

Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.

ISSN 1936-2854 (Print)

ISSN 2767-3898 (Online)

Measuring Job Loss during the Pandemic Recession in Real Time with Twitter Data

**Anbar Aizenman, Connor M. Brennan, Tomaz Cajner, Cynthia Doniger, Jacob
Williams**

2023-035

Please cite this paper as:

Aizenman, Anbar, Connor M. Brennan, Tomaz Cajner, Cynthia Doniger, and Jacob Williams (2023). "Measuring Job Loss during the Pandemic Recession in Real Time with Twitter Data," Finance and Economics Discussion Series 2023-035. Washington: Board of Governors of the Federal Reserve System, <https://doi.org/10.17016/FEDS.2023.035>.

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.

Measuring Job Loss during the Pandemic Recession in Real Time with Twitter Data*

Anbar Aizenman

Anbar.Aizenman@gmail.com

Connor M. Brennan

Connor.M.Brennan@frb.gov

Tomaz Cajner

Tomaz.Cajner@frb.gov

Cynthia Doniger

Cynthia.L.Doniger@frb.gov

Jacob Williams

Jacob.M.Williams@frb.gov

May 22, 2023

Abstract

We present an indicator of job loss derived from Twitter data, based on a fine-tuned neural network with transfer learning to classify if a tweet is job-loss related or not. We show that our Twitter-based measure of job loss is well-correlated with and predictive of other measures of unemployment available in the official statistics and with the added benefits of real-time availability and daily frequency. These findings are especially strong for the period of the Pandemic Recession, when our Twitter indicator continues to track job loss well but where other real-time measures like unemployment insurance claims provided an imperfect signal of job loss. Additionally, we find that our Twitter job loss indicator provides incremental information in predicting official unemployment flows in a given month beyond what weekly unemployment insurance claims offer.

Keywords: Job Loss, Natural Language Processing, Neural Networks.

JEL Classification: J63.

*We thank Elizabeth Vrankovich for help in obtaining Twitter data and Anderson Monken for providing useful insight regarding machine learning implementation. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors.

1 Introduction

Official statistics on unemployment and job displacement lag behind the actual state of the labor market. The Current Population Survey (CPS), which provides detailed information on job loss in the United States as well as underlies the official statistics on employment, is typically released during the first week after the end of each month, and thus lags about four weeks behind economic conditions, and only shows a snapshot of when the survey was conducted. In addition, statistics may be distorted by sudden catastrophic events, such as the shutdowns at the onset of the COVID-19 pandemic. For example, [Ward and Edwards \(2020\)](#) and [Rothbaum and Bee \(2021\)](#) document that shutdowns non-randomly altered patterns of non-response to official government surveys, such as the CPS. State unemployment insurance (UI) claims are released weekly with little lag and at a weekly frequency. However, UI claims likely provided a distorted picture of job losses during the Pandemic Recession because of processing delays, errant claim duplication, and fraud ([Cajner, Figura, Price, Ratner and Weingarden, 2020](#)).

Meanwhile, knowledge about the state of the labor market in as close to real time as possible is valuable to policy makers and markets. For example, official unemployment statistics are legislated as triggers in “automatic stabilization” policies, such as UI expansions. As such, prescience regarding the likely realization of these statistics can help policy makers ensure that appropriate resources are in place. As another example, [Boyd, Hu and Jagannathan \(2005\)](#) document movements in financial markets associated with surprises in these official releases.

In this paper, we construct low latency and high-frequency proxies for job loss derived from Twitter data. We then evaluate the performance of these measures over the period from January 1, 2015 to March 18, 2023, which includes the Pandemic Recession. Twitter is a popular social media platform that allows users to post short messages known as “tweets”. We hypothesize that tweets discussing job loss and unemployment can provide timely information on the current state of the labor market.

The approach is well preceded. Specifically, we build on the methodology

of [Antenucci, Cafarella, Levenstein, Ré and Shapiro \(2014\)](#), who used job-related phrases in tweets to create indexes of labor market flows. In particular, they use job-related phrases in Tweets to create indexes of job loss, job search, and job posting and find that the job loss index can track and predict initial claims for unemployment insurance better than official data and can capture the effects of economic shocks in real time. Similarly, [Proserpio, Counts and Jain \(2016\)](#) use keyword searches refined with a dictionary-based approach to develop a behavioral macroeconomic model that predicts levels of the U.S. unemployment. They found that their psychological well-being measures were leading indicators, predicting economic indices weeks in advance with higher accuracy than traditional forecast techniques.

However, we show that the signal contained in tweet data can be improved by refining tweets to those germane to actual job losses of the tweeting individuals and their social network. Considering tweets as job-loss related based on containing keywords alone (such as “lost job”, “laid off”, or “pink slip”) is not sufficient as they may contain posts unrelated or not reflective of real job losses, such as jokes or the job losses of celebrities. Our method contributes to the literature by using machine learning to filter out this noise. We use natural language processing (NLP) techniques to identify and analyze tweets related to job loss from 2015 to 2023, covering the period before the Pandemic Recession, the recession, and its aftermath. Our Twitter-based measures of job loss demonstrate a high degree of correlation with official statistics from the CPS, as well as UI claims and JOLTS layoffs and could be constructed in near to real time. We then test if our Twitter-based measure of job loss can provide additional explanatory power in predicting current-month employment-to-unemployment flows in CPS above what initial claims—the typical real-time indicator—provide.

Our main contributions are as follows: First, we corroborate the feasibility and validity of using Twitter data to measure job loss. Second, we provide new insights into the dynamics, heterogeneity, and real-time traceability of job loss by gender. Third, we explore the potential applications and limitations of our approach for future research and policy.

2 Data

With the access to the historical Twitter data through the Twitter Enterprise Full-Archive Search API, we obtained a sample of tweets based on search terms included in a tweet’s text. In particular, our sample is based on the following two search criteria, which are intentionally kept simple: i) “laid off” or “layoff”, and ii) “lost job”. For the first criterion we match on the exact keywords, while for the second we match tweets where “lost” and “job” are no more than six tokens apart (e.g., “lost my job”). Our data are daily and go from January 1, 2015 to March 18, 2023.¹ We restrict the sample to tweets from the U.S., that is, tweets that tagged a place in the U.S. or tweets from a user who listed the U.S. as where they reside in their Twitter profile. We exclude retweets, because we are primarily interested in people reporting their own personal job loss. Table 1 shows the count of tweets containing the job-loss related keywords (Table 2 shows the exact Twitter API queries used). In total, we have 2.1 million tweets.

Table 1: Frequency of Different Twitter Queries (January 1, 2015 to March 18, 2023)

Query Term	Queried Tweets
laid off or layoff	1,220,164
laid off	901,878
layoff	318,286
lost job	919,369

Source: Twitter and authors’ calculations.

We further refine our initial query by filtering out tweets posted by users with a follower-to-following ratio either greater than 10 or lower than 0.1. The rationale is that we would like to focus on tweets of “typical” individuals and thus we aim to remove tweets likely posted by large corporations and bots. After applying this restriction, our sample size is reduced from 2.1 million to 1.7 million.

¹While Twitter data have a few more years of earlier history, we found that the period before 2015 has a more limited number of tweets available and is thus less suitable for our analysis.

Table 2: Twitter API Queries

Group	Query
Lay off	("laid off" OR layoff) (profile_country:us OR place_country:us) -is:retweet
Lost job proximity	("lost job"~6) (profile_country:us OR place_country:us) -is:retweet

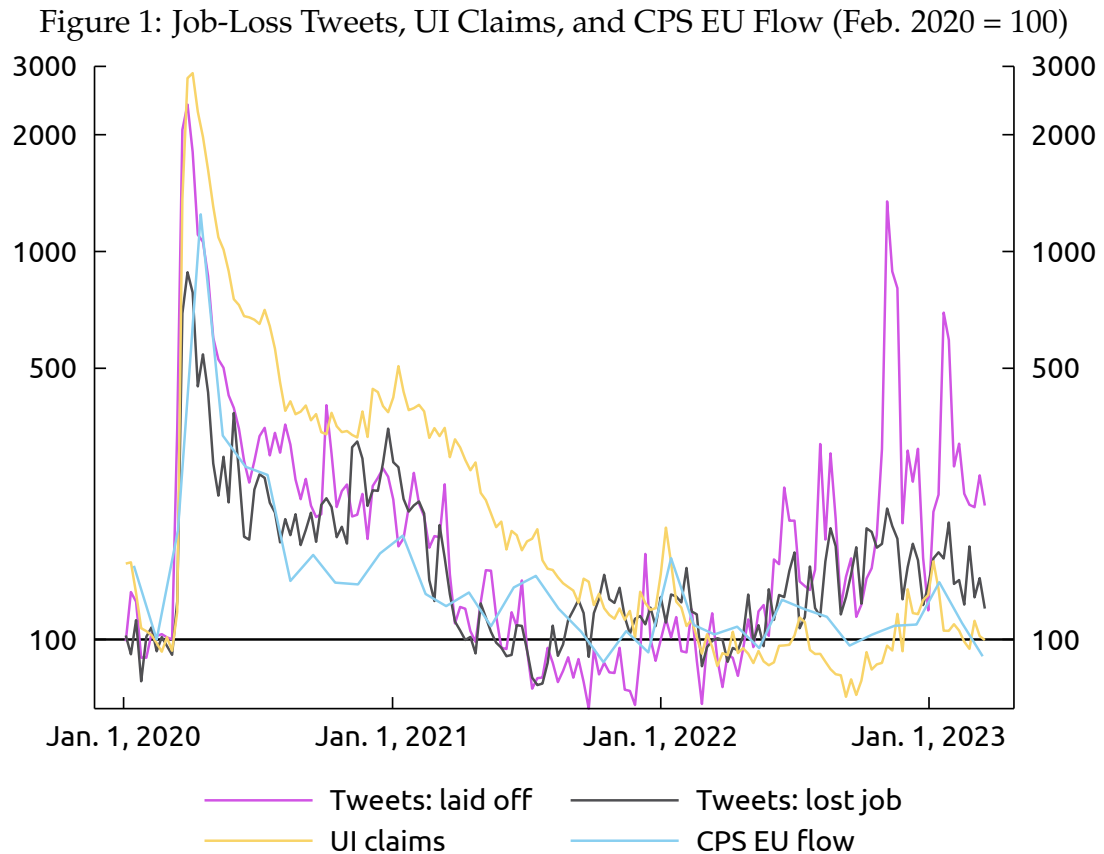
3 Analysis of Raw Job-Loss Twitter Data

We proceed with the analysis of the raw job-loss Twitter data. In Figure 1, we plot the weekly number of tweets for our two search criteria, together with official statistics on initial claims for unemployment insurance (UI claims) and labor market flows from employment-to-unemployment (EU flow) as measured in the Current Population Survey (CPS) data. For a more direct comparison, we use non-seasonally adjusted data.²

As can be seen in Figure 1, UI claims spiked roughly by a factor of 30, rising from levels around 200,000 in February 2020 to an unprecedented 6.2 million in the week ending April 4, 2020. Interestingly, the number of tweets containing the phrase "laid off" rose by almost exactly the same relative amount, while the number of tweets with "lost job" rose somewhat less. As Twitter data are available in real-time, they provided some time advantage when compared to UI claims. In particular, the UI claims data for the week ending March 21 — which showed an unprecedented tenfold increase to 3.3 million — were published on March 26. In contrast, the rapid increase in job loss was evident in Twitter data already as of Monday, March 16, providing a time advantage of 10 days. Moreover, UI claims were overstating true job losses, especially in the period from April 2020 to about

²The large swings observed during the COVID pandemic introduced several challenges for seasonally adjusting data, especially for methods that rely on multiplicative adjustment. Additionally, typical seasonal patterns were swamped by the observed movements during the COVID pandemic.

April 2021, as evidenced by CPS EU flow being substantially below UI claims. Interestingly, both Twitter-based job-loss measures followed CPS EU flow data relatively well during that period, suggesting they provide useful information about job loss.



Note: Non-seasonally adjusted data plotted. Ratio scale used. Job-loss tweets are weekly sum of tweets, indexed to their average in February 2020 being equal to 100.
 Source: Twitter, BLS, CPS, and authors' calculations.

3.1 Gender Analysis

Next, we construct Twitter-based job loss measures by gender. To do so, we first download counts of names by gender and year going back to 1900 for the U.S.

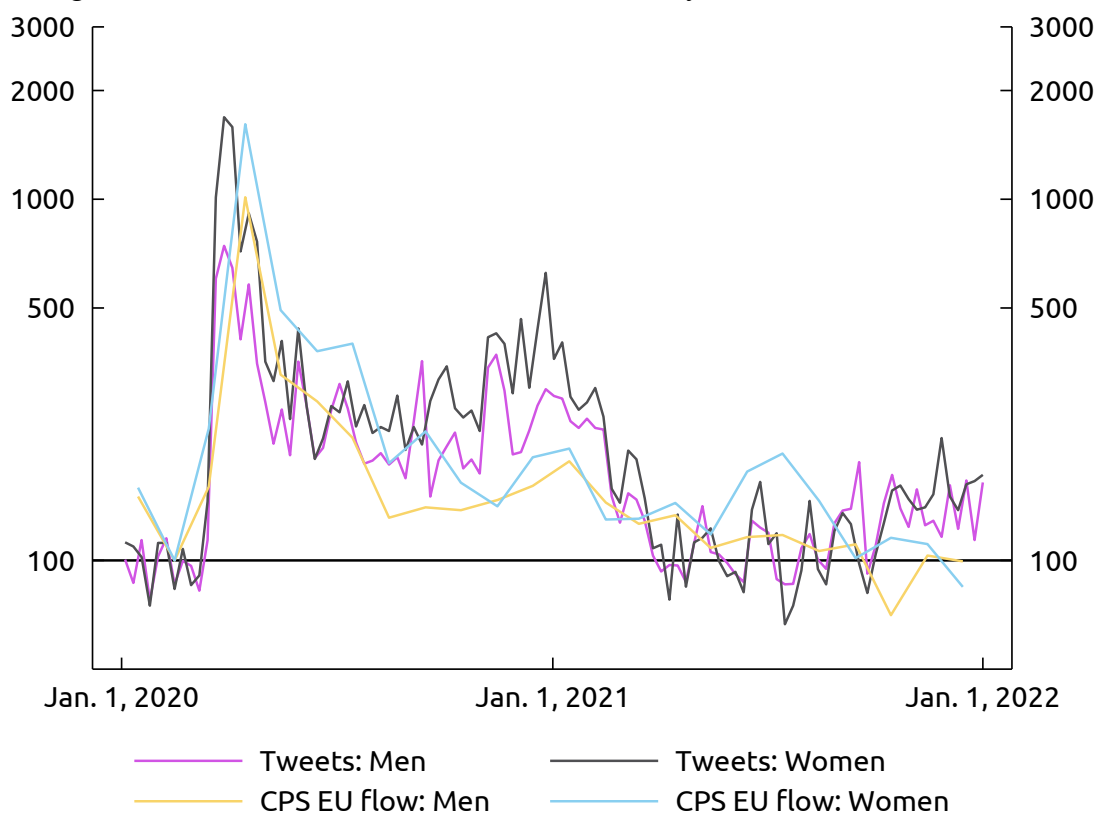
from the Social Security Administration website.³ The data contain all names except for those with fewer than five occurrences in any given year. We sum names by gender from 1900 to 2021 and to exclude names that are used for both genders, we only keep names that have a 95 percent or greater occurrence in a single gender following the method from [Mislove, Lehmann, Ahn, Onnela and Rosenquist \(2011\)](#). To clean the Twitter name field, we remove titles from the name that may be mistaken for first names as well as replace characters often used in place of letters in names.⁴ We then extract the first word from the name field and match it to the 95-percent threshold name gender list from above. Overall, we were able to assign gender to about 50 percent of tweets in our sample.

One notable empirical pattern observed during the pandemic recession was that the increase in the unemployment rate was larger for women than men, in sharp contrast with the typical pattern observed during recessions from 1980 to 2010. In particular, while the unemployment rate for both men and women equaled 3.5 percent in January 2020, it increased in April 2020 to 16.2 percent for women and 13.5 percent for men. By December 2020, the unemployment rate for both genders was again equal at 6.7 percent. In order to examine whether the same pattern exists also in Twitter-based job loss data, [Figure 2](#) plots job-loss tweets by gender (because the sample is now smaller given that we are able to assign gender to only about 50 percent of tweets, we join tweets for both of our search terms together). Interestingly, job loss tweets from women jumped more than those from men in the early stages of pandemic, consistent with CPS EU flow data by gender. Additionally, job loss tweets for both genders were roughly equal from early 2021 onward, consistent with the convergence in CPS EU flow data as well (and the convergence in published unemployment rates).

³See <https://www.ssa.gov/oact/babynames/names.zip>.

⁴The excluded words include: king, prince, the, dr, sir, mr, ms, mrs, lil, rev, fr, father, queen, princess, lord, mr., ms., mrs., brother, sister, little, doc, sir, and professor. The characters and substitutions include: \$:s, @:a, 4:a, !:l, 8:b, 0:o, 3:e, /:l, and |:l.

Figure 2: Job-Loss Tweets and CPS EU Flow by Gender (Feb. 2020 = 100)



Note: Non-seasonally adjusted data plotted. Ratio scale used. Job loss tweets are weekly sum of tweets, indexed to their average in February 2020 being equal to 100.
 Source: Twitter, CPS, and authors' calculations.

4 Using Machine Learning to Obtain More Precise Job Loss Twitter Measure

So far we have been using raw counts of job-loss tweets. While these raw counts based on relatively simple search terms do a pretty good job at measuring job loss in real time as argued in the previous section, the Twitter queries we gathered might not yield exclusively tweets that are capturing job loss. To obtain a more precise job-loss Twitter measure, we proceed by using a machine learning model for NLP called “BERT” or Bidirectional Encoder Representations from Transform-

ers. BERT is a pre-trained neural network developed by researchers at Google in 2018 and described in [Devlin, Chang, Lee and Toutanova \(2018\)](#). BERT’s key technical innovation is applying bidirectional training to language modelling: instead of reading text sequentially (left-to-right or right-to-left), BERT reads the entire sequence of words at once, allowing the model to learn the context of a word based on all of its surroundings. Understanding the context of a word is crucial for determining whether a tweet is actually related to job loss or not despite containing a job-loss related phrase. For example, consider the tweet: “No one lost their job”. Context is needed to recognize that “no one” is important when interpreting “lost their job”. Otherwise, the model would incorrectly classify the tweet as related to an actual job loss when it is not. In addition to its bidirectional training, BERT also comes pre-trained on all of Wikipedia (2.5 billion words) and Book Corpus (800 million words), giving it a fair understanding of sentence structure and writing.

We fine-tune the BERT model with an additional output layer that will classify tweets as related to an actual job loss or not. We create training and validation data by labeling 6,011 randomly selected tweets from our query. In addition to binary labeling whether tweets generally relate to job loss or not, we also mark whether a tweet is present-tense, meaning relating to a job loss that just happened, and whether a tweet relates to a celebrity rather than a everyday worker (often people tweet about famous individuals losing their job and we would like to exclude such tweets). Tweets are labeled a “real” job loss if they are both related to job loss and present tense but not related to a famous person, for a total of 37.5 percent of our labeled data. We then split 70 percent of our data to a training set, 15 percent to a test set, and 15 percent to a validation set.

Figure 3 shows both the model’s decreasing loss over successive rounds of training and a confusion matrix of our model’s classifications on our validation set. Over each epoch of training, we retain a given epoch of training if the validation loss of that epoch is better than the validation loss of the previously retained epoch. Early stopping of the training is executed after 15 consecutive epochs of non-decreasing validation loss relative to the last retained epoch of training. In total, we train 118 epochs into our outer layer, and the model is able to correctly

predict approximately two-thirds of the time. Table 3 presents the model’s classification report. With accuracy of approximately 70 percent, our model is considered acceptable by most machine learning standards. Given that our groupings are a bit unbalanced, our reasonable F1-scores give us confidence in our model’s performance.

Figure 3: Model Training and Evaluation

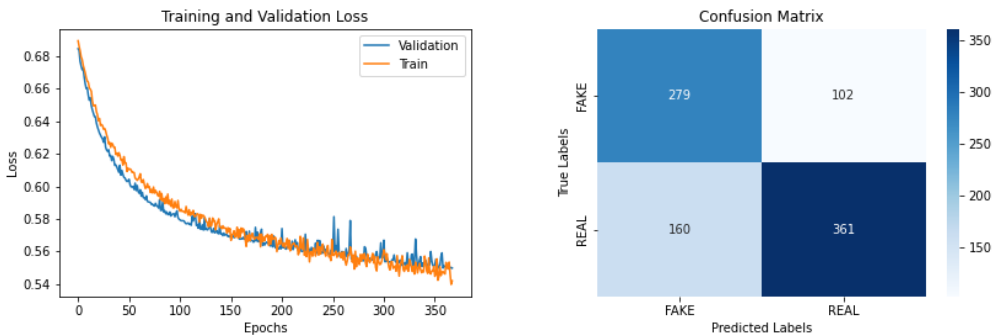
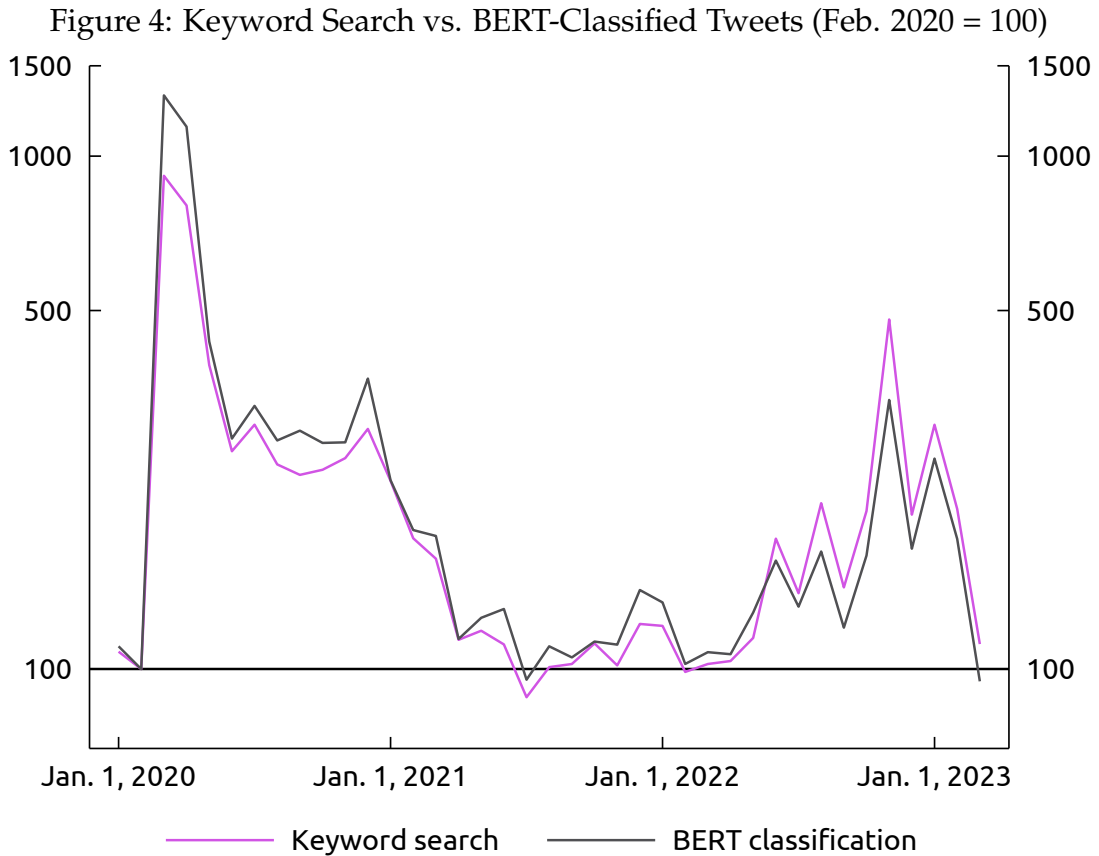


Table 3: Model Classification Report

True Label	Precision	Recall	F1 Score	Support
0	0.760	0.687	0.722	521
1	0.622	0.703	0.660	381
Macro Average	0.691	0.695	0.691	902
Weighted Average	0.702	0.694	0.696	902
Accuracy	0.694			

Figure 4 shows raw job-loss tweet counts together with the data based on the BERT classification algorithm. Note that in absolute numbers (not shown), the BERT-classified job-loss tweet series are always below raw job-loss tweets counts (by about one third). However, Figure 4 plots the data indexed to February 2020 being equal to 100. The results show that the BERT-classified series actually jumped more during the Pandemic Recession than the raw data. On the other hand, during the period of late 2022 and early 2023, the BERT-classified series is below the raw data. Note that a spike in Twitter-based job loss measure during

that period to a large extent reflects layoffs in the tech sector. The BERT algorithm correctly predicts that many of those tweets are not about individual people reporting their own job loss, but rather about many tweets discussing relatively isolated layoff events in the tech sector. However, even the BERT-classified series seems somewhat elevated to information we have from UI claims and CPS data, suggesting that further improvements in making the Twitter-based job-loss measure more precise through a further fine-tuned BERT model are possible.



Note: Non-seasonally adjusted data plotted. Ratio scale used. Value are indexed to their average in February 2020 being equal to 100.
 Source: Twitter and authors' calculations.

Table 4 shows how well our Twitter job-loss indicator correlates to official measures of unemployment flows compared to counts of tweets containing keywords

for job loss. Over our whole sample, we note that our BERT-classified tweets are markedly more correlated to the CPS, JOLTS, and UI claims data. This suggests that our indicator more stably tracks job loss over a wider horizon. We also note that these correlations jump during the COVID pandemic, suggesting that Twitter may provide a useful alternative measure of job losses during crises, when other data like UI claims may have issues.

Table 4: Correlations of Twitter job-loss indicators to official job-loss indicators

	BERT Classified Tweets		Tweets by Job Loss Keywords	
	Whole Sample	COVID	Whole Sample	COVID
CPS	0.943	0.986	0.886	0.983
JOLTS	0.859	0.898	0.814	0.893
UI claims	0.948	0.971	0.906	0.980

Notes: *CPS* refers to the Employed to Unemployed flow from the Current Population Survey, from the week containing the twelfth of the month from last month to the week containing the twelfth of the month in the current month. *JOLTS* refers to monthly layoffs in JOLTS. *UI claims* refers to unemployment insurance numbers released weekly. *COVID* refers to correlations over the period January 2020 to March 2020 while *Whole Sample* refers to correlations over the period January 2015 to March 2023. Data are of monthly frequency. Source: Twitter, BLS, CPS, and authors' calculations.

5 Forecasting EU Flows With Job-Loss Twitter Data

We estimate four regression models with the monthly data from January 2015 to March 2023 to examine the relationship between employed-to-unemployed CPS flows and our Twitter job-loss measure. The predictors include initial claims, tweets by job-loss keywords, the lag of employed-to-unemployed CPS flows, and tweets classified by BERT as job-loss related. The first model uses only UI claims data and lagged EU flows. The second model includes all predictors except the BERT-classified tweets. The third model includes all predictors. The fourth model includes the BERT-classified tweets but excludes the tweets by job-loss keywords. Results are presented in Table 5. We find that the tweets by job-loss keywords are

insignificant in the second model, suggesting that they contain too much noise from unrelated topics and are dominated by initial claims as a predictor. However, in the third model, both the BERT classified tweets and the tweets by job-loss keywords are significant with the coefficient of BERT-classified tweets being positive whereas the coefficient for job-loss keywords is negative, indicating that the BERT-classification extracts the job-loss signal from the noisy tweets and improves the forecasting power of CPS flows. In contrast, in the fourth model, only the BERT-classified tweets are significant and positive, confirming their informational value in predicting employed-to-unemployed CPS flows.

Table 5: Entire Sample (January 2015 to March 2023)

	(1) Excl. Twitter	(2) Excl. BERT-classified	(3) All	(4) Excl. keywords
UI claims	0.602*** (19.81)	0.578*** (15.88)	0.495*** (11.75)	0.549*** (14.45)
Lag of EU flows	-0.328*** (-4.61)	-0.327*** (-4.77)	-0.323*** (-6.69)	-0.325*** (-5.10)
Keyword tweets		0.0369 (1.03)	-0.514*** (-2.92)	
BERT-classified			0.606*** (3.29)	0.0722* (1.91)
Constant	-0.00266 (-0.15)	-0.00367 (-0.21)	-0.00570 (-0.34)	-0.00470 (-0.27)
Observations	97	97	97	97
R^2	0.715	0.717	0.743	0.721
AIC	-65.88	-64.54	-71.86	-66.14

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

6 Conclusions

This paper develops a measure of job loss using Twitter data, based on a machine learning algorithm called BERT that classifies Tweets as job-loss related or not. We have shown that this measure can track job loss and predict EU flows at high frequency and in real time, and can capture the effects of economic shocks such as the COVID-19 pandemic.

Our paper contributes to the literature on using social media data to measure and analyze labor market, and provides a valuable tool for policy makers and researchers who need timely and accurate indicators of job loss. However, there are several limitations and challenges that need to be addressed in future research. For example:

- How to improve the identification and extraction of job-related phrases from tweets, using natural language processing and machine learning techniques
- How to extend the analysis to other countries and regions, taking into account the differences in language, culture, and labor market institutions
- How to incorporate other sources of social media data, such as Facebook, LinkedIn, or Reddit, to create a more comprehensive picture of labor market dynamics.

References

- ANTENUCCI, D., CAFARELLA, M., LEVENSTEIN, M., RÉ, C. and SHAPIRO, M. D. (2014). *Using Social Media to Measure Labor Market Flows*. NBER Working Paper 20010, University of Michigan.
- BOYD, J. H., HU, J. and JAGANNATHAN, R. (2005). The stock market's reaction to unemployment news: Why bad news is usually good for stocks. *The Journal of Finance*, **60** (2), 649–672.

- CAJNER, T., FIGURA, A., PRICE, B. M., RATNER, D. and WEINGARDEN, A. (2020). *Reconciling Unemployment Claims with Job Losses in the First Months of the Covid-19 Crisis*. Feds working paper, Federal Reserve Board of Governors.
- DEVLIN, J., CHANG, M.-W., LEE, K. and TOUTANOVA, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- MISLOVE, A., LEHMANN, S., AHN, Y.-Y., ONNELA, J.-P. and ROSENQUIST, J. (2011). Understanding the Demographics of Twitter Users. *Fifth International AAAI Conference on Weblogs and Social Media*, 5 (1).
- PROSERPIO, D., COUNTS, S. and JAIN, A. (2016). The psychology of job loss: using social media data to characterize and predict unemployment. In *WEB-Sci'16: Proceedings of the 8th ACM Conference on Web Science*, WEBSCI, Seattle, WA: ACM, pp. 223–232.
- ROTHBAUM, J. and BEE, A. (2021). *Coronavirus Infects Surveys, Too: Survey Nonresponse Bias and the Coronavirus Pandemic*. Tech. rep., US Census.
- WARD, J. and EDWARDS, K. A. (2020). *Statistics in the Time of Coronavirus: COVID-19-Related Nonresponse in the CPS Household Survey*. Tech. rep., Rand.