Laboratoire d'Economie de Dauphine

WP n°2/2020

Document de travail

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

Andrea Bassanini
Eve Caroli
Bruno Chaves Ferreira
Antoine Reberioux

Pôle Laboratoire d'Economie et de Gestion des Organisations de Santé (LEGOS)
Place du Maréchal de Lattre de Tassigny 75775 Paris Cedex 16
Tél (33) 01 44 05 44 46 Fax (33) 01 44 05 40 67
Site : www.legos.daupine.fr

Don't Downsize This!
Social Reactions to Mass Dismissals on Twitter[1]

Andrea Bassanini
(OECD and IZA)

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

Bruno Chaves Ferreira
(Governance Analytics, PSL)

Antoine Reberioux
(Université de Paris, LADYSS)

April 2020

Abstract

We study the reactions to job destructions on Twitter. We use information on mass dismissals and other 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 trigger numerous and strongly negative reactions. On the day of the announcement of mass dismissals, the number of tweets and first-level replies sharply increases as does the negativity of the sentiments of the posted tweets. These reactions are systematically more important than reactions to job creations. Our findings suggest that job destructions are likely to harm firms' reputation to the extent that they induce a strong negative buzz involving the company name.

Keywords: reputation, job destructions, job creations, adjustment costs, social media, sentiment analysis.

JEL codes: J63, L82, M14, 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). We are also indebted to Eric Brousseau, Kevin Geay and Julien Jourdan for useful comments and suggestions. All remaining errors are ours.

1. Introduction

Company reputation is well known to be one of the most important firm strategic assets (Kreps, 1990; Fombrun, 1996; Tadelis, 1999). To the extent that it takes time to build and is hard to imitate, reputation generates a stream of rents (Milgrom and Roberts, 1992) and hence positively affects firm performance (Roberts and Dowling, 2002; Raithel and Schwaiger, 2015).

Mass layoffs are likely to damage firm reputation since they are highly visible and often perceived as unfair (Charness and Levine, 2000; Hallock, 2009). In addition, the way workers are treated is part of the credence attributes that consumers value (Baron, 2011). Consistent with these observations, the literature in management shows that mass layoffs negatively affect firms' reputation as assessed by senior executives and outside directors in the America's Most Admired Corporations (AMAC) survey – see Flanagan and O'Shaughnessy (2005); Love and Kraatz (2009); Schulz and Johann (2018). In contrast, evidence of social reactions to dismissals from outside the business community 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, with the spreading of social media, these negative social reactions may rapidly become viral, since online social networks have become 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 a strong negative buzz involving the company names on Twitter. We focus on Twitter since it is one of the most important social-media platforms. In addition, on Twitter, reactions to information are almost instantaneous, which permits clear identification of the impact of mass-dismissal announcements on social buzz.

We rely on the Restructuring Events database, which is part of the European Restructuring Monitor (ERM) and provides information on job-creation and job-destruction announcements for a large number of EU companies since the early 2000s. We consider announcements made by companies in the United Kingdom 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 run an event analysis showing 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. However, when running a difference-in-difference estimation, we find that the increase in the number of tweets is significantly larger following job-destruction announcements than following job-creation 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). 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, as well as a more refined score constructed by using the contextual VADER algorithm and capturing the overall positivity or negativity of each tweet.[2] Our event analysis shows that the average negativity of the tweets significantly increases following job-destruction announcements. We also provide evidence that the number of negative words in the tweets following job destructions increases by a much larger amount than the number of positive words in the tweets following job creations. Similarly, we show that the fall in positive sentiments following job destructions is of greater magnitude than the decrease in negative sentiments following job creations. Finally, we show that the reduction in the VADER score following job destructions is larger in absolute value than the increase in this score following job creations. We interpret these results as indicating that job destruction generates a negative buzz involving the company's name, which is likely to damage the image of the firm in the society and to impact its reputation with stakeholders.

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; Van Dick et al., 2016; Sucher and Gupta, 2018; Heinz et al., 2020), thereby raising unit labour costs. Eventually, a more recent literature shows that firms refrain from cutting jobs close to headquarters and hypothesises that this strategy could aim at avoiding the reputational cost that dismissals of its members

[2] This score has been used in the literature in economics by Shapiro and Wilson (2019) and Shapiro et al. (2019), among others.

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 dismissal announcements trigger negative reactions on social media, thereby generating reputational costs 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 growing role on product markets. On the one hand, they provide a new source of information for consumers, and should therefore be integrated in firms' marketing strategies (Chen and Xie, 2008). On the other hand, social media (Facebook and Twitter in particular) are increasingly used by consumers to get companies to do what they think would be fair (Hendel et al., 2017). Social media also play an increasing role in financial markets. Siganos et al. (2014) and Deng et al. (2018) show that they reflect the sentiments of the investors' community; in addition, Nguyen et al. (2019) and Chen et al. (2014) provide evidence that professional institutional investors use data scraping and data mining to capture social-media sentiments, and trade in accordance with these sentiments. More generally, a burgeoning literature focuses on the role of social media in the formation of firm reputation (see Etter et al., 2019, for a review of the literature). We speak to this literature by using an original dataset and showing that human resource management decisions entail a significant buzz involving companies' name on social media, which contributes to shape firm reputation and may therefore affect firm performance.

Eventually, our paper also speaks to a more limited literature focusing on the differential way in which media react to dismissals' announcements as compared to job creations'. Heinz and Swinnen (2015) document a media slant against dismissals in Germany. Based on the review of daily articles in a leading German newspaper over 8 years, they find 20 times as many articles reporting on job destructions as articles reporting on job creations. Friebel and Heinz (2014) show that media slant against dismissals is particularly strong in the case of foreign firms. We provide evidence that dismissals also trigger many more reactions than job creations on one of the most important social media, i.e. Twitter.

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 empirical results. Section 5 concludes.

2. Data

The first dataset that we use is the ERM Restructuring Events database.[3] 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.[4] We consider restructuring events reported in the United Kingdom over 2013-2018. 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,[5] 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. 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 1,047 useful events, corresponding to 813 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 events is slightly larger than that of job creations with, on average 469

[3] Available at http://www.eurofound.europa.eu/observatories/emcc/erm/factsheets.

[4] According to OECD (2018) mass dismissals reported in the ERM Restructuring Events database account for 15% of all dismissals (including dismissals for personal reasons and individual redundancies) in the UK. As mass dismissals are only a small proportion of all dismissals, this dataset likely covers most large-scale dismissal events.

[5] 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.

jobs destroyed in job-destruction events as compared to 402 jobs created in job-creation events. Similarly, the median size of job destructions is slightly larger than that of job creations (237 vs 200).

For each tweet, we have information on the username, the exact date of the post, the number of first-level replies[6] 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, while we are interested in social reactions to job destructions rather than in information disclosed by companies. Overall, our database contains 11,697,816 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), we first consider a number of word-count variables: 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). Alternatively, we use a score constructed using the contextual VADER algorithm of sentiment analysis.[9] 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" or "so 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).

[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 post 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.

[9] https://github.com/cjhutto/vaderSentiment.

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*, *RatioDiff* and of the VADER score 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, with a daily average of 233 tweets between $t = 0$ and $t = +10$, as compared to 190 between $t = -45$ and $t = -1$. A much smaller difference is observed for job creations with a daily average number of tweets of 205 after the announcement as compared to 200 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* and the VADER score which capture the overall positivity of the sentiments expressed also go down from 0.417 to 0.141 for the former and from 0.146 to 0.053 for the latter – see Appendix Table A.3. As regards job creations, the ratios of negative (resp. positive) to total words remain almost stable after the announcement has taken place – see Appendix Table A.4. In contrast, *RatioDiff* and the VADER score both increase – although by a much smaller amount than their decrease following job destructions – from 0.500 to 0.562 for the former and from 0.183 to 0.198 for the latter.

3. Empirical Model

Our main goal 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.

3.1 Event study

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

$$Y_{jt} = \sum_{t=-45}^{t=-4} \alpha_t D_t + \sum_{t=-2}^{t=10} \beta_t D_t + \mu_f + \mu_y + \mu_m + \mu_{wd} + \varepsilon_{jt} \quad (1)$$

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

D_t are dummy variables measuring the time distance in days from the date of the event, i.e. of the announcement ($t = 0$). 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 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) 2 days before. Similarly, the downsizing of 150 persons in Belfast by Legal and General in September 2017 was officially confirmed by a company spokesperson 2 days after the staff received an email and the information leaked to the press. In such cases, taking $t = -1$ as a reference would wrongly underestimate the true effect of the announcement.

We include year (μ_y), month (μ_m) and weekday (μ_{wd}) fixed effects to account for the fact that Twitter activity is unevenly distributed over time.[10] μ_f is a firm fixed effect capturing the fact that some companies attract more tweets than others in normal times, i.e. immediately before the event takes place. Standard errors are clustered at the company level.

Our key parameter of interest is β_0, which yields the magnitude of the change in the outcome variable from $t = -3$ to $t = 0$. 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

[10] 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. There are also fewer tweets in summer and, as discussed above, the number of twitter users has increased over time, although only slowly since 2013.

expect it to be significantly negative for *RatioPos*, *RatioDiff* and the VADER score 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 and significant (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* and the VADER score. By contrast, we expect it to be significantly 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. Moreover, if reactions on Twitter continue over several days, $\hat{\beta}_t$ (for $t > 0$) will carry the same sign and significance as $\hat{\beta}_0$. This would be a noteworthy result since previous studies have shown that reactions on Twitter are immediate, and their intensity fades away quickly over time even if related to persistent phenomena or changes in public opinions (see e.g. O'Connor et al., 2011; Stautz et al., 2017; Stolee and Caton, 2018; Yousefinaghani et al., 2019).

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

As a second step, we want to gauge the 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:

$$Y_{jt} = \sum_{t=-45}^{t=-4} \alpha'_t D_t * JbD_j + \sum_{t=-2}^{t=10} \beta'_t D_t * JbD_j + \sum_{t=-45}^{t=-4} \alpha_t D_t + \sum_{t=-2}^{t=10} \beta_t D_t + \delta JbD_j + \mu_f + \mu_y + \mu_m + \mu_{wd} + \varepsilon_{jt} \quad (2)$$

where JbD_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. δ captures the difference in the value of the outcome variable between job destructions and job creations at $t = -3$.[11] Our main parameter of interest is ${\beta'}_0$ 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. Equation (2) is similar to a difference-in-difference (DID) model in which we compare the effects of two different treatments (job-destruction and job-creation announcements) rather than a treated and a control group. To make sure that our comparison is meaningful, as in a DID model, we need to check that pre-event trends and levels are not significantly different between job destructions and job creations, i.e. $\widehat{\alpha'}_t$ for $t \in [-45; -4]$ and $\hat{\delta}$ 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: $\widehat{\beta'}_0$ would be positive for the ratio of negative to total words and it would be negative for the ratio of positive to total words, *RatioDiff* and the VADER score.

To obtain a meaningful comparison, we consider new, modified dependent variables. The first one is defined as follows:

$$Y_{jt}^{(1)} = JbD_j * RatioNeg_{jt} + JbC_j * RatioPos_{jt}$$

where JbC_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. $Y_{jt}^{(1)}$ 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 $Y_{jt}^{(1)}$, $\widehat{\beta'}_0$ (and $\widehat{\beta'}_{t>0}$) 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. Note that, while it is still essential to check that pre-event trends are parallel, the condition on levels (i.e. $\hat{\delta}$=0) will not be satisfied since, in our data, there are more positive than negative words in the pre-event period.

[11] Because our model includes a firm fixed effect, δ is identified only on firms with at least one job-destruction and one job-creation event over our sample period.

Similarly, we define $Y_{jt}^{(2)}$ 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, $\widehat{\beta'}_0$ (and $\widehat{\beta'}_{t>0}$) will be negative and significant if job destructions reduce positive sentiments more than job creations reduce negative sentiments.

The third variable we consider is:

$$Y_{jt}^{(3)} = -JbD_j * RatioDiff_{jt} + JbC_j * RatioDiff_{jt}$$

where, as defined in Section 2:

$$RatioDiff_{jt} = \frac{Positive\ Words_{jt} - Negative\ Words_{jt}}{Positive\ Words_{jt} + Negative\ Words_{jt}}$$

$Y_{jt}^{(3)}$ is therefore equal to the excess number of positive words (standardised by the sum of positive and negative words) in the case of job creations and to the excess number of negative words (standardised by the sum of positive and negative words) in the case of job destructions.

As regards the VADER score, we define the following variable:

$$Y_{jt}^{(4)} = -JbD_j * VADER_{jt} + JbC_j * VADER_{jt}$$

When estimating equation (2) on either $Y_{jt}^{(3)}$ or $Y_{jt}^{(4)}$, $\widehat{\beta'}_0$ will be negative and significant if negative reactions in case of job destructions are of greater magnitude than positive reactions in case of job creations.

4. Results

4.1 Number of Tweets and First-Level Replies

We first estimate equation (1) for each of our outcome variables separately. The regression coefficients and standard errors are plotted against the time distance to the announcement – see Figures 1 to 6. $t = -3$ 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, whatever the outcome we consider.

As evidenced on Figure 1 – 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: +293, significant at the 1% level, as compared to an average of 176 at $t = -3$. It is still significantly higher at $t = +1$ and $t = +2$, although by a smaller amount (+124 and +68, 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 estimated number of tweets also increases, but by a much smaller amount: +58 at $t = 0$ and +22 at $t = +1$, as compared to an average of 185 at $t = -3$.

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 a DID model – see equation (2) – with the number of tweets as the dependent variable. The results are presented in Table 1. As can be seen on the first line of column (1), the number of tweets associated to job-destruction events is not significantly different from that associated to job-creation events at the reference date ($t = -3$). Moreover, 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 247 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 $t = +1$ than at the reference date (+105), while it becomes insignificant at later dates. As a consequence, we only report the estimated coefficients until $t = +5$.[12] Overall, these results indicate that job-destruction announcements trigger many more reactions on Twitter than job-creations'.

We then consider 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. Since they are tweets themselves, some of them are included in the number of tweets we considered above. However, this is the case only if the reply mentions the name of the company that made the announcement. Now, many replies do not include the company name and are therefore not captured by our scraping algorithm. So, we consider the overall number of replies separately

[12] 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 in Table 1.

as a complementary indicator of the social reactions triggered by job-destruction (resp. job-creation) announcements. This number 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: +82 as compared to an average of 49 replies at $t = -3$. The number of replies is also higher in the following days, 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. One caveat is that, 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*][13] announced some job creations. 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 A.1.

Coming to the DID estimates 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. The estimated number of replies associated to job-destruction events is not significantly different from that associated to job-creation events at the reference date ($t = -3$), and 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 $t = -3$ (+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.

As emphasised above, the dynamics of first-level replies is mechanically short-termed since they carry the date of the original tweet to which they are attached. However, the dynamics of the number of tweets is also quite short-lived since the sharp increase observed on the day of a

[13] The company name has been suppressed for confidentiality reasons.

job-destruction announcement fades away within 3 days. One could be concerned that such short-term buzz might not carry sizable consequences for firms. The literature in finance, however, suggests that this is not the case. Mao et al. (2012) indeed show that the daily number of tweets that mention S&P 500 stocks helps predicting whether the closing price will go up or down, by making the estimation more precise. Sprenger et al. (2014a) also show that the daily number of tweets mentioning the dollar-tagged ticker symbol of an S&P 500 company is reflected in stock prices. Moreover, many papers in this literature find that taking Twitter sentiment into account substantially improves the accuracy of stock price predictions – see Ranco et al. (2015), Leitch and Sherif (2017), Sprenger et al. (2014b). This is an additional reason to analyse the impact of job-destruction and job-creation announcements on the sentiments expressed on Twitter, which we present in the next subsection.

4.2 Sentiment Analysis

4.2.1 Word-count variables: RatioNeg, RatioPos and RatioDiff

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* more than doubles: +3.4 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 6 days after the event: at $t = +6$, it is still 0.7 percentage points higher than at $t = -3$ (significant at the 5% level). Visual inspection of Figure 3 – Panel A suggests that *RatioNeg* is already marginally higher at $t = -1$ than at $t = -3$, thus suggesting that some anticipation could take place. This is not impossible due to potential leakages in the press before the planned job destructions get officially announced by the company – see Section 3. However, testing for the difference in the regression coefficients across $t = -1$ and $t = 0$ yields unambiguous results: the increase in negativity at $t = 0$ is much larger than at $t = -1$ with a difference significant at the 1% level. Moreover, for all dates before $t = -1$, the estimated ratio of negative to total words is never significantly different from that at $t = -3$. 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 2.8 percentage points at $t = 0$ (significant at the 1% level), as compared to a baseline level of 7% at $t = -3$. As for *RatioNeg*, the deviation from the reference value lasts for several days: by

$t = +5$, *RatioPos* is still 0.8 percentage points lower than at baseline, significant et the 1% level. Here again, the ratio of positive to total words is already marginally lower at $t = -1$ than at $t = -3$, but we check that this reduction is significantly smaller than the one taking place at the time when job destructions are announced ($t = 0$). The difference in coefficients is indeed significant at the 1% level. This pattern is confirmed by our third indicator, *RatioDiff* – see Sections 2 and 3. As evidenced on Figure 5 – Panel A, the estimated excess number of positive over negative words (standardised by the sum of positive and negative words) sharply decreases following a job-destruction announcement: -62.9 percentage points (significant at the 1% level) as compared to a baseline level of 41.5% at $t = -3$. This reduction is quite long lasting since it only fades away 10 days after the event. As evidenced in the graph, some anticipation takes place at $t = -1$, with a reduction of *RatioDiff* by 11.6 percentage points. However, here again, this decrease is of much smaller magnitude than the one taking place at $t = 0$, with the difference in coefficients being significant at the 1% level. 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 event studies also suggest 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 (-0.47 percentage points at $t = 0$ as compared to 2.4% at $t = -3$, significant at the 1% level) and that this reduction fades away within 2 days. Consistent with this decrease in negativity, the ratio of positive to total words significantly increases at $t = 0$ as compared to $t = -3$ (+1.1 percentage points at $t = 0$ as compared to 8.1% at $t = -3$, significant at the 1% level), as does the excess number of positive words, *RatioDiff* (+9.2 and +6 percentage points at $t = 0$ and $t = +1$ respectively, as compared to a baseline value of 53.2% 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 3 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. The coefficient on the *Job Destruction* variable (-0.039, significant at the 1% level) suggests that, at $t = -3$, there are more positive words in the tweets mentioning the names of the companies that are about to announce job creations than negative words in the tweets mentioning the names of the companies that are about to announce job destructions. This is not surprising since, no matter the type of event, tweets systematically contain more positive than negative words in pre-event periods – see Appendix Tables A.3 and A.4. 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 – and this pattern lasts until $t = +3$.[14] 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. This differential impact lasts longer than for $Y^{(1)}$, since it only fades away after $t = +5$.[15] 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. Note that for all outcomes in Table 3, there is some anticipation effect either at $t = -1$ – for $Y^{(1)}$ and $Y^{(3)}$ – or $t = -2$ – for $Y^{(2)}$. However, the tests of the difference in coefficients presented in Table 2 confirm that, for all outcomes, the change between $t = 0$ and $t = -3$ is much larger than the change between $t = -1$ (or $t = -2$) and $t = -3$ – with the difference being significant at conventional levels.[16] This implies that taking $t = -1$ (or $t = -2$) as a reference, the negativity of the sentiments still significantly increases on the day of the announcement. Overall, our findings support the idea that there exists an asymmetry between job-destruction

[14] Interactions terms between *Job Destruction* and time-to-event dummies from t = -45 to t = -11 on the one hand and between *Job Destruction* and time-to-event dummies from t = +6 to t = +10 on the other hand are included in the regression although not reported since none of them are significant at conventional levels.

[15] In this case, the point estimate on *Job Destruction* is positive (0.043), meaning that at t=-3, there are more positive words in the tweets mentioning the names of the companies that are about to announce job destructions than negative words in the tweets mentioning the names of the companies that are about to announce job creations. Given this gap, one could worry that the larger reduction (in absolute value) in the share of positive words that we find following a job-destruction event does not correspond to a larger percentage change in this share (since it is initially larger). However, estimating equation (2) by replacing the dependent variable $Y^{(2)}$ with log(1+$Y^{(2)}$) yields similar results (with, at t=0, a point estimate equal to -0.021 and a standard error of 0.003, and significant coefficients up to t=+5): this indicates that the decrease in the share of positive words following job destructions is proportionally larger than the decrease in the share of negative words following job creations.

[16] This also holds for $t \in [1,3]$ for $Y^{(1)}$, $t \in [1,2]$ for $Y^{(2)}$, and $t \in [1,5]$ for $Y^{(3)}$. Detailed results are available upon request.

and job-creation announcements: the negative buzz generated by the former is significantly stronger than the positive buzz generated by the latter.

4.2.2 VADER score

Measuring sentiments based on a simple word count is, of course, very crude. As an alternative indicator of sentiments expressed in the tweets, we use the score computed using the contextual VADER algorithm which takes into account punctuation, negation, capital letters and the use of intensifiers – see Shapiro and Wilson (2019) and Shapiro et al. (2019). The event study for this outcome is presented on Figure 6. As evidenced in Panel A, the positive sentiments expressed by the tweets collapse following a job-destruction announcement: the estimated VADER score decreases by 0.198 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 $t = +10$. As can be seen in the chart, some anticipation takes place at $t = -1$ and $t = -2$, but as for the other outcomes, tests for the difference across the regression coefficients show that the reduction taking place at $t = 0$ is significantly larger than at prior dates. In contrast, the change in the VADER score following job creations is much smaller (+0.026 as compared to 0.190 at $t = -3$, significant at the 5% level) and it only takes place on the very day of the announcement – see Panel B of Figure 6.

To make sure that this differential effect of job creations and job destructions is statistically significant, we estimate equation (2) for output $Y^{(4)}$. This is equal to the VADER score in case of job creations and to the opposite of the VADER score in case of job destructions. The results are presented in Table 3. The decrease in the VADER score following a job-destruction announcement is much larger than the increase following a job-creation announcement and this gap remains positive and significant until 8 days after the event takes place.[17] Some anticipation takes place at $t = -1$ and $t = -2$, but as shown in the Table, the difference in regression coefficients is significantly larger at $t = 0$ than at earlier dates. These findings confirm that job-destruction announcements trigger a negative buzz which is much more important than the positive buzz generated by job creations. Overall, this suggests that firm reputation is negatively affected by job destructions and that this effect is quite long lasting by Twitter standards since it lasts for at about one week.

[17] Complete results available upon request.

4.3 Robustness checks

Our empirical strategy essentially relies on comparing social reactions to job-destruction and job-creation announcements. However, in most cases, both types of events are not initiated by the same companies. One could be concerned that social-network reactivity could be systematically greater regarding specific companies. If these firms have a greater propensity to announce job destructions, our results could be driven by this reactivity bias. To overcome this problem, we re-estimate equation (2) for our 6 outcome variables including firm-by-time-to-event dummies. These account for the fact that the time pattern of social reactions could be firm specific. The results are presented in Appendix Table A.5. They are similar to those reported in Tables 1, 2 and 3: job-destruction announcements trigger significantly more reactions on Twitter than job-creation announcements – see columns (1) and (2) – and the sentiments expressed in the corresponding tweets are not only more negative, but also intensified with respect to job creations – see columns (3) to (6). This suggests that our main results are not due to firm heterogeneity in social-media reactivity.

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 (2) separately for 2 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 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 Table A.6 – column (1), the estimated total number of users increases by a larger amount following job-destruction announcements as compared to job-creation ones: +178, significant at the 1% level. Interestingly, out of those 178 additional users, 96 are multiple users – see column (2) – thus suggesting that more than half of those individuals who account for the difference between job creations and job destructions have reacted to several events and are hence unlikely to do so only because they have been personally impacted.

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. To make sure that this is not driving our results, we re-estimate equation (2) for our 6 main outcome variables on 2 subsamples: the first one contains job-creation and job-destruction events the size of which is above the median size of all events while the second subsample contains the events whose size is below the median – see Appendix Table A.7. We do so because Panel C of Appendix Table A.2 shows that, within the group of events with size above median, the average size of job creations is slightly larger than that of job destructions (781 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 on Panel A of Appendix Table A.7, large-size job destructions trigger more reactions than large-size job creations: they attract altogether more tweets, more replies and more acute sentiments – see columns (1) to (6). 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.[18]

5. Conclusion

We have shown 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 of mass dismissals, 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, the excess number of positive over negative words and the VADER score significantly decrease. The negative effect of job-destruction announcements on the sentiments expressed by the tweets is surprisingly long-lived with respect to Twitter standards (almost one week for most outcomes), whereas the impact in terms of buzz (number of tweets and replies) is much shorter (3 days for tweets and only 1 day for replies). The reactions to job-destruction events are systematically larger than the reactions to job-creation events. The latter trigger fewer tweets and replies and, interestingly, they also trigger weaker changes in sentiments: the

[18] Comparing Panels A and B of Appendix Table A.7 also suggests that small-size job destructions trigger fewer reactions (as compared to small-size job creations) than larger ones: the impact on the number of tweets is indeed substantially lower in Panel B than in Panel A, while the effect on the number of replies in Panel B is insignificant at conventional levels for small-size events. As regards the negativity of reactions to job destructions, it is slightly stronger for large-size events than for smaller ones, but the gap is much smaller than for the number of tweets and replies – see columns (3) to (6) of Panels A and B. This suggests that when job destructions involve a relatively small number of employees, fewer tweets are posted, but these still contain negative reactions that are more acute than the positive reactions triggered by small-scale job creations.

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.

Our findings therefore suggest that job destructions are likely to harm firms' reputation to the extent that they induce a negative buzz involving the company name. This extends the existing results regarding reputation with peers and with the general public, by emphasising the role of human-resource management decisions in the making of company reputation. It also raises 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 loss? If job destructions are announced and social reactions are particularly fierce, do they reduce the scope of their original downsizing plan?

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

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.

Baron, David. 2011. Credence attributes, voluntary organizations, and social pressure, *Journal of Public Economics*, 95(11-12): 1331-1338.

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 Doulkeridis. 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.

Charness, Gary and David Levine. 2000. Are Layoffs Acceptable? Evidence from a Quasi-Experient. *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.

Chen, Hailiang, Prabuddha De, Yu Hu and Byoung-Hyoun Hwang. 2014. Wisdom of Crowds: The Value of Stock Opinions Transmitted Through Social Media, *The Review of Financial Studies*, 27(5): 1367-1403.

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.

Deng, Shuyuan, Zhijian Huang, Atish Sinha and Huimin Zhao. 2018. The Interaction between Microblog Sentiment and Stock Return: An Empirical Examination, *MIS quarterly*, 42(3): 895-918.

Drzensky, Frank and Matthias Heinz. 2016. The Hidden Costs of Downsizing, *The Economic Journal*, 126(598): 2324-2341.

Etter, Michael, Davide Ravasi and Elanor Colleoni. 2019. Social Media and the Formation of Organizational Reputation, *Academy of Management Review*, 44(1): 28-52.

Flanagan, David and Kenneth O'Shaughnessy. 2005. The Effect of Layoffs on Firm Reputation, *Journal of Management*, 31(3): 445-463.

Fombrun, Charles. 1996. *Reputation. Realizing Value from the Corporate Image*, Boston: Harvard Business School Press, 441p.

Friebel, Guido and Matthias Heinz. 2014. Media Slant Against Foreign Owners: Downsizing, *Journal of Public Economics*, 120, 97-106.

Gentzkow, Matthew, Bryan Kelly and Matt Taddy. 2019. Text as Data, *The Journal of Economic Literature*, 57(3): 535-574.

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.

Heinz, Matthias and Johan Swinnen. 2015. Media slant in economic news: A factor 20, *Economics Letters*, 132(C): 18-20.

Heinz, Matthias, Sabrina Jeworrek, 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.

Kramarz, Francis and Marie-Laure Michaud. 2010. The Shape of Hiring and Separation Costs in France, *Labour Economics*, 17(1): 27-37.

Kreps, David M. 1990. Corporate Culture and Economic Theory, in James Alt and Kenneth A. Shepsle, eds. *Perspectives on positive political economy*. New York, Cambridge University Press: 90-143.

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.

Leitch, Darren and Mohamed Sherif. 2017. Twitter mood, CEO succession announcements and stock returns, *Journal of Computational Sciences*, 21: 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.

Mao, Yuexin, Wei Wei, Bing Wang and Benyuan Liu. 2012. Correlating S&P 500 Stocks with Twitter Data, 10.1145/2392622.2392634.

Milgrom, Paul and John Robert. 1992. *Economics, Organization and Management*, Englewood Cliffs, New Jersey, Prentice Hall, 621p.

Nguyen, Hang, Roger Calantone and Ranjani Krishnan. 2019. Influence of Social Media Emotional Word of Mouth on Institutional Investors' Decisions and Firm Value, *Management Science*, 66(2): 887-910.

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.

O'Connor, Brendan, Ramnath Balasubramanyan, Bryan Routledge and Noah Smith. 2011. From Tweets to polls: Linking text sentiment to public opinion time series, in *Proceedings of the Fourth International AAAI Conference on Weblogs and Social Media*, Menlo Park, California, The AAAI Press: 122-129.

Raithel, Sascha and Manfred Schwaiger. 2015. The Effect of Corporate Reputation Perceptions of the General Public on Shareholder Value, *Strategic Management Journal*, 36(6): 945-956.

Ranco, Gabriele, Darko Aleksovski, Guido Caldarelli, Miha Grcar and Igor Mozetic. 2015. The Effect of Twitter Sentiments on Stock Price Returns, *PLoS One*, 10(9), doi: 10.1371/journal.pone.0138441.

Roberts, Peter and Grahame Dowling. 2002. Corporate Reputation and Sustained Superior Financial Performance, *Strategic Management Journal*, 23(12): 1077-1093.

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. 2019. Taking the Fed at its Word: A New Approach to Estimating Central Bank Objectives using Text Analysis, Federal Reserve Bank of San Francisco Working Paper 2019-02, R&R *Review of Economic Studies*.

Shapiro, Adam, Moritz Sudhof and Daniel Wilson. 2019. Measuring News Sentiment, Federal Reserve Bank of San Francisco Working Paper 2017-01. R&R *Journal of Econometrics*.

Siganos, Antonios, Evangelos Vagenas-Nanos and Patrick Verwijmeren. 2014. Facebook's daily sentiment and international stock markets, *Journal of Economic Behavior & Organization*, 107(Part B): 730-743.

Sprenger, Timm, Philipp Sandner, Andranik Tumasjan and Isabell Welpe. 2014a. News or Noise? Using Twitter to Identify and Understand Company-specific News Flow, *Journal of Business Finance and Accounting*, 41(7&8): 791-830.

Sprenger, Timm, Andranik Tumasjan, Philipp Sandner and Isabell Welpe. 2014b. Tweets and Trades: the Information Content of Stock Microblogs, *European Financial Management*, 20(5): 926–957.

Stautz, Kaidi, Giacomo Bignardi, Gareth Hollands and Theresa Marteau. 2017. Reactions on Twitter to updated alcohol guidelines in the UK: a content analysis, *British Medical Journal Open*, 7, doi:10.1136/bmjopen-2016-015493.

Stolee, Galen and Steve Caton. 2018. Twitter, Trump, and the Base: A Shift to a New Form of Presidential Talk?, *Signs and Society*, 6(1): 147-165.

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.

Tadelis, Steven. 1999. What's in a Name? Reputation as a Tradable Asset, *The American Economic Review*, 89(3): 548-563.

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.

Van Dick, Rolf, Frank Drzensky and Matthias Heinz. 2016. Goodbye or Identify: Detrimental Effects of Downsizing on Identification and Survivor Performance, *Frontiers in Psychology*, 7(771): 1-9.

Yousefinaghani, Samira, Rozita Dara, Zvonimir Poljak, Theresa Bernardo and Shayan Sharif, 2019. The Assessment of Twitter's Potential for Outbreak Detection: Avian Influenza Case Study, *Nature: Scientific Reports*, 9(18147).

Figures

Figure 1. Event Study – Number of Tweets

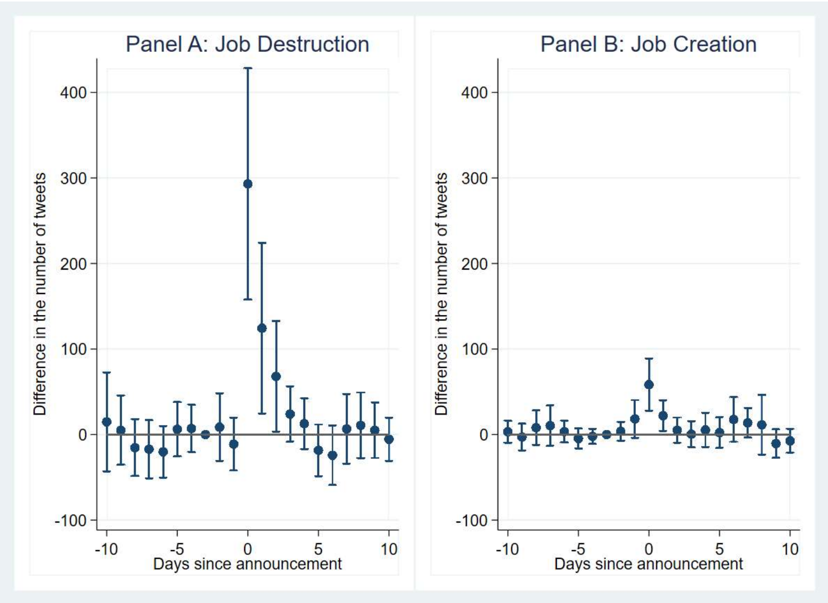

Notes: This graph reports estimated coefficients on time-to-event dummies (in days) from t=-10 to t=+10, with t=-3 used as a reference, obtained by estimating equation (1) by OLS. The dependent variable is the daily number of tweets mentioning the company name in the text but not in the username. Regressions also include firm dummies, time-to-event dummies (in days) from t=-45 to t=-11, as well as year, month and weekday dummies. Standard errors are clustered at the company level. Error bars correspond to 95% confidence intervals.

Figure 2. Event Study – Number of First-Level Replies

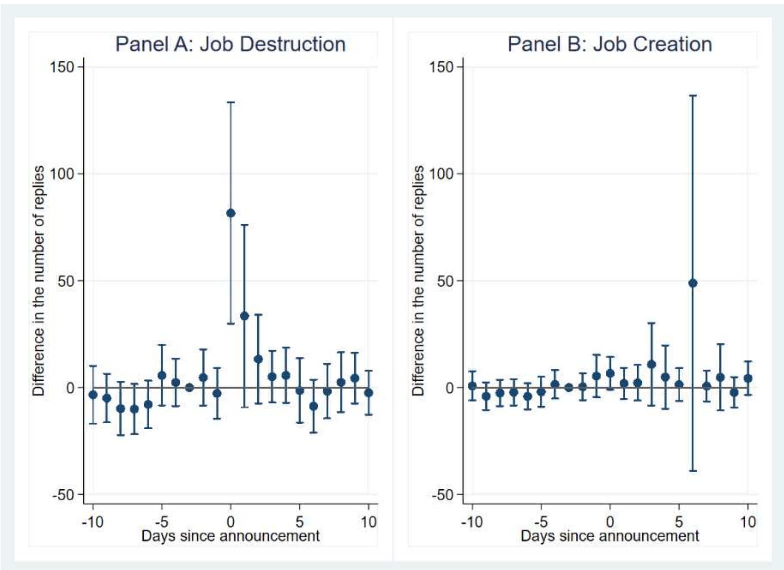

Notes: This graph reports estimated coefficients on time-to-event dummies (in days) from t=-10 to t=+10, with t=-3 used as a reference, obtained by estimating equation (1) by OLS. The dependent variable is the daily number of first-level replies. Regressions also include firm dummies, time-to-event dummies (in days) from t=-45 to t=-11, as well as year, month and weekday dummies. Standard errors are clustered at the company level. Error bars correspond to 95% confidence intervals.

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

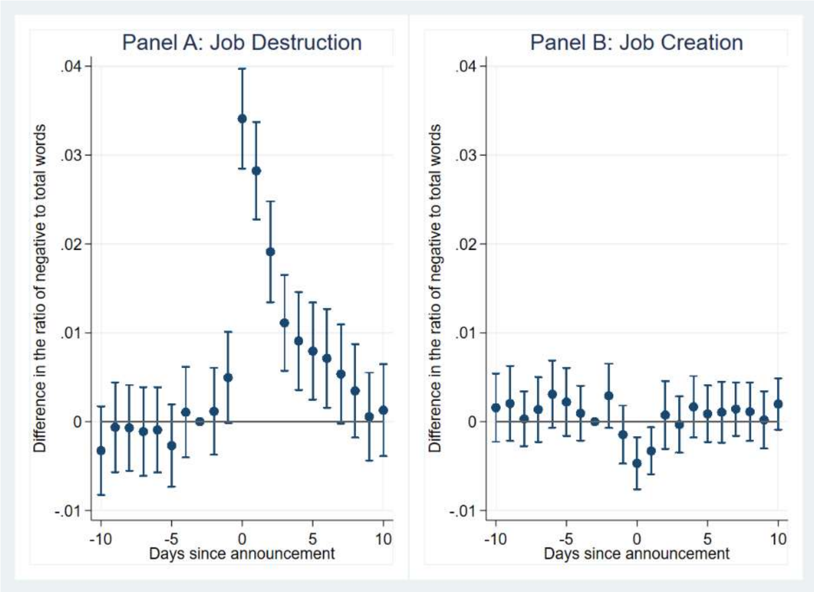

Notes: This graph reports estimated coefficients on time-to-event dummies (in days) from *t*=-10 to *t*=+10, with *t*=-3 used as a reference, obtained by estimating equation (1) by OLS. The dependent variable is the ratio of negative to total words. Regressions also include firm dummies, time-to-event dummies (in days) from *t*=-45 to *t*=-11, as well as year, month and weekday dummies. Standard errors are clustered at the company level. Error bars correspond to 95% confidence intervals.

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

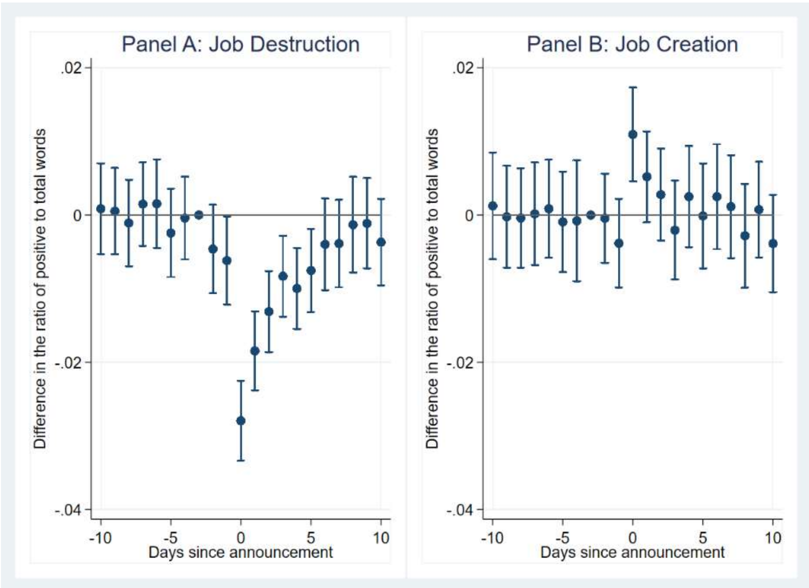

Notes: This graph reports estimated coefficients on time-to-event dummies (in days) from *t*=-10 to *t*=+10, with *t*=-3 used as a reference, obtained by estimating equation (1) by OLS. The dependent variable is the ratio of positive to total words. Regressions also include firm dummies, time-to-event dummies (in days) from *t*=-45 to *t*=-11, as well as year, month and weekday dummies. Standard errors are clustered at the company level. Error bars correspond to 95% confidence intervals.

Figure 5. Event Study – *RatioDiff*

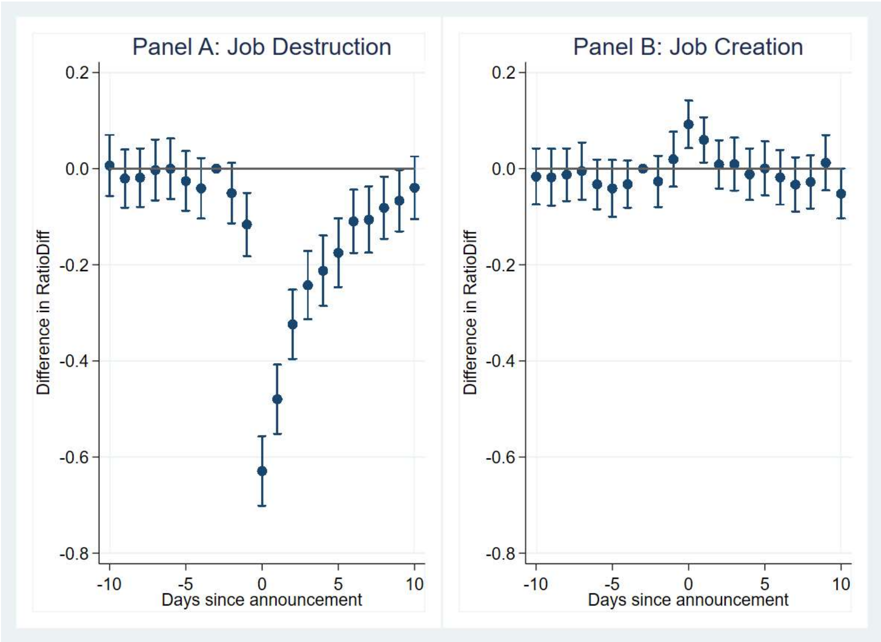

Notes: This graph reports estimated coefficients on time-to-event dummies (in days) from *t*=-10 to *t*=+10, with *t*=-3 used as a reference, obtained by estimating equation (1) by OLS. The dependent variable is the ratio of the difference between the number of positive and negative words to the sum of positive and negative words. Regressions also include firm dummies, time-to-event dummies (in days) from *t*=-45 to *t*=-11, as well as year, month and weekday dummies. Standard errors are clustered at the company level. Error bars correspond to 95% confidence intervals.

Figure 6. Event Study – VADER Score

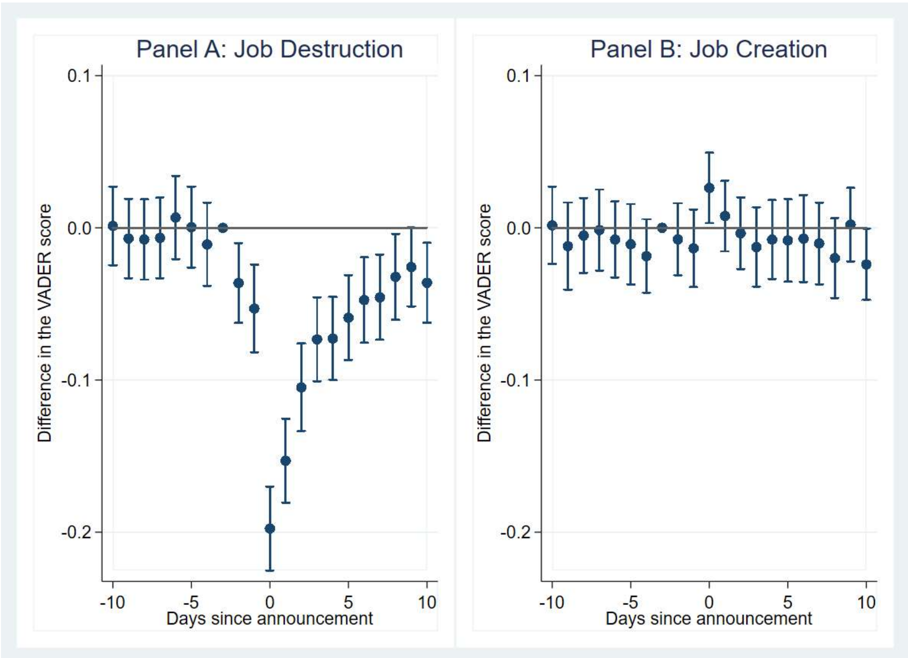

Notes: This graph reports estimated coefficients on time-to-event dummies (in days) from *t*=-10 to *t*=+10, with *t*=-3 used as a reference, obtained by estimating equation (1) by OLS. The dependent variable is the VADER score. Regressions also include firm dummies, time-to-event dummies (in days) from *t*=-45 to *t*=-11, as well as year, month and weekday dummies. Standard errors are clustered at the company 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	9.94	11.56
	(70.37)	(25.27)
Job Destruction*t_0	246.80***	77.57***
	(71.72)	(26.63)
Job Destruction*t_{+1}	105.21**	30.78
	(53.18)	(22.68)
Job Destruction*t_{+2}	57.18	9.01
	(34.86)	(11.88)
Job Destruction*t_{+3}	17.96	-7.15
	(17.82)	(11.56)
Job Destruction*t_{+4}	7.23	0.77
	(17.21)	(9.99)
Job Destruction*t_{+5}	-13.48	-1.11
	(15.22)	(8.37)
Job Destruction*t_{-1}	-17.11	-4.90
	(16.17)	(7.40)
Job Destruction*t_{-2}	12.04	6.05
	(19.87)	(7.27)
Job Destruction*t_{-3}	*ref*	*ref*
	-	-
Job Destruction*t_{-4}	3.67	-0.42
	(14.02)	(6.60)
Job Destruction*t_{-5}	5.10	5.56
	(16.60)	(8.25)
Job Destruction*t_{-6}	-20.68	-4.55
	(15.79)	(6.06)
Job Destruction*t_{-7}	-15.26	-5.30
	(18.90)	(6.22)
Job Destruction*t_{-8}	-11.01	-4.12
	(17.88)	(6.58)
Job Destruction*t_{-9}	15.13	0.78
	(22.85)	(6.85)
Job Destruction*t_{-10}	11.60	-4.11
	(29.65)	(7.77)
Job Destr.*$(t_0 - t_{-1})$: p-value	0.0001	0.0009
Job Destr.*$(t_0 - t_{-2})$: p-value	0.0009	0.0070
Observations	58,632	58,632
R-squared	0.503	0.355

Notes: Models are estimated by OLS. 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. $t_{[-10;+5]}$ denote the time distance in days from the announcement date (t_0). Regressions include firm dummies, time-to-event dummies (in days) from t_{-45} to t_{+10}, with t_{-3} used as a reference, as well as year, month and weekday 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 company level in parentheses. *** p<0.01, ** p<0.05.

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	-0.039***	0.043***	-0.879***
	(0.005)	(0.004)	(0.060)
Job Destruction*t_0	0.022***	-0.022***	0.519***
	(0.004)	(0.003)	(0.042)
Job Destruction*t_{+1}	0.022***	-0.014***	0.404***
	(0.004)	(0.003)	(0.041)
Job Destruction*t_{+2}	0.016***	-0.013***	0.308***
	(0.004)	(0.003)	(0.043)
Job Destruction*t_{+3}	0.013***	-0.008**	0.231***
	(0.004)	(0.003)	(0.045)
Job Destruction*t_{+4}	0.007	-0.012***	0.224***
	(0.004)	(0.003)	(0.045)
Job Destruction*t_{+5}	0.008	-0.008**	0.171***
	(0.005)	(0.003)	(0.045)
Job Destruction*t_{-1}	0.008**	-0.004	0.088**
	(0.004)	(0.003)	(0.044)
Job Destruction*t_{-2}	0.002	-0.007**	0.074
	(0.004)	(0.003)	(0.041)
Job Destruction*t_{-3}	*ref*	*ref*	*ref*
	-	-	-
Job Destruction*t_{-4}	0.002	-0.001	0.070
	(0.005)	(0.003)	(0.039)
Job Destruction*t_{-5}	-0.002	-0.004	0.057
	(0.004)	(0.004)	(0.042)
Job Destruction*t_{-6}	-0.003	-0.001	0.015
	(0.004)	(0.004)	(0.040)
Job Destruction*t_{-7}	-0.002	-0.001	-0.009
	(0.004)	(0.003)	(0.044)
Job Destruction*t_{-8}	-0.001	-0.001	0.021
	(0.004)	(0.003)	(0.039)
Job Destruction*t_{-9}	-0.000	-0.001	0.036
	(0.004)	(0.004)	(0.043)
Job Destruction*t_{-10}	-0.005	-0.001	0.009
	(0.004)	(0.004)	(0.043)
Job Destr.*$(t_0 - t_{-1})$: p-value	0.0003	0.0000	0.0000
Job Destr.*$(t_0 - t_{-2})$: p-value	0.0000	0.0000	0.0000
Observations	39,974	39,974	37,193
R-squared	0.471	0.448	0.612

Notes: Models are estimated by OLS. 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 (standardised by the sum of positive and negative words) in case of job creations and to the excess number of negative words (standardised 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. $t_{[-10;+5]}$ denote the time distance in days from the announcement date (t_0). Regressions include firm dummies, time-to-event dummies (in days) from t_{-45} to t_{+10}, with t_{-3} used as a reference, as well as year, month and weekday 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 company level in parentheses. *** p<0.01, ** p<0.05.

Table 3 – Differential impact of job-destruction vs job-creation announcements on the VADER score

	$Y^{(4)}$ − JobDestr*VADER + JobCreat*VADER
Job Destruction	-0.297***
	(0.024)
Job Destruction*t_0	0.165***
	(0.018)
Job Destruction*t_{+1}	0.137***
	(0.017)
Job Destruction*t_{+2}	0.104***
	(0.018)
Job Destruction*t_{+3}	0.084***
	(0.019)
Job Destruction*t_{+4}	0.080***
	(0.019)
Job Destruction*t_{+5}	0.067***
	(0.019)
Job Destruction*t_{-1}	0.063***
	(0.019)
Job Destruction*t_{-2}	0.043**
	(0.018)
Job Destruction*t_{-3}	*ref*
	-
Job Destruction*t_{-4}	0.028
	(0.018)
Job Destruction*t_{-5}	0.006
	(0.018)
Job Destruction*t_{-6}	-0.008
	(0.019)
Job Destruction*t_{-7}	0.001
	(0.019)
Job Destruction*t_{-8}	0.010
	(0.018)
Job Destruction*t_{-9}	0.019
	(0.019)
Job Destruction*t_{-10}	-0.004
	(0.019)
Job Destr.*$(t_0 - t_{-1})$: p-value	0.0000
Job Destr.*$(t_0 - t_{-2})$: p-value	0.0000
Observations	39,974
R-squared	0.518

Notes: The model is estimated by OLS. The dependent variable $Y^{(4)}$ is the VADER score in case of job creation and the opposite of the VADER score in case of job destruction. *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. $t_{[-10;+5]}$ denote the time distance in days from the announcement date (t_0). Regressions include firm dummies, time-to-event dummies (in days) from t_{-45} to t_{+10}, with t_{-3} used as a reference, as well as year, month and weekday 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 regression although not reported here. Standard errors clustered at the company level in parentheses. *** p<0.01, ** p<0.05.

Appendix Figures

Figure A1. Event Study – Number of First-Level Replies

Correcting for 1 outlier unrelated to the event

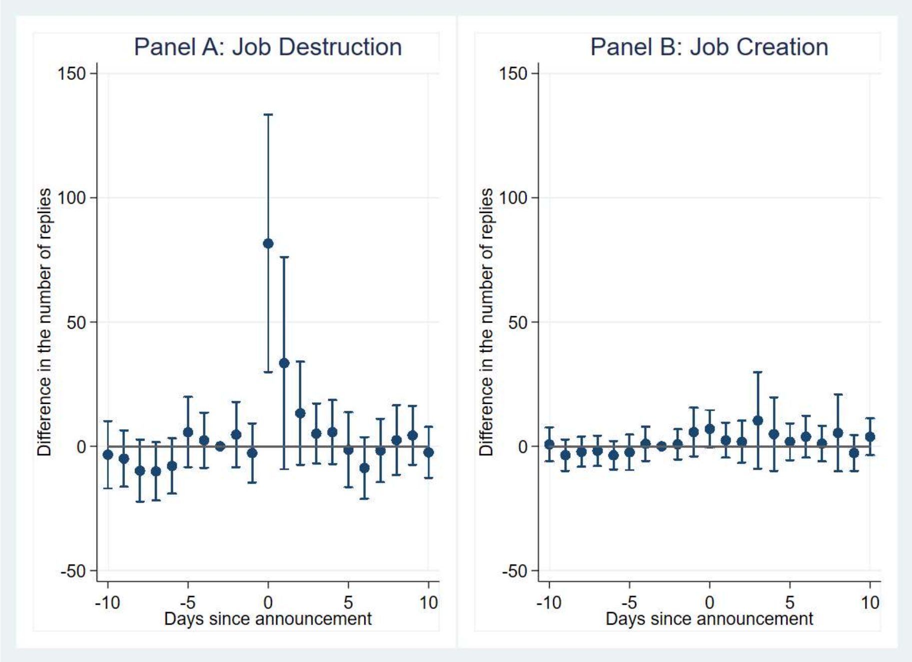

Notes: This graph reports estimated coefficients on time-to-event dummies (in days) from *t*=-10 to *t*=+10, with *t*=-3 used as a reference, obtained by estimating equation (1) by OLS. The dependent variable is the daily number of first-level replies. One tweet with 23,648 replies has been dropped at *t*=+6. Regressions also include firm dummies, time-to-event dummies (in days) from *t*=-45 to *t*=-11, as well as year, month and weekday dummies. Standard errors are clustered at the company level. Error bars correspond to 95% confidence intervals.

Appendix Tables

Table A.1: Number of events and firms

	Number of events	Number of distinct firms
All events	1,047	813
Job destructions	532	452
Job creations	515	403

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	435.76	837.92	36	130	200	400	11,000
Job destructions	468.73	781.40	36	142	236.5	410	11,000
Job creations	401.70	888.72	40	116	200	350	9,400
Panel B : equal or below median size of all events							
All events	135.74	39.69	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	1120.94	203	300	400	689	11,000
Job destructions	742.98	1130.64	203	283	400	678	11,000
Job creations	780.88	1109.83	205	300	430	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	190.0	1026.51	0	93,827
Number of first level replies	23,940	51.64	382.41	0	33,223
RatioNeg	15,808	0.030	0.034	0	0.5
RatioPos	15,808	0.072	0.054	0	0.513
RatioDiff	14,612	0.417	0.470	-1	1
VADER score	15,808	0.146	0.213	-0.960	0.970
Panel B : t = 0 to t = +10					
Number of tweets	5,852	232.81	937.32	0	26,580
Number of first level replies	5,852	62.39	327.22	0	12,096
RatioNeg	4,618	0.047	0.044	0	0.4
RatioPos	4,618	0.058	0.050	0	0.484
RatioDiff	4,341	0.141	0.574	-1	1
VADER score	4,618	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 average ratio of negative to total words, computed for each tweet. *RatioPos* is equal to the average ratio of positive to total words, computed for each tweet. *RatioDiff* is equal to the average ratio 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	23,175	199.59	602.25	0	26,373
Number of first level replies	23,175	57.52	199.70	0	7,194
RatioNeg	15,269	0.026	0.029	0	0.333
RatioPos	15,269	0.082	0.059	0	0.667
RatioDiff	14,260	0.500	0.423	-1	1
VADER score	15,269	0.183	0.212	-0.929	0.977
Panel B : t = 0 to t = +10					
Number of tweets	5,665	205.02	576.27	0	10,376
Number of first level replies	5,665	65.61	396.60	0	25,226
RatioNeg	4,279	0.023	0.027	0	0.333
RatioPos	4,279	0.085	0.060	0	0.597
RatioDiff	3,980	0.562	0.413	-1	1
VADER score	4,279	0.198	0.210	-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 average ratio of negative to total words, computed for each tweet. *RatioPos* is equal to the average ratio of positive to total words, computed for each tweet. *RatioDiff* is equal to the average ratio 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 – Controlling for Firm*Time-to-event FE

	(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	(6) $Y^{(4)}$ – JobDestr*VADER + JobCreat*VADER
Job Destruction	24.97	8.95	-0.037***	0.036***	-0.845***	-0.277***
	(70.70)	(20.30)	(0.007)	(0.004)	(0.080)	(0.028)
Job Destruction*t_0	554.01**	128.89***	0.019**	-0.014***	0.461***	0.127***
	(230.46)	(46.38)	(0.010)	(0.004)	(0.108)	(0.039)
Job Destruction*t_{+1}	118.78	14.79	0.026***	-0.014***	0.454***	0.136***
	(63.06)	(16.11)	(0.008)	(0.003)	(0.096)	(0.033)
Job Destruction*t_{+2}	-34.03	1.99	0.017**	-0.010**	0.357***	0.100***
	(37.62)	(12.08)	(0.008)	(0.004)	(0.092)	(0.036)
Job Destruction*t_{-1}	-38.18	-17.90	-0.004	0.015	-0.010	-0.024
	(39.47)	(10.31)	(0.008)	(0.008)	(0.111)	(0.034)
Job Destruction*t_{-2}	6.77	17.66	-0.007	0.010	-0.088	0.007
	(17.61)	(23.50)	(0.007)	(0.006)	(0.072)	(0.024)
Job Destruction*t_{-3}	*ref*	*ref*	*ref*	*ref*	*ref*	*ref*
	-	-	-	-	-	-
Job Destruction*t_{-4}	-34.11	-4.07	0.013	0.001	0.043	0.055
	(40.46)	(14.32)	(0.007)	(0.009)	(0.064)	(0.045)
Job Destruction*t_{-5}	-28.73	24.84	-0.002	0.006	0.000	0.003
	(35.12)	(17.47)	(0.005)	(0.007)	(0.088)	(0.024)
Job Destruction*t_{-6}	-73.87	-22.66	-0.008	0.014	-0.033	-0.037
	(57.70)	(12.01)	(0.007)	(0.007)	(0.074)	(0.034)
Job Destruction*t_{-7}	-20.73	-10.21	-0.006	0.013	-0.121	-0.021
	(35.06)	(12.44)	(0.007)	(0.005)	(0.109)	(0.032)
Job Destruction*t_{-8}	29.99	-7.11	-0.007	0.008	0.030	-0.012
	(44.70)	(19.01)	(0.007)	(0.007)	(0.063)	(0.030)
Job Destr.*$(t_0 - t_{-1})$: p-value	0.0159	0.0058	0.0017	0.0003	0.0001	0.0001
Job Destr.*$(t_0 - t_{-2})$: p-value	0.0217	0.0437	0.0071	0.0000	0.0000	0.0005
Observations	58,632	58,632	39,974	39,974	37,193	39,974
R-squared	0.649	0.509	0.579	0.441	0.609	0.548

Notes: Models are estimated by OLS. 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 (standardised by the sum of positive and negative words) in case of job creations and to the excess number of negative words (standardised by the sum of positive and negative words) in case of job destructions. In column (6), the dependent variable $Y^{(4)}$ is equal to the VADER score in case of job creations and to the opposite of the VADER score 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. $t_{[-8;+2]}$ denote the time distance in days from the announcement date (t_0). Regressions include firm*distance-to-event dummies (in days) from t_{-45} to t_{+10} with t_{-3} used as a reference, as well as year, month and weekday dummies. Interactions terms between *JD* and time-to-event dummies from t_{-45} to t_{-9} on the one hand and between *JD* and time-to-event dummies from t_{+3} to t_{+10} on the other hand are included in the regressions although not reported here. Standard errors clustered at the company level in parentheses. *** p<0.01, ** p<0.05.

Table A.6 – Differential impact of job-destruction vs job-creation announcements on the number users

	(1) Total Number of Users	(2) Number of Multiple Users
Job Destruction	-12.40	6.58
	(55.99)	(13.49)
Job Destruction*t_0	178.28***	96.48***
	(49.07)	(19.73)
Job Destruction*t_{+1}	81.78	28.09***
	(42.14)	(8.45)
Job Destruction*t_{+2}	40.55	10.41
	(26.85)	(6.84)
	(12.80)	(4.52)
Job Destruction*t_{-1}	-11.00	-0.75
	(13.36)	(5.70)
Job Destruction*t_{-2}	12.49	9.59
	(15.99)	(8.93)
Job Destruction*t_{-3}	*ref*	*ref*
	-	-
Job Destruction*t_{-4}	3.26	0.52
	(11.85)	(4.45)
Job Destruction*t_{-5}	9.14	0.34
	(14.08)	(5.02)
Job Destruction*t_{-6}	-13.55	-5.14
	(12.80)	(5.18)
Job Destruction*t_{-7}	-7.10	0.73
	(15.43)	(6.16)
Job Destruction*t_{-8}	-6.75	2.17
	(14.67)	(6.17)
Job Destr.*$(t_0 - t_{-1})$ – p-value	0.0000	0.0000
Job Destr.*$(t_0 - t_{-2})$ – p-value	0.0005	0.0000
Observations	58,632	58,632
R-squared	0.518	0.409

Notes: Models are estimated by OLS. In column (1), the dependent variable is the daily number of individuals who have tweeted mentioning the company name. In column (2), the dependent variable is the daily number of individuals 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. *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. $t_{[-45;+10]}$ denote the time distance in days from the announcement date (t_0). Regressions include firm dummies, time-to-event dummies (in days) from t_{-45} to t_{+10}, with t_{-3} used as a reference, as well as year, month and weekday dummies. Interactions terms between *JD* and time-to-event dummies from t_{-45} to t_{-9} on the one hand and between *JD* and time-to-event dummies from t_{+3} to t_{+10} on the other hand are included in the regression although not reported here. Standard errors clustered at the company level in parentheses. *** p<0.01, ** p<0.05.

Table A.7 – Differential impact of job-destruction vs job-creation announcements – According to size of event.

	(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	(6) $Y^{(4)}$ – JobDestr*VADER + JobCreat*VADER
Panel A – Above median						
Job Destruction*t_0	382.62***	128.42***	0.028***	-0.020***	0.522***	0.183***
	(128.67)	(47.95)	(0.006)	(0.004)	(0.056)	(0.023)
Observations	28,224	28,224	21,529	21,529	20,322	21,529
R-squared	0.4692	0.3145	0.4299	0.4087	0.5811	0.4870
Panel B – At or below median						
Job Destruction*t_0	41.29**	10.97	0.015**	-0.024***	0.489***	0.136***
	(17.76)	(8.08)	(0.006)	(0.005)	(0.068)	(0.028)
Observations	30,408	30,408	18,445	18,445	16,871	18,445
R-squared	0.754	0.751	0.512	0.494	0.653	0.556

Notes: Models are estimated by OLS. 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 (standardised by the sum of positive and negative words) in case of job creations and to the excess number of negative words (standardised by the sum of positive and negative words) in case of job destructions. In column (6), the dependent variable $Y^{(4)}$ is equal to the VADER score in case of job creations and to the opposite of the VADER score 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. $t_{[-45;+10]}$ denote the time distance in days from the announcement date (t_0). Regressions include a dummy variable for Job Destruction, firm dummies, time-to-event dummies (in days) from t_{-45} to t_{+10}, with t_{-3} used as a reference, as well as year, month and weekday 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 included in the regressions although not reported here. Standard errors clustered at the company level in parentheses. *** p<0.01, ** p<0.05.