweets and replies and weaker changes in sentiments: the increase in the positivity of the tweets following job creations is significantly smaller than the increase in the negativity of the tweets following job destructions. All in all, our study documents a strong asymmetry between job-

destruction and job-creation announcements in terms of buzz and sentiments expressed by individuals on social media.

For the subset of these job-destruction events involving companies that are listed on the London Stock Exchange, we have also shown that negative buzz, as measured by Twitter reactions to job-destruction announcements, triggers a significant negative adjustment in the market value of the downsizing companies. This effect appears substantial in the short run and, even more so in the medium run. An increase in our indicator of negative buzz on Twitter by 10% of one standard deviation indeed generates an estimated reduction by 4.4% (resp. 16.2%) of one standard deviation in one-week (resp. three-month) cumulative abnormal returns. This suggests that the damage to the firms' image triggered by the negative social reactions following job destructions entail a financial cost for firms that adds to the standard adjustment costs of dismissals.

Our findings raise the question of whether and to what extent managers anticipate and/or subsequently adapt their downsizing plans to this form of social pressure. Do they sometimes give up restructuring projects for fear of reputational and associated financial losses? If job destructions are announced and social reactions are particularly fierce, do they reduce the scope of their original downsizing plan? Answering these questions would advance our knowledge on how social media may influence the decisions of economic actors, in general, and human resource decisions of employers, in particular.

Our study covers a period of relative economic stability in which mass dismissals were the exception, rather than the rule. The social reactions to job destructions may, of course, be quite different in times of large-scale economic crisis such as the massive recession induced by the Covid-19 pandemics. The public could indeed consider job destructions as more justified or, alternatively, reactions could be stronger insofar as they resonate with general negative sentiment. This is likely to be another promising avenue for future research.

References


- Abraham, Filip, Tim Goesaert and Josef Konings. 2014. Staying Home or Moving Away? Restructuring Efforts within Multinational Enterprises, *The World Economy*, 37(6): 765-782.
- Abowd, John, George Milkovich, and John Hannon. 1990. The effects of human resource management decisions on shareholder value, *Industrial and Labor Relations Review*, 43(3): 203s-236s.
- Ahern, Kenneth and Amy Dittmar. 2012. The Changing of the Boards. The Impact on Firm Valuation of Mandated Female Board Representation, *The Quarterly Journal of Economics*, 127(1): 137-197.
- Alkhatib, Akram and Murad Harasheh. 2018. Performance of Exchange Traded Funds during the Brexit Referendum: An Event Study, *International Journal of Financial Studies*, 6(3), <https://doi.org/10.3390/ijfs6030064>.
- Asur, Sitaram and Bernardo Huberman. 2010. Predicting the Future with Social Media. Proceedings - 2010 IEEE/WIC/ACM International Conference on Web Intelligence, WI 2010. 1. 10.1109/WI-IAT.2010.63.
- Bassanini, Andrea, Giorgio Brunello and Eve Caroli. 2017. Not in my Community: Social Pressure and the Geography of Dismissals, *Journal of Labor Economics*, 35(2): 429-483.
- Baziotis, Christos, Nikos Pelekis and Christos Doukeridis. 2017. DataStories at SemEval-2017 Task 4: Deep LSTM with Attention for Message-level and Topic-based Sentiment Analysis, in *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, Association for Computational Linguistics, Vancouver, Canada: 747-754.
- Bertola, Giuseppe. 1992. Labor Turnover Costs and Average Labor Demand, *Journal of Labor Economics*, 10(4): 389-411.
- Boeri, Tito and Jan Van Ours. 2013. *The Economics of Imperfect Labor Markets*, Princeton University Press.
- Brady, William, Julian Wills, John Jost, Joshua Tucker and Jay Van Bavel. 2017. Emotion shapes the diffusion of moralized content in social networks, *Proceedings of the National Academy of Sciences*, 114(28): 7313-7318.
- Callaway, Brantly and Pedro Sant'Anna. 2021. Difference-in-differences with multiple time periods, *Journal of Econometrics*, 225: 200-230.
- de Chaisemartin, Clément and Xavier d'Haultfœuille. 2021. Difference-in-differences estimators of intertemporal treatment effects. arXiv preprint arXiv:2007.04267.
- de Chaisemartin, Clément and Xavier d'Haultfœuille. 2023. Two-Way Fixed Effects and Differences-in-Differences with Heterogeneous Treatment Effects: A Survey. *Econometrics Journal*, forthcoming.
- Charness, Gary and David Levine. 2000. Are Layoffs Acceptable? Evidence from a Quasi-Experiment. *Industrial and Labor Relations Review*, 53(3): 381-400.
- Chen, Yubo and Jinhong Xie. 2008. Online Consumer Review: Word-of-Mouth as a New Element of Marketing Communication Mix, *Management Science*, 54(3): 477-491.
- Datta, Deepak, James Guthrie, Dynah Basuil and Alankrita Pandey. 2010. Causes and Effects of Employee Downsizing: A Review and Synthesis, *Journal of Management*, 36(1): 281-348.

- Drzensky, Frank and Matthias Heinz. 2016. The Hidden Costs of Downsizing, *The Economic Journal*, 126(598): 2324-2341.
- Durante, Ruben and Brian Knight. 2012. Partisan Control, Media Bias, and Viewer Responses: Evidence from Berlusconi's Italy, *Journal of the European Economic Association*, 10(3): 451-481.
- Farber, Henry and Kevin Hallock. 2009. The Changing Relationship between Job Loss Announcements and Stock Prices: 1970-1999, *Labour Economics*, 16(1) : 1-11.
- Flanagan, David and Kenneth O'Shaughnessy. 2005. The Effect of Layoffs on Firm Reputation, *Journal of Management*, 31(3): 445-463.
- Friebel, Guido and Matthias Heinz. 2014. Media Slant Against Foreign Owners: Downsizing, *Journal of Public Economics*, 120, 97-106.
- Gathmann, Christina, Ines Helm and Uta Schönberg. 2020. Spillover Effects of Mass Layoffs, *Journal of the European Economic Association*, 18(1): 427-468.
- Gentzkow, Matthew and Jesse Shapiro. 2010. What Drives Media Slant? Evidence from U.S. Daily Newspapers, *Econometrica*, 78(1), 35-71.
- Gentzkow, Matthew, Bryan Kelly and Matt Taddy. 2019. Text as Data, *The Journal of Economic Literature*, 57(3): 535-574.
- Hallock, Kevin. 1998. Layoffs, top executive pay and firm performance, *American Economic Review*, 88(4): 711-723.
- Hallock, Kevin. 2009. Job Loss and the Fraying of the Implicit Employment Contract, *Journal of Economic Perspectives*, 23(4): 69-93.
- Hamermesh, Daniel. 1995. Labour Demand and the Source of Adjustment Costs, *The Economic Journal*, 105(430): 620-634.
- Hansen, Nele, Ann-Kristin Kupfer and Thorsten Hennig-Thurau. 2018. Brand crises in the digital age: The short- and long-term effects of social media firestorms on consumers and brands, *International Journal of Research in Marketing*, 35(4): 557-574.
- Heinz, Matthias and Johan Swinnen. 2015. Media slant in economic news: A factor 20, *Economics Letters*, 132(C): 18-20.
- Heinz, Matthias, Sabrina Jeworek, Vanessa Mertins, Heiner Schumacher and Matthias Sutter. 2020. Measuring the Indirect Effects of Adverse Employer Behavior on Worker Productivity – A Field Experiment, *Economic Journal*, forthcoming.
- Hendel, Igal, Saul Lach and Yossi Spiegel. 2017. Consumers Activism: the Cottage Cheese Boycott, *RAND Journal of Economics*, 48(4): 972-1003.
- Hutto, Clayton and Eric Gilbert. 2014. VADER: A Parsimonious Rule-based Model for Sentiment Analysis of Social Media Text, *Proceedings of the Eighth International AAAI Conference on Weblogs and Social Media*, Ann Arbor, Michigan.
- Kothari, S P and Warner. 2007. Econometrics of event studies, in Eckbo B. Espen ed, *Handbook of Corporate Finance: Empirical Corporate Finance*, Volume 2, pp. 3-36, North-Holland.
- Kramarz, Francis and Marie-Laure Michaud. 2010. The Shape of Hiring and Separation Costs in France, *Labour Economics*, 17(1): 27-37.
- Landier, Augustin, Vinay Nair and Julie Wulf. 2009. Trade-offs in staying close: corporate decision making and geographic dispersion, *Review of Financial Studies*, 22(3): 1119-1148.

- Law, Cherry, Laura Cornelsen, Jean Adams, Tarra Penney, Harry Rutter, Martin White and Richard Smith. 2020. An analysis of the stock market reaction to the announcements of the UK Soft Drinks Industry Levy, *Economics and Human Biology*, 37: 1-10.
- Love, Geoffrey and Matthew Kraatz. 2009. Character, Conformity or the Bottom Line? How and Why Downsizing Affected Corporate Reputation, *Academy of Management Journal*, 52(2): 314-335.
- MacKinlay, Craig. 1997. Event studies in economics and finance, *Journal of Economic Literature*, XXXV: 13–39.
- Majumdar, Adrija and Indranil Bose. 2019. Do tweets create value? A multi-period analysis of Twitter use and content of tweets for manufacturing firms, *International Journal of Production Economics*, 216 : 1-11.
- Nickell, Stephen. 1986. Dynamic Models of Labour Demand, in Ashenfelter Orley and David Card, *Handbook of Labor Economics*, North Holland, Volume 1, Chapter 9: 473-524.
- OECD. 2018. Back to work: Lessons from nine country case studies of policies to assist displaced workers, in *Employment Outlook*, Paris: OECD, Chapter 4: 125-186.
- Puglisi, Riccardo and James Snyder Jr. 2015. The Balanced US Press, *Journal of the European Economic Association*, 13(2): 240–264.
- Rosini, Lucrezia and Vijay Shenai. 2020. Stock returns and calendar anomalies on the London Stock Exchange in the dynamic perspective of the Adaptive Market Hypothesis: A study of FTSE100 & FTSE250 indices over a ten year period, *Quantitative Finance and Economics*, 4(1): 121-147.
- Rui, Huaxia, Yizao Liu and Andrew Whinston. 2013. Whose and what chatter matters? The effect of tweets on movie sales. *Decision Support Systems*, 55: 863-870.
- Schulz, Ann-Christine and Sarah Johann. 2018. Downsizing and the fragility of corporate reputation: An analysis of the impact of contextual factors, *Scandinavian Journal of Management*, 34(1): 40-50.
- Shapiro, Adam and Daniel Wilson. 2022. Taking the Fed at its Word: A New Approach to Estimating Central Bank Objectives using Text Analysis, *Review of Economic Studies*, 89(5): 2768–2805.
- Shapiro, Adam, Moritz Sudhof and Daniel Wilson. 2021. Measuring News Sentiment, *Journal of Econometrics*, forthcoming.
- Sucher, Sandra and Shalene Gupta. 2018. Layoffs That Don't Break Your Company: Better Approaches to Workforce Transition, *Harvard Business Review*, 96(3): 122-129.
- Sun, Liyang and Sarah Abraham. 2021. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects, *Journal of Econometrics*, 225(2): 175-199.
- Tetlock, Paul, Maytal Saar-Tsechansky and Sofus Macskassy. 2008. More Than Words: Quantifying Language to Measure Firms' Fundamentals, *The Journal of Finance*, 63(3): 1437-1467.


Figures

Figure 1. Number of Tweets




Notes: This graph reports estimated coefficients on time-to-event dummies (in days) from $t=-10$ to $t=+10$, with $t=-3$ used as a reference, obtained by estimating equation (1) using the DID_1 estimator of de Chaisemartin and D'Haultfœuille (2021; 2023). The dependent variable is the daily number of tweets mentioning the company name in the text but not in the username. Regressions also include time-to-event dummies (in days) from $t=-45$ to $t=-11$, as well as event and calendar-day fixed effects. Standard errors are clustered at the event level. Error bars correspond to 95% confidence intervals.

Figure 2. Number of First-Level Replies




Notes: This graph reports estimated coefficients on time-to-event dummies (in days) from $t=-10$ to $t=+10$, with $t=-3$ used as a reference, obtained by estimating equation (1) using the DID_1 estimator of de Chaisemartin and D'Haultfœuille (2021; 2023). The dependent variable is the daily number of first-level replies. Regressions also include time-to-event dummies (in days) from $t=-45$ to $t=-11$, as well as event and calendar-day fixed effects. Standard errors are clustered at the event level. Error bars correspond to 95% confidence intervals.

Figure 3. Ratio of Negative to Total Words (*RatioNeg*)




Notes: This graph reports estimated coefficients on time-to-event dummies (in days) from $t=-10$ to $t=+10$, with $t=-3$ used as a reference, obtained by estimating equation (1) using the DID_1 estimator of de Chaisemartin and D'Haultfœuille (2021; 2023). The dependent variable is the ratio of negative to total words. Regressions also include time-to-event dummies (in days) from $t=-45$ to $t=-11$, as well as event and calendar-day fixed effects. Standard errors are clustered at the event level. Error bars correspond to 95% confidence intervals.

Figure 4. Ratio of Positive to Total Words (*RatioPos*)



Notes: This graph reports estimated coefficients on time-to-event dummies (in days) from $t=-10$ to $t=+10$, with $t=-3$ used as a reference, obtained by estimating equation (1) using the DID_1 estimator of de Chaisemartin and D'Haultfœuille (2021; 2023). The dependent variable is the ratio of positive to total words. Regressions also include time-to-event dummies (in days) from $t=-45$ to $t=-11$, as well as event and calendar-day fixed effects. Standard errors are clustered at the event level. Error bars correspond to 95% confidence intervals.

Figure 5. *RatioDiff*



Notes: This graph reports estimated coefficients on time-to-event dummies (in days) from $t=-10$ to $t=+10$, with $t=-3$ used as a reference, obtained by estimating equation (1) using the DID_1 estimator of de Chaisemartin and D'Haultfœuille (2021; 2023). The dependent variable is the ratio of the difference between the number of positive and negative words to the sum of positive and negative words. Regressions also include time-to-event dummies (in days) from $t=-45$ to $t=-11$, as well as event and calendar-day fixed effects. Standard errors are clustered at the event level. Error bars correspond to 95% confidence intervals.

Tables

Table 1 – Differential impact of job-destruction vs job-creation announcements on the number of tweets and first-level replies

	(1) Number of Tweets	(2) Number of First Level Replies
Job Destruction* t_0	233.58*** (64.38)	78.26*** (27.16)
Job Destruction* t_{+1}	108.84** (53.28)	30.34 (22.48)
Job Destruction* t_{+2}	66.34* (36.69)	8.92 (12.23)
Job Destruction* t_{+3}	3.21 (32.65)	-16.86 (16.11)
Job Destruction* t_{+4}	-16.48 (44.13)	-8.46 (17.12)
Job Destruction* t_{+5}	-27.06 (25.91)	-4.29 (10.65)
Job Destruction* t_{-1}	-31.70 (24.65)	-9.35 (9.33)
Job Destruction* t_{-2}	5.79 (25.99)	2.01 (9.06)
Job Destruction* t_{-3}	<i>ref</i> -	<i>ref</i> -
Job Destruction* t_{-4}	10.23 (17.85)	-1.81 (7.20)
Job Destruction* t_{-5}	8.31 (22.23)	4.62 (9.80)
Job Destruction* t_{-6}	-41.51 (31.27)	-14.24 (12.23)
Job Destruction* t_{-7}	-45.93 (29.72)	-20.01 (12.60)
Job Destruction* t_{-8}	-25.72 (42.36)	-12.02 (16.19)
Job Destruction* t_{-9}	5.68 (39.66)	-10.91 (14.44)
Job Destruction* t_{-10}	16.15 (43.20)	-3.65 (12.02)
Job Destr.*($t_0 - t_{-1}$) : p-value	0.0000	0.0005
Job Destr.*($t_0 - t_{-2}$) : p-value	0.0002	0.0037
Observations	58,576	58,576

Notes: Models are estimated using the DID_1 estimator of de Chaisemartin and D'Haultfœuille (2021; 2023). In column (1), the dependent variable is the daily number of tweets mentioning the company name in the text but not in the username. In column (2), the dependent variable is the daily number of first-level replies to the previous tweets, dated at the date at which each tweet was posted. *Job Destruction (JD)* is a dummy variable taking value 1 if the event is a job-destruction announcement and 0 if it is a job-creation announcement. t_n denotes the time distance n in days from the announcement date t_0 . Regressions include event dummies, time-to-event dummies (in days) from t_{-45} to t_{+10} , with t_{-3} used as a reference, as well calendar-day dummies. Interactions terms between *JD* and time-to-event dummies from t_{-45} to t_{-11} on the one hand and between *JD* and time-to-event dummies from t_{+6} to t_{+10} on the other hand are included in the regressions although not reported here. Standard errors clustered at the event level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 2 – Differential impact of job-destruction vs job-creation announcements on sentiments as measured by word count

	Y ⁽¹⁾ JobDestr*RatioNeg + JobCreat*RatioPos	Y ⁽²⁾ JobDestr*RatioPos + JobCreat*RatioNeg	Y ⁽³⁾ – JobDestr*RatioDiff + JobCreat*RatioDiff
Job Destruction*t ₀	0.021*** (0.004)	-0.018*** (0.004)	0.441*** (0.049)
Job Destruction*t ₊₁	0.019*** (0.005)	-0.015*** (0.004)	0.378*** (0.049)
Job Destruction*t ₊₂	0.012** (0.005)	-0.012*** (0.004)	0.276*** (0.051)
Job Destruction*t ₊₃	0.012** (0.005)	-0.010** (0.004)	0.215*** (0.057)
Job Destruction*t ₊₄	0.008 (0.006)	-0.011** (0.004)	0.141** (0.062)
Job Destruction*t ₊₅	0.001 (0.005)	-0.010** (0.005)	0.111* (0.061)
Job Destruction*t ₋₁	0.003 (0.005)	-0.002 (0.004)	0.045 (0.054)
Job Destruction*t ₋₂	0.008 (0.005)	-0.007 (0.004)	0.040 (0.050)
Job Destruction*t ₋₃	ref -	ref -	ref -
Job Destruction*t ₋₄	0.001 (0.005)	-0.000 (0.004)	0.013 (0.052)
Job Destruction*t ₋₅	0.004 (0.005)	-0.002 (0.004)	0.060 (0.055)
Job Destruction*t ₋₆	-0.000 (0.005)	0.001 (0.004)	0.053 (0.055)
Job Destruction*t ₋₇	0.002 (0.006)	0.002 (0.004)	0.005 (0.055)
Job Destruction*t ₋₈	-0.006 (0.006)	-0.002 (0.004)	-0.005 (0.054)
Job Destruction*t ₋₉	-0.006 (0.005)	-0.001 (0.004)	0.026 (0.050)
Job Destruction*t ₋₁₀	-0.006 (0.006)	-0.001 (0.005)	-0.020 (0.057)
Job Destr.*(t ₀ – t ₋₁) : p-value	0.0002	0.0000	0.0000
Job Destr.*(t ₀ – t ₋₂) : p-value	0.0060	0.0124	0.0000
Observations	40,350	40,350	37,573

Notes: Models are estimated using the DID_1 estimator of de Chaisemartin and D'Haultfœuille (2021; 2023). In column (1), the dependent variable $Y^{(1)}$ is equal to the ratio of negative to total words in case of job destructions and to the ratio of positive to total words in case of job creations. In column (2), the dependent variable $Y^{(2)}$ is equal to the ratio of positive to total words in case of job destructions and to the ratio of negative to total words in case of job creations. In column (3), the dependent variable $Y^{(3)}$ is equal to the excess number of positive words (standardized by the sum of positive and negative words) in case of job creations and to the excess number of negative words (standardized by the sum of positive and negative words) in case of job destructions. *Job Destruction (JD)* is a dummy variable taking value 1 if the event is a job-destruction announcement and 0 if it is a job-creation announcement. t_n denotes the time distance n in days from the announcement date t_0 . Regressions include event dummies, time-to-event dummies (in days) from t_{-45} to t_{+10} , with t_{-3} used as a reference, as well as calendar-day dummies. Interaction terms between JD and time-to-event dummies from t_{-45} to t_{-11} on the one hand and between JD and time-to-event dummies from t_{+6} to t_{+10} on the other hand are included in the regressions although not reported here. Standard errors clustered at the event level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 3: Negative buzz and Cumulative Abnormal Returns

	(1) 1-week <i>CAR</i>	(2) 3-month <i>CAR</i>
$\Delta NegBuzz$	-0.0055*** (0.0012)	-0.0202** (0.0077)
Size of Job Destruction	yes	yes
Industry dummies	yes	yes
Time dummies	yes	yes
Observations	97	97
Adjusted R-squared	0.312	0.208

Notes: Models are estimated by OLS. The dependent variable is the cumulative abnormal stock returns (*CAR*) of companies listed on the London Stock Exchange announcing a job-destruction event at t_0 . 1-week *CAR* are computed using a 7-day time window around the event $[t_{-2}; t_{+4}]$. 3-month *CAR* are computed using a 92-day time window around the event $[t_{-2}; t_{+90}]$. Normal returns used to compute *CAR* are predicted based on an estimation performed on a time window $[t_{-60}; t_{-31}]$ using the FTSE 250 Index for market returns. Net negative buzz (*NegBuzz*) is measured by the difference between the sum of negative and positive words over all tweets of the day in our database. $\Delta NegBuzz$ is computed as the difference between the average of *NegBuzz* over the $[t_{-2}; t_0]$ time window and its value at t_{-3} . Regressions include industry, year, month and weekday dummies and a control for event size (i.e. number of jobs destroyed). Standard errors clustered at the company level in parentheses. *** $p < 0.01$.

Appendix A – Tables

Table A.1: Number of events and firms

	Number of events	Number of distinct firms
All events	1,046	766
Job destructions	532	430
Job creations	514	380

Table A.2: Size of job-destruction and job-creation events
Number of jobs destroyed/created

	Mean	Std Dev.	Min	P25	Median	P75	Max
Panel A: full sample							
All events	435.83	838.32	36	130	200	400	11,000
Job destructions	468.73	888.72	36	142	236.5	410	11,000
Job creations	401.78	782.16	40	116	200	350	9,400
Panel B : equal or below median size of all events							
All events	135.73	39.65	36	100	130	163	200
Job destructions	137.58	38.56	36	109	136	168	200
Job creations	134.26	40.50	40	100	130	160	200
Panel C : above median size of all events							
All events	759.80	1,121.91	203	300	400	700	11,000
Job destructions	742.98	1,130.64	203	283	400	678	11,000
Job creations	782.88	1,112.07	205	300	440	700	9,400

Notes: Std Dev. denotes the standard deviation, P25 the 25th percentile and P75 the 75th percentile of the event size distributions.

Table A.3
Number of tweets and first-level replies, sentiment ratios and VADER score per event*day
Job-destruction events

	Obs	Mean	Std Dev.	Min	Max
Panel A : t = -45 to t = -1					
Number of tweets	23,940	192.34	1027.08	0	93,827
Number of first-level replies	23,940	52.85	383.08	0	33,223
<i>RatioNeg</i>	15,985	0.030	0.034	0	0.500
<i>RatioPos</i>	15,985	0.072	0.054	0	0.513
<i>RatioDiff</i>	14,799	0.422	0.476	-1	1
VADER score	15,985	0.145	0.212	-0.960	0.970
Panel B : t = 0 to t = +10					
Number of tweets	5,852	234.60	937.71	0	26,580
Number of first-level replies	5,852	63.35	327.75	0	12,096
<i>RatioNeg</i>	4,627	0.047	0.044	0	0.4
<i>RatioPos</i>	4,627	0.058	0.049	0	0.484
<i>RatioDiff</i>	4,351	0.130	0.591	-1	1
VADER score	4,627	0.053	0.222	-0.895	0.940

Notes: Descriptive statistics at the event*day level. Panel A covers a period ranging from 45 days ($t=-45$) to 1 day ($t=-1$) before the announcement. Panel B covers a period ranging from the announcement day ($t=0$) to 10 days after the announcement ($t=+10$). *Number of tweets* is the daily number of tweets including the company name in the text but not in the username. *Number of first-level replies* is the daily number of first-level replies to the previous tweets, dated at the date at which each tweet was posted. *RatioNeg* is equal to the daily average of the ratios of negative to total words computed for each tweet. *RatioPos* is equal to the daily average of the ratios of positive to total words computed for each tweet. *RatioDiff* is equal to the daily average of the ratios of the difference between the number of positive and negative words to the sum of positive and negative words, computed for each tweet. *VADER score* is the average score computed using the contextual VADER algorithm.

Table A.4
Number of tweets and first-level replies, sentiment ratios and VADER score per event*day
Job-creation events

	Obs	Mean	Std Dev.	Min	Max
Panel A : t = -45 to t = -1					
Number of tweets	23,130	206.63	642.32	0	26,463
Number of first-level replies	23,130	59.67	208.80	0	7,228
<i>RatioNeg</i>	15,436	0.026	0.030	0	0.333
<i>RatioPos</i>	15,436	0.082	0.059	0	0.667
<i>RatioDiff</i>	14,420	0.509	0.426	-1	1
VADER score	15,436	0.183	0.213	-0.929	0.977
Panel B : t = 0 to t = +10					
Number of tweets	5,654	210.86	588.59	0	10,376
Number of first level-replies	5,654	67.22	397.99	0	25,226
<i>RatioNeg</i>	4,302	0.023	0.027	0	0.333
<i>RatioPos</i>	4,302	0.085	0.060	0	0.597
<i>RatioDiff</i>	4,003	0.572	0.415	-1	1
VADER score	4,302	0.197	0.209	-0.883	0.961

Notes: Descriptive statistics at the event*day level. Panel A covers a period ranging from 45 days ($t=-45$) to 1 day ($t=-1$) before the announcement. Panel B covers a period ranging from the announcement day ($t=0$) to 10 days after the announcement ($t=+10$). *Number of tweets* is the daily number of tweets including the company name in the text but not in the username. *Number of first-level replies* is the daily number of first-level replies to the previous tweets, dated at the date at which each tweet was posted. *RatioNeg* is equal to the daily average of the ratios of negative to total words computed for each tweet. *RatioPos* is equal to the daily average of the ratios of positive to total words computed for each tweet. *RatioDiff* is equal to the daily average of the ratios of the difference between the number of positive and negative words to the sum of positive and negative words computed for each tweet. *VADER Score* is the average score computed using the contextual VADER algorithm.

Table A.5 – Differential impact of job-destruction vs job-creation announcements – Above-median-size events.

	(1) Number of Tweets	(2) Number of First Level Replies	(3) $Y^{(1)}$ JobDestr*RatioNeg + JobCreat*RatioPos	(4) $Y^{(2)}$ JobDestr*RatioPos + JobCreat*RatioNeg	(5) $Y^{(3)}$ – JobDestr*RatioDiff + JobCreat*RatioDiff
Job Destruction* t_0	375.943** (157.323)	138.161** (58.214)	0.019** (0.008)	-0.015** (0.006)	0.432*** (0.092)
Observations	28,168	28,168	21,768	21,768	20,573

Notes: Models are estimated using the DID_1 estimator of de Chaisemartin and D'Haultfœuille (2021; 2023). In column (1), the dependent variable is the daily number of tweets mentioning the company name in the text but not in the username. In column (2), the dependent variable is the daily number of first-level replies to the previous tweets, dated at the date at which each tweet was posted. In column (3), the dependent variable $Y^{(1)}$ is equal to the ratio of negative to total words in case of job destructions and to the ratio of positive to total words in case of job creations. In column (4), the dependent variable $Y^{(2)}$ is equal to the ratio of positive to total words in case of job destructions and to the ratio of negative to total words in case of job creations. In column (5), the dependent variable $Y^{(3)}$ is equal to the excess number of positive words (standardized by the sum of positive and negative words) in case of job creations and to the excess number of negative words (standardized by the sum of positive and negative words) in case of job destructions. *Job Destruction (JD)* is a dummy variable taking value 1 if the event is a job-destruction announcement and 0 if it is a job-creation announcement. Regressions include event dummies, time-to-event dummies (in days) from t_{-45} to t_{+10} , with t_{-3} used as a reference, as well as calendar day dummies. Interactions terms between *JD* and time-to-event dummies from t_{-45} to t_{-1} on the one hand and between *JD* and time-to-event dummies from t_{+1} to t_{+10} on the other hand are also included in the regressions although not reported here. Standard errors clustered at the event level in parentheses. *** $p < 0.01$, ** $p < 0.05$.

Table A.6 – Descriptive statistics on Cumulative Abnormal Returns

	Obs.	Mean	Std Dev.	Min	P25	Median	P75	Max
1-week <i>CAR</i>	97	0.30	5.48	-22.20	-2.70	0.23	2.73	14.20
3-month <i>CAR</i>	97	1.66	21.87	-63.91	-10.03	-0.59	15.54	66.84

Notes: *CAR* are the cumulative abnormal stock returns of companies listed on the London Stock Exchange announcing a job-destruction event at t_0 . 1-week *CAR* are computed using a 7-day time window around the event $[t_{-2}; t_{+4}]$. 3-month *CAR* are computed using a 92-day time window around the event $[t_{-2}; t_{+90}]$. Normal returns used to compute *CAR* are predicted based on an estimation performed on a time window $[t_{-60}; t_{-31}]$ using the FTSE 250 Index for market returns.


Table A.7 – Negative buzz and Cumulative Abnormal Returns – Placebo tests

	(1)	(2)
	1-week <i>CAR</i>	3-month <i>CAR</i>
$\Delta NegBuzz$	0.0001 (0.0006)	-0.0008 (0.0013)
Observations	2,716	2,716
Adjusted R-squared	0.024	0.140

Notes: Models are estimated by OLS. The dependent variable is the cumulative abnormal stock returns (*CAR*) of companies listed on the London Stock Exchange announcing a "true" job-destruction event at t_0 . For each event, we define a series of fictive events taking place at dates t_F for $t_F \in [t_{-42}; t_{-15}]$ where $[t_{-42}; t_{-15}]$ is defined with respect to the day of the "true" event t_0 . 1-week *CAR* are computed using a 7-day time window around the fictive event $[t_F-2; t_F+4]$. 3-month *CAR* are computed using a 92-day time window around the fictive event $[t_F-2; t_F+90]$. Normal returns used to compute *CAR* are predicted based on an estimation performed on a time window $[t_F-60; t_F-31]$ using the FTSE 250 Index for market returns. Net negative buzz (*NegBuzz*) is measured by the difference between the sum of negative and positive words over all tweets of the day in our database. $\Delta NegBuzz$ is computed as the difference between the average of *NegBuzz* over the $[t_F-2; t_F]$ time window and its value at $t = t_F-3$. Regressions include industry, year, month and weekday dummies and a control for event size (i.e. the number of jobs destroyed). Standard errors clustered at the company level in parentheses. *** $p < 0.01$.


Appendix B – Figures

Figure B1. Number of First-Level Replies
Correcting for 1 outlier unrelated to the event



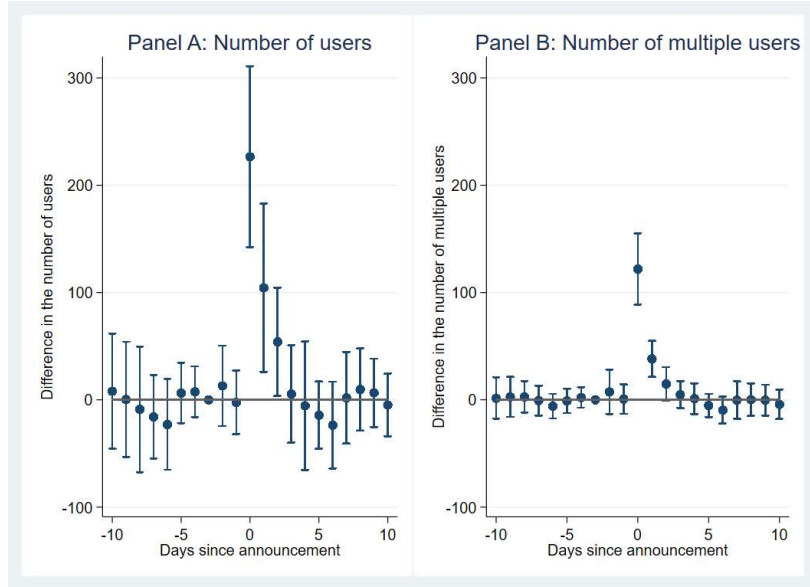
Notes: This graph reports estimated coefficients on time-to-event dummies (in days) from $t=-10$ to $t=+10$, with $t=-3$ used as a reference, obtained by estimating equation (1) using the DID_1 estimator of de Chaisemartin and D'Haultfœuille (2021; 2023). The dependent variable is the daily number of first-level replies. One tweet with 23,648 replies has been dropped at $t=+6$. Regressions also include time-to-event dummies (in days) from $t=-45$ to $t=-11$, as well as event and calendar-day fixed effects. Standard errors are clustered at the event level. Error bars correspond to 95% confidence intervals.

Figure B2. VADER Score



Notes: This graph reports estimated coefficients on time-to-event dummies (in days) from $t=-10$ to $t=+10$, with $t=-3$ used as a reference, obtained by estimating equation (1) using the DID_1 estimator of de Chaisemartin and D'Haultfœuille (2021; 2023). The dependent variable is the VADER score. Regressions also include time-to-event dummies (in days) from $t=-45$ to $t=-11$, as well as event and calendar-day fixed effects. Standard errors are clustered at the event level. Error bars correspond to 95% confidence intervals.

Figure B3. Total Number of Users and Multiple Users



Notes: This graph reports estimated coefficients on time-to-event dummies (in days) from $t=-10$ to $t=+10$, with $t=-3$ used as a reference, obtained by estimating equation (1) using the DID_1 estimator of de Chaisemartin and D'Haultfoeuille (2021; 2023). The sample is restricted to job destruction events. In Panel A, the dependent variable is the total number of users, defined as the number of distinct user identifiers who have tweeted the company name. In Panel B, the dependent variable is the number of multiple users defined as users who have tweeted mentioning the company name and for whom we have tweets mentioning at least another company for which we have an event in the same quarter in our dataset. Regressions also include time-to-event dummies (in days) from $t=-45$ to $t=-11$, as well as event and calendar-day fixed effects. Standard errors are clustered at the event level. Error bars correspond to 95% confidence intervals.