

The Psychology of Job Loss: Using Social Media Data to Characterize and Predict Unemployment

Davide Proserpio
Boston University, USA
dproserp@bu.edu

Scott Counts
Microsoft Research, USA
counts@microsoft.com

Apurv Jain
Microsoft Research, USA
apurvj@microsoft.com

ABSTRACT

Using data from social media, we study the relationship between the macroeconomic shock of employment instability and psychological well-being. We analyze more than 1.2B Twitter posts from over 230,000 U.S. users who either lost a job or gained a new job over a period spanning five years, from 2010 to 2015. First we quantify the magnitude and length of effects of job loss/gain on psychological variables such as anxiety, sadness, and anger. We then define a behavioral macroeconomic model that leverages these changes in psychological state to predict levels of unemployment in the U.S. Our results show that our psychological well-being measures are leading indicators, predicting economic indices weeks in advance with higher accuracy than baseline models. Taken together, these findings suggest that by capturing the human experience of a shock like job loss, social media data can augment current economic models to generate a better understanding of the overall causes and consequences of macroeconomic performance.

CCS Concepts

•Applied computing → Law, social and behavioral sciences; Economics; Psychology;

Keywords

Social science, social media, economics, unemployment

1. INTRODUCTION

The psychological impact of job loss can be severe, with the unemployed exhibiting higher likelihood of depression and anxiety, social alienation, and hopelessness, as well as a variety of associated physical symptoms including being more sick, taking more medications, and making more doctor visits than employed counterparts [27]. While these effects have been studied dating back to the 1930s [17], connecting aggregate measures with individual effects has remained a challenge. For instance, while positive correlations have been found between the unemployment rate and mortality, heart disease, and heavy drinking [30], it is difficult to show that the people who lost jobs are those contributing to the increases

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WebSci '16, May 22 - 25, 2016, Hannover, Germany

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-4208-7/16/05...\$15.00

DOI: <http://dx.doi.org/10.1145/2908131.2913008>

in these aggregate measures of mental, physical, and behavioral health issues. Further, in the presumed interaction between macroeconomic shock and personal impact, the effect of the individual scale psychological damage of job loss on the aggregated unemployment rate is not yet understood. It makes intuitive sense that those depressed after losing their job would struggle even harder to find a job, but quantifying such effects at population scale has remained an outstanding challenge given the scope of data required. This paper addresses precisely this question: how do the psychological consequences of job loss (or gain) in individuals predict future population level of unemployment?

Thus the value of this work lies in the ability to more precisely estimate and predict the unemployment rate. Such macroeconomic statistics are critical to governments and private industries when making decisions that affect populations at national and even global scale. For instance, a current problem related to unemployment for many governments, including the U.S. is the minimum wage dilemma. Minimum wages are strictly correlated with the unemployment rate (see [9] for a review) and the government therefore relies on the unemployment index when controlling minimum wages. A wrong decision to increase minimum wages, perhaps caused by an underestimation of the unemployment rate, might in fact lead to higher rates of unemployment. Unfortunately, economic indices are difficult to measure. During the recent great recession, the Bureau of Economic Analysis initially projected that the economy shrank at an annual rate of 3.8% in the last quarter of 2008, but months later it almost doubled that rate to 6.2%. Similar difficulties can be found in trying to measure the labor force participation rate, which in turn can radically alter unemployment rate estimates. Again, one can see the role the psychological state of the population of those who have lost a job could play on the work force participation, as the extremely discouraged may simply drop out of the labor force altogether. In general, monetary policy changes require timely, precise measurements of the economy, yet these measurements are today largely based on laborious survey techniques unchanged over many decades.

Behavioral economists underscore how the “realism of the psychological underpinnings of economic analysis will improve economics on its own terms generating theoretical insights, making better predictions of field phenomena, and suggesting better policy” [12]. These psychological micro-foundations of macroeconomics [4], may be increasingly measurable. In this work, we measure the psychological impact of job loss or gain across hundreds of thousands of individuals using data from social media. These data provide a rich characterization, including measures like anxiety and focus on self that are known manifestations of negative psychological states [15]. We first characterize descriptively how these measures are distributed across the population of people who lost

(or gained) a job, showing for example how tentativeness and a focus on causation spike around the job loss event. We then show how measurements of these changes to the psychological state of individuals who lost or gained a job explain more variance in future official population scale unemployment rates than various baseline models. We conclude by returning to the discussion of how such data might improve monetary policy setting.

2. BACKGROUND

2.1 Unemployment and psychological well-being

Macroeconomic shocks such as job loss can be devastating to one's psychological well-being. In a first study of the impact of unemployment on mental health, [17] catalogued the multidimensional and unpleasant emotional consequences of unemployment, showing increases in depression and anxiety, along with strained personal relations. Reviews of the literature on the psychological impact of unemployment e.g., [5] and large-scale meta-analyses [30], [33] document many more effects including psychological effects such as lowered self-esteem, and physiological symptoms including an increased mortality rate. Matched sample [27] and longitudinal studies [40] show similar effects, with the unemployed elevated in negativity, depression, dissatisfaction, hopelessness, and life satisfaction.

Social psychological research on the relation between employment, joblessness and psychological well-being suggests a variety of psychological drivers of the effects. For example [8] and [32] show that unemployment can create a feeling that life is not under one's control, which in turn can affect mental health. Aspects of identity theory [3] may play a role as moderators of the effects of job loss, as certain types of people may view their job as more central to their identity [35] and thus suffer more psychological damage when that job is taken away. Jahoda in [24] and [25] suggests that employment creates latent benefits, such as time structure and connecting individuals to goals that, when taken away, are psychologically destructive. Moreover, the effects of exposure to unemployment can have substantial repercussion on tests for work and job search strategies, which in turn can affect re-employment probability, or productivity should they find a job. For a complete review of the literature the reader can refer to [18].

Economists, especially those from labor economics who focus on the interplay of wages, earnings, employment and productivity, have also examined the role of psychological well-being in this complex equation. In [13], the authors give a sense of how psychological literature can be used by economists to create a behavioral model of the macroeconomy that accounts for the relationship between employment outcomes and psychological well-being. Current Federal Reserve chair Janet Yellen's papers [42] and [1] show that workers efforts are affected by anger, jealousy, and gratitude, and that these behavioral factors can help explain cross-sectional wage differentials and unemployment patterns.

2.2 Economic Measurement

Current macroeconomic statistics are based on data collection methods dating to the beginning of the previous century ([29], [11]). Inflation, for example, is calculated by the Bureau of Labor Statistics by manually collecting 80,000 prices each month. U.S. unemployment is calculated monthly by using the Current Population Survey (CPS) of about 60,000 households. Besides being costly to collect, these data suffer sample limitations, and are difficult to repeat on an ongoing basis with frequency. Surveying tens of thousands of households on a daily basis, for instance,

would be nearly impossible. Further, the statistics computed from these surveys methods often require revisions for months after their first release, primarily because data continues to roll in even after the official data release date ([10], [19], [26]). In summary, current macroeconomic measurement methods suffer in terms of timeliness and breadth of population coverage, reducing the precision with which the study of the impact of economic shocks can be executed.

Augmenting existing measurement processes with new data sources may alleviate these issues. Social media data, in particular from Twitter, because of its free and public nature have shown promise ([2], [28]). For example in [2] the authors created the *Social Media Job Loss Index* which tracks initial claims for unemployment insurance at medium and high frequencies, and predicts 15 to 20 percent of the variance of the prediction error of the consensus forecast for initial claims. The authors do so by creating a dictionary of words correlated with unemployment and tracking the tweets containing such words. Others, such as [38], exploit mobile phone call logs to create a predictive model for unemployment that outperforms traditional forecasts. In this case, the authors take advantage of a natural experiment, a mass layoff due the closure of a large manufacturing plant in a Spanish city, allowing them to measure the effect of unemployment on call activities.

In this paper we leverage social media data first to provide a detailed picture of the psychological impact of job loss and gain at the level of individuals, and then show how these factors can explain variance in the prediction of the macroeconomic measurement of unemployment.

3. DATA

3.1 Twitter dataset

We leverage data from Twitter, a popular online social media network that enables users to send and read 140-character messages called tweets. We chose Twitter because of the large user base of people within the United States who make their tweets publicly available, and because it has been shown to be a good source of information for analyzing human behavior, emotion, and important population events (e.g., [43], [15], [36]). The full Twitter stream was made available to us through an agreement with Twitter. As our study focuses on the U.S. population, we consider only English language Tweets sent from U.S. time zones. We further restrict the sample of users to those that have been identified as gaining a new job or losing a job.

3.1.1 Unemployment data

To identify users that lost their job we search for tweets containing one of the following text strings: "I got laid off" or "I lost my job". To verify that these tweets do in fact reflect people who lost a job, three Mechanical Turk master workers rated 1000 randomly selected tweets for whether the tweet indicated job loss or not. 91% of tweets were rated by at least 2 raters as being about job loss; 87% were rated as about the loss of the job of the tweet author; 2.7% of tweets were unanimously not about the job loss of the tweet author. We then collect two years of tweets history around the candidate job loss declaration tweets for all the users identified as unemployed. The final unemployment dataset contains 51,158 users and about 346M tweets spanning a period of five years, from 2010 to 2015. Note the number of users reporting job loss is only about half of those reporting job gain (below), which may reflect a bias toward reporting good news on social media, but also corresponds with a generally low, and falling, unemployment rate over the time period.

3.1.2 Employment data

To identify a user as part of the employed group we search for tweets containing one of the following phrases: “got a new job”, “started a new job”, “started my new job” and “starting my new job”. As with the unemployment tweets, we had three Amazon Mechanical Turk master workers rate a sample of 1,000 tweets selected at random, indicating whether or not the tweet did reflect a person starting a job. Results showed that 99% of our employment tweets were rated by at least two of the three raters as about the start of a job; 75% of tweets were rated as about the tweet author posting about the start of his or her own job; 18% of tweets were unanimously not about job gain of the tweet author. Compared to the unemployment tweets, these were more difficult to narrow down to just those about the tweet author, primarily because many tweets that were in fact about the author did not specify it in a reliable, automatically detectable manner. For instance, “they closed got a new job tho” and “i moved houses and started a new job” are likely about the authors’ own new jobs, while “kinda miss Warren. guy moved on up and got a new job on me” appears to be about Warren getting a new job. Although not perfect, the percent of tweets that truly reflect job gain is fairly high and given the size of the data, should largely swamp any noise due to the small percent of false positives. After identifying the candidate tweets, we collected the tweets history of these users, one year before and one year after the date at which the candidate tweet was posted. The final employment dataset contains 180,594 users and about 865M tweets spanning a period of five years, from 2010 to 2015.

3.1.3 Random sample

Finally, we collected a random sample of about 250,000 US Twitter users and their entire tweet history for a total of about 221M posts. Table 1 provides summary statistics of the datasets used in this paper.

| Dataset | Users | Tweets |
|--------------|---------|-------------|
| Unemployment | 51,158 | 346,357,649 |
| Employment | 180,594 | 865,307,413 |
| Random | 254,569 | 221,788,781 |

Table 1: Datasets description

3.2 Measures

Our goal is to measure peoples’ psychological well-being by analyzing the content of the Twitter posts collected. We quantify changes in the psychological state of people in our dataset by leveraging the LIWC (Linguistic Inquiry Word Count) dictionaries developed by [34]. Given a Twitter post, we perform a regular expression match to determine the fraction of words in the tweet that match the words in one of our LIWC dictionaries. For a given measure we compute this fraction for every person and take the mean over all tweets that person posted in a given day, yielding a daily value for each person. We repeat the process for each of the measures shown in Table 2, which provides means of our measures for the three sets of people: those who lost a job, those who gained a job, and the random sample.

The LIWC measures generally are self-explanatory: the anxiety measure is a dictionary of words validated to capture anxiety, tentativeness to capture tentativeness, and so on. We did not use all LIWC measures, but selected a subset deemed relevant to (un)employment, including sadness, work, and first person pronouns which are known to detect depressive mood states

[15]. LIWC-based measures have been used to extract psychological variables from Twitter data, including correlating those measures to population scale phenomena (e.g., [43]). For broader validation of LIWC, we refer the reader to [34] or to the LIWC website (liwc.wpengine.com).

| Measures | Employment | Unemployment | Random |
|---------------|------------|--------------|--------|
| Achievement | 0.0138 | 0.0135 | 0.0143 |
| Anger | 0.0130 | 0.0138 | 0.0093 |
| Anxiety | 0.0028 | 0.0026 | 0.0021 |
| Causation | 0.0114 | 0.0117 | 0.0099 |
| Certainty | 0.0125 | 0.0121 | 0.0103 |
| Discrepancy | 0.0185 | 0.0182 | 0.0139 |
| Exclusion | 0.0208 | 0.0206 | 0.0156 |
| Family | 0.0040 | 0.0038 | 0.0033 |
| First Person | 0.0517 | 0.0482 | 0.0355 |
| Friend | 0.0022 | 0.0019 | 0.0018 |
| Home | 0.0045 | 0.0041 | 0.0037 |
| Inclusion | 0.0245 | 0.0239 | 0.0227 |
| Ingroup | 0.0012 | 0.0012 | 0.0014 |
| Leisure | 0.0158 | 0.0157 | 0.0180 |
| Money | 0.0055 | 0.0062 | 0.0067 |
| Sadness | 0.0042 | 0.0043 | 0.0035 |
| Tentativeness | 0.0176 | 0.0178 | 0.0141 |
| Work | 0.0135 | 0.0130 | 0.0135 |

Table 2: Mean values of the psychological variables extracted from the Twitter posts

4. IMPACT OF JOB LOSS OR GAIN ON PSYCHOLOGICAL STATE

4.1 Economic shock effect

We start by analyzing the trend of the measures reported in Table 2. To compare pre-treatment trends to post-treatment trends, we partition time around the month each user lost (or gained) a job in monthly intervals, taking the offset of the economic shock to be 0. Thus the interval 0 represents the month in which the user lost (or gained) his job, the interval -1 the month just before, and the interval 1 the month just after. We then plot the average (over all treated users) interval values for all the measures for a period of 24 months, corresponding to a period of one year before the economic shock, and one year after it. The resulting plots are reported in Figure 1. For several measures, we observe drastic changes around interval 0 (the time of the economic shock). For example anxiety increases for the job losers, while it decreases for the job gainers. Sadness shows a short, but very strong spike at interval 0 for those users that lost a job, while the trend for users that gained a new job is completely flat. Money increases for both types of users. Somewhat paradoxically a number of measures show similar patterns across the employed and unemployed, which can be interpreted to mean that the economic shock, either positive or negative, brings to focus certain psychological factors. This is reflected by the massive spike in the *work* measure precisely at the (un)employment event time.

Interestingly, for some measures, the trends start changing *before* the shock occurred. While we had not anticipated such effects, a bit of reflection offers a potential explanation. In the case of gaining a job, our measures may be detecting confidence in finding a job leading up to the actual employment event. Job search confidence or self-efficacy is strongly associated with the capacity to change one’s situation [7]. In this sense, self-efficacy can be viewed as a

catalyst for an increase in well being, as observed in our metrics. In the case of job loss, deterioration of one's psychological well-being is correlated with poor productivity and negative attitude toward work [14]. It is possible then, that in the period leading up to job loss, an individual would show a change in metrics related to his or her well-being.

Broadly, the clearly notable changes observed in monthly behavior of individuals with a high probability of being laid off or obtaining a new job is compelling evidence that, using Twitter data, we are able to identify people that experienced an economic shock, and that such shock appears to have an effect on their psychological well-being.

4.2 Difference in differences

We generalize the previous findings by analyzing psychological changes associated with employment or unemployment at the individual level relative to a control group. To do so we compare changes in the psychological measures for people exposed to job loss or job gain (the treated group) between the post-treatment and pre-treatment period with respect to a baseline of changes for a randomly selected set of Twitter users (the control group) over the same period of time.

Note that there is an important selection issue we have to take into account to ascribe a causal interpretation of the results. The Twitter users selection into the treatment - the person gaining or losing a job - is not random. Therefore convincingly estimating the average treatment effect (ATE), which would be the impact of macroeconomic shock on a randomly selected population, is not possible. Thus, we focus on estimating the average treatment on the treated effect (ATT). The ATT can be consistently estimated when the treatment is not assigned randomly, and particularly when there is correlation between the treatment and potential outcomes. The question the ATT answers is the following: conditional on a user being fired (hired), what is the impact of this action on the psychological well-being of the user? While the ATE seems more informative, the ATT is more relevant in our setting given that the treatment is never administered randomly ([21], [22]).

A problem that occurs when estimating the ATT is that endogeneity can arise if the treatment is correlated with any time-varying unobserved factor that also affects the outcome variable (in our case the measures related to the users psychological well-being). In order to alleviate such concern we employ a difference-in-differences (DD) empirical strategy. Our DD strategy identifies the effect of employment or unemployment by comparing differences in the psychological variables for people affected by the shock before and after the shock occurs, with respect to a baseline of differences in psychological variables of unaffected people over the same period of time. The specification takes the following form:

$$\begin{aligned} Measure_{it} &= \beta_0 + \beta_1 Treated_{it} + \beta_2 After_{it} \\ &+ \beta_3 Treated_{it} \times After_{it} + X_{it} + u_{it} + \tau_t + \epsilon_{it} \end{aligned} \quad (1)$$

The dependent variable is the monthly value (the average fraction of words belonging to a certain category) of one of the psychological measures defined in Table 2. The variable $After_{it}$ is an indicator of whether the observations are in the post period (after the users lost or gained a job). $Treated_{it}$ is a dummy variable indicating the users that lost or gained a job. The coefficient of interest is $Treated_{it} \times After_{it}$, which measures the difference in psychological variables before and after the users lost or gained a job, against a baseline of changes in the same variables for a set of random users over the same period time. Since our dependent variable is

a fraction, β_1 can be interpreted as the $\beta_1 * 100$ percentage points change in that variable after the user lost or gained a job. Our model includes user fixed effect u_{it} to control for time invariant users' unobservable characteristics, and calendar-month fixed effects τ_t to control for transient shocks to the psychological measures that are common to all the users. Finally, X_{it} is a vector of controls in which we include, as is common in DD analyses, treated specific linear time-trends as an additional safeguard against a possible difference in trends between treated and control units. We estimate the model in Equation 1 using OLS. To account for serial correlation in our dependent variable, we cluster errors at the user level ([16], [6]).

We present our results in Table 3. In column 1 we report the results for the employment shock, and in column 3 those for the unemployment shock. For the employment shock, the estimated coefficients for the interaction term $Treated_{it} \times After_{it}$ are statistically significant for all measures but causation, discrepancy, and sadness. For the unemployment shock, the estimated coefficients are statistically significant for all measures but anger, anxiety, friend, home, and money. Further, the coefficients make intuitive sense, reinforcing our claim that what we are measuring is indeed the effect of gaining or losing a job. We find that anger and anxiety decrease for employment. Sadness increases for unemployment, but it is not statistically significant for employment. Tentativeness decreases for employment, and it increases for unemployment. Work related Twitter posts increase for both employment and unemployment. Friends and ingroup decrease for employment.

4.2.1 Reducing the time window

To further limit the influence of unobserved factors that could affect users' psychological reaction to an employment shock, we borrow an idea from regression discontinuity designs, and we limit our estimation sample to three months before and three months year after the treatment. We estimate Equation 5 using this sub-sample of our data and report the results in column 2 (for employment) and in column 4 (for unemployment). The estimates are generally consistent with those reported in columns 1 and 3. Many coefficient are larger suggesting that the effect of employment and unemployment dissipate over time (as one would expect).

4.2.2 Falsification check

The key assumption for any DD strategy is that the outcome in the treatment and control groups would follow the same time trend in the absence of the treatment. While this assumption is untestable, we can assess the robustness of our results by performing a falsification test. We implement the test as follows: we limit our data to the pre-treatment period, i.e., before the users lost or gained a new job, and then for every user, we move the treatment a few months earlier creating in this way a placebo treatment. The assumption behind this test is that if pre-treatment trends are indeed parallel between the treated group (employed or unemployed users) and the control group, the coefficient of the placebo treatment should not be statistically significant. We estimate the following model:

$$\begin{aligned} Measure_{it} &= \beta_0 + \beta_1 Treated_{it} + \beta_2 Placebo_{it} \\ &+ \beta_3 Treated_{it} \times Placebo_{it} + X_{it} + u_{it} + \tau_t + \epsilon_{it} \end{aligned} \quad (2)$$

where $Placebo_{it}$ is the placebo treatment as defined and the remaining variables are as per Equation 1. We report the result of the robustness test using a placebo treatment that starts six months prior to the true treatment in column 5 and 6 of Table 3. We observe that for most of the variables and for both shocks, the coefficients of interest, $Treated_{it} \times Placebo_{it}$, are not statistically significant, reinforcing the credibility of our results.

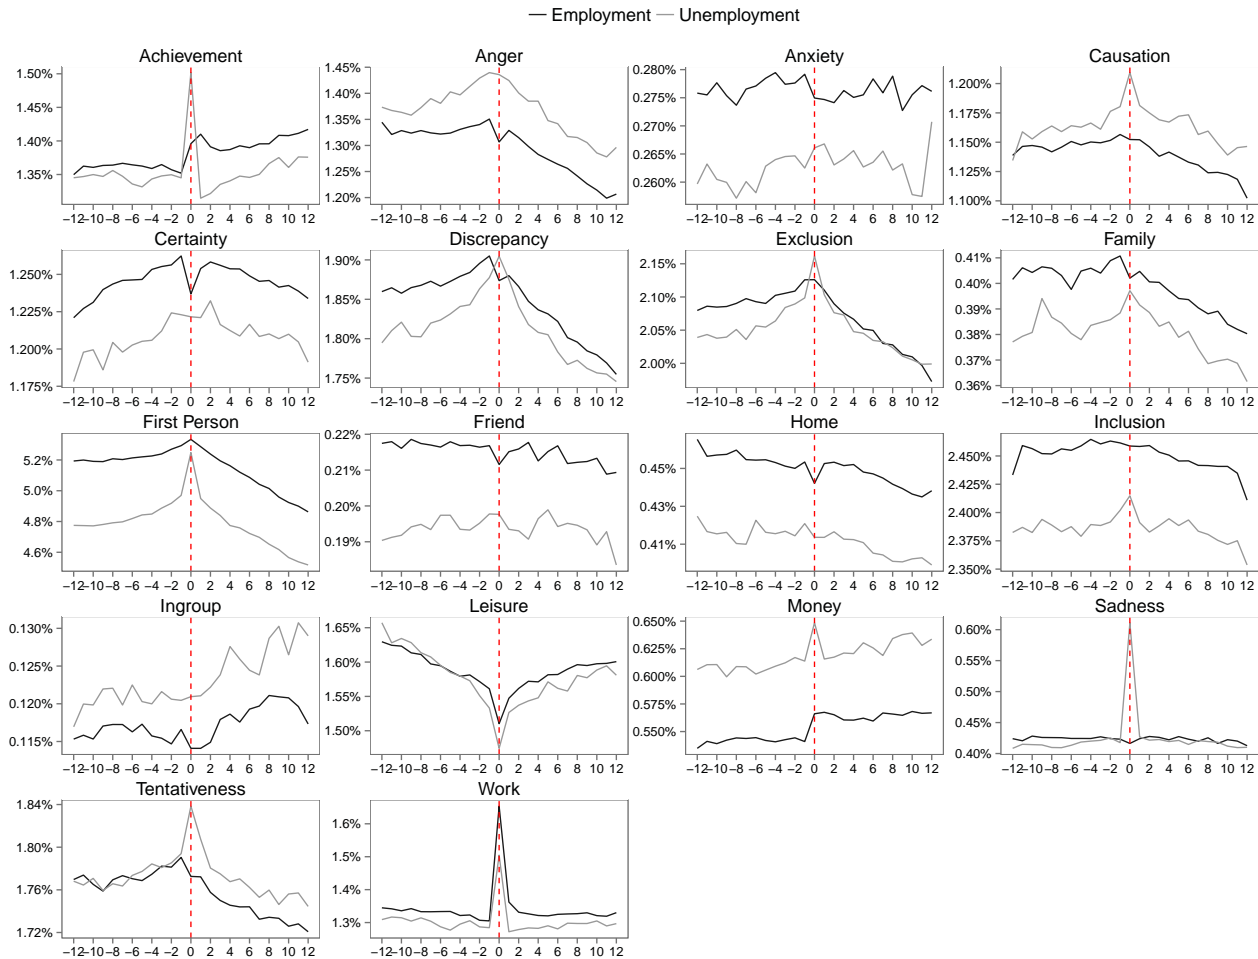


Figure 1: Comparing users' psychological variables before and after the economic shock.

4.2.3 Magnitude of the coefficients

For a sense of the magnitude of the effects observed, we can align our estimates with the mean of the psychological variables reported in Table 2. For example the mean of anxiety for job gainers is 0.0028 and the effect we observe -0.0000473 (column 1 of Table 1). Then, the decrease in anxiety with respect to the mean corresponds to about 1.7% ($(-0.0000473/0.0028) * 100$). This effect reaches 2.3% during the first three months. Similarly, achievement for job losers increases by about 2.4% ($(0.000321/0.0135) * 100$) with respect to its mean. This effect grows up to 9% when we look at the 3 months period. Sadness, increases for job losers by about 14% (39% if we look at the three months period).

5. IMPACT OF PSYCHOLOGICAL STATE ON FUTURE UNEMPLOYMENT

The previous section established the impact of job loss and gain on the psychological state of those affected. Thus a macro-scale phenomenon like unemployment exhibits micro-scale effects in the form of changes to the psychological state of individuals. We turn now to the reverse: how these micro-scale psychological states can predict the national unemployment rate.




Figure 2: Users distributions for employment, unemployment and random user.

5.1 Data

We limit our data sample to tweets posted after the year 2011. This was necessary given the small number of tweets in the earlier years of twitter, with only hundreds of observations in 2010 and 2011, compared to the many thousands of observations in later years (see Figure 2). Further, within each user we consider only data from the time of the (un)employment shock onward. This isolates the effect on future unemployment of the psychological effects caused by the initial shock. Therefore, in our predictions we use the

Table 3: Difference in differences: the effect of economic shocks on users' psychological variables

| | (1) Emp. | (2) Emp. \pm 3 months | (3) Unemp. | (4) Unemp. \pm 3 months | (5) Emp. Placebo | (6) Unemp. Placebo |
|---------------|--------------------------|----------------------------|-------------------------|------------------------------|-------------------------|-------------------------|
| Achievement | 0.000347*** (14.68) | 0.000422*** (10.45) | 0.000321*** (7.39) | 0.00122*** (18.55) | -0.0000197 (-0.50) | -0.000164** (-2.08) |
| Anger | -0.000152*** (-6.75) | -0.000224*** (-5.48) | 0.0000265 (0.59) | 0.0000866 (1.49) | 0.00000389 (0.01) | -0.000110 (-1.52) |
| Anxiety | -0.0000473*** (-5.30) | -0.0000648*** (-4.37) | 0.0000243 (1.60) | 0.0000476** (2.18) | -0.0000163 (-1.04) | -0.00000300 (-0.11) |
| Causation | 0.0000222 (1.28) | -0.00000300 (-0.11) | 0.000134*** (4.23) | 0.000257*** (5.48) | -0.000000822 (-0.03) | -0.0000463 (-0.83) |
| Certainty | -0.0000744*** (-4.06) | -0.0000958*** (-2.86) | 0.0000823** (2.50) | 0.0000692 (1.48) | -0.0000584* (-1.81) | -0.0000422 (-0.66) |
| Discrepancy | -0.0000314 (-1.44) | -0.000116*** (-3.42) | 0.000192*** (4.79) | 0.000502*** (8.74) | -0.000110*** (-2.98) | -0.0000548 (-0.80) |
| Exclusion | 0.000116*** (5.08) | 0.000176*** (4.94) | 0.000350*** (8.25) | 0.000748*** (12.49) | -0.000119*** (-3.20) | -0.0000327 (-0.46) |
| Family | -0.0000100 (-0.84) | -0.0000539*** (-3.18) | 0.0000877*** (4.25) | 0.000109*** (3.62) | -0.0000352* (-1.88) | -0.0000522 (-1.44) |
| First Person | 0.000768*** (20.31) | 0.000891*** (15.07) | 0.00123*** (17.40) | 0.00275*** (28.03) | -0.0000960 (-1.59) | -0.0000146 (-1.32) |
| Friend | -0.0000239*** (-2.69) | -0.0000486*** (-4.07) | -0.00000569 (-0.38) | 0.0000107 (0.48) | -0.00000170 (-0.12) | 0.0000124 (0.50) |
| Home | 0.0000380*** (3.05) | -0.0000218 (-1.27) | -0.00000990 (-0.44) | -0.0000365 (-0.98) | -0.000000786 (-0.04) | 0.0000990*** (2.68) |
| Inclusion | 0.0000509** (2.03) | 0.0000459 (1.21) | 0.000133*** (2.91) | 0.000161** (2.45) | 0.0000178 (0.42) | -0.0000705 (-0.89) |
| Ingroup | -0.0000135** (-2.06) | -0.0000263*** (-3.04) | 0.00000484 (0.41) | 0.00000458 (0.28) | -0.00000302 (-0.26) | 0.00000524 (0.25) |
| Leisure | -0.000260*** (-9.91) | -0.000413*** (-10.74) | -0.000279*** (-5.82) | -0.000521*** (-8.09) | 0.0000697 (1.57) | 0.0000589 (0.72) |
| Money | 0.000150*** (9.65) | 0.000198*** (8.87) | 0.0000467 (1.51) | 0.000213*** (4.79) | 0.0000146 (0.58) | -0.0000625 (-1.27) |
| Sadness | -0.0000164 (-1.31) | -0.0000299 (-1.13) | 0.000629*** (22.81) | 0.00169*** (36.07) | -0.00000831 (-0.34) | 0.0000271 (0.69) |
| Tentativeness | -0.0000676*** (-3.11) | -0.0000144 (-0.41) | 0.000259*** (6.45) | 0.000621*** (10.63) | -0.0000701* (-1.93) | 0.0000411 (0.60) |
| Work | 0.00136*** (49.17) | 0.00243*** (56.67) | 0.000618*** (12.97) | 0.00191*** (27.61) | -0.0000313 (-0.78) | -0.000231*** (-2.81) |
| N | 5020793 | 3346746 | 3080526 | 2610938 | 3653171 | 2689547 |

Note: All models include user and time fixed effects, and treatment specific linear time trends. Cluster robust t-statistics (at the user level) are shown in parentheses.

Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

users' tweet history starting from the month just after the economic shock, and up to one year after such date.

5.2 Variable selection

We used stepwise regression for variable selection. Stepwise regression is a semi-automated process of building a statistical model by successively adding or removing variables based only on the t-statistics of their estimated coefficients. Such an approach is especially useful for going through a large number of potential independent variables in order to fine-tune a model by excluding those variables that do not add much explanatory power. We run the stepwise regression on the model described in Equation 3, starting with the full model, i.e., the model that includes all the measures as per Table 2. The process then proceeds backward, removing one variable at a time, depending on the parameters specified by the user. In our case we set the significance level for removal from the model to be $p - value > 0.05$. We repeated the stepwise regression for the three dataset samples: employed, unemployed and random users. The output of the stepwise regression is reported in Table 4. In

column 1 we show the results for the unemployed sample. The process selected eight variables: achievement, family, home, inclusion, sadness, and tentativeness. In column 2 we report the output for the employed sample. In this case the generated model includes six variables: achievement, family, home, inclusion, tentativeness, and work. Finally, in column 3 we present the model for the random sample which includes 11 variables: achievement, anxiety, certainty, family, first person pronouns, inclusion, ingroup, leisure, money, sadness, tentativeness, and work.

5.3 Model

We use the selected features to build a model to predict country level monthly unemployment trends. First we create an aggregate version of our data, where for every sample (employment, unemployment and random) and every measure (see Table 2) we create a monthly average over all the users. After aggregating the metrics at the monthly level, we assess their power in predicting the U.S. unemployment rate by regressing monthly unemployment rates on

Table 4: Stepwise regression

| | (1) Unemployment | (2) Employment | (3) Random |
|--------------------|-------------------------|-------------------------|------------------------|
| Achievement | 600.835*** (8.16) | 350.611** (2.35) | 348.951*** (3.92) |
| Anxiety | | | -1169.309** (-2.72) |
| Certain | | -1038.535*** (-5.50) | -586.785*** (-3.58) |
| Family | 1086.756*** (5.71) | 846.562*** (6.42) | 364.399** (2.18) |
| First person pron. | | | 264.897*** (9.76) |
| Home | 1386.581*** (6.96) | 1250.563*** (3.51) | |
| Inclusion | -1185.369*** (-5.28) | -1096.914*** (-5.62) | -306.114*** (-2.89) |
| Ingroupp | | | 1312.405*** (2.92) |
| Leisure | | | 379.556*** (5.43) |
| Money | | | -683.149*** (-3.00) |
| Sadness | -950.447** (-2.60) | | -593.855** (-2.65) |
| Tentativeness | 1037.967*** (8.25) | 1224.703*** (9.13) | 644.435*** (3.97) |
| Work | | -173.750*** (-3.94) | |
| N | 39 | 39 | 39 |
| R ² | 0.91 | 0.94 | 0.95 |

Note: All the models contains observations form January 2012 up to March 2015. Model 1 and 2 include the observations after the users are exposed to either job loss or job gain. Cluster robust t-statistics are shown in parentheses.

Significance levels: * p<0.1, ** p<0.05, *** p<0.01.

the selected psychological variables from the previous months. The model we estimate is as follows:

$$Urate_t = \beta_0 + \beta_1 Measures_{t-1} + \varepsilon_t \quad (3)$$

Where $Urate_t$ is the official U.S. unemployment rate at month t and $Measures_{t-1}$ is a vector containing a subset of the psychological variables defined in Table 2, and lagged by one month. To perform the prediction for the month $t+1$ we estimate Equation 3 using the set of monthly observations from the period $[t-24, t-1]$, that is, we use the previous 23 month of our observed psychological variables. Then, we feed the estimated model with the measures at time t to predict the unemployment rate at time $t+1$. We then shift our 23 months window and repeat the process. Our earliest prediction is for the unemployment rate of January 2014, while our latest for April 2015. Therefore, the dataset of monthly psychological variables used for the predictions spans a period of 39 months, from January 2012 to March 2015. Finally, to deal with heteroscedasticity, we estimate robust standard errors ([23], [39]).

5.4 Comparison models

To evaluate the ability of our predictive model, we compare it with a variety of alternative models.

First we compare it to an autoregressive model:

$$Urate_t = \beta_0 + \beta_1 Urate_{t-1} + \varepsilon_t \quad (4)$$

Where the dependent variable is the U.S. unemployment rate, and the independent variable $Urate_{t-1}$ is the one month lagged U.S. unemployment rate. Given the strong serial correlation between consecutive unemployment rates, this model can be considered the baseline to beat.

Second, we compare our approach to a model similar to the one used in [2]. In [2] the authors predict labor market conditions by looking at mentions of certain key terms on Twitter and creating aggregate counts of those candidate tweets at a monthly level. These counts are then used to predict U.S. unemployment rates. In a similar fashion we use the counts of our initial set of (un)employment tweets (containing the text “got a new job”, “started a new job” for the employment sample, and “I got laid off” or “I lost my job” for the unemployment sample). We use these counts to predict the unemployment rate using Equation 3, but substituting the variable $Measures_{t-1}$ for $Counts_{t-1}$, i.e., the variable representing the raw tweet counts.

Third, for each of our three user datasets (unemployment, employment, random), we predict the unemployment rate using their respective sets of psychological variables as determined by the output of the stepwise regressions reported in Table 4. This last comparison is particularly important because it will allow us to understand how much improvement (in the predictions) one can get by selecting a certain type of users (in our case job losers and job gainers) rather than a random sample of the Twitter users.

5.5 Results

We report the prediction root mean square error (RMSE) of our models in Table 5. In the first row, we report the results obtained with the autoregressive model (our baseline). The RMSE obtained is 0.40. Next we report the results using the (un)employment tweet counts (rows 2-5). Rows 2 and 3 show RMSEs for these models based on the data starting in 2012. The errors are 0.86 for the employment dataset and 0.83 for the unemployment dataset. As noted above, because the tweet counts are very low in early years (see Figure 2), we also report (in rows 4 and 5) the prediction errors using data starting in January 2013. As expected, the RMSE for employed improves by about 25%, from 0.86 to 0.64, while the results for the unemployed sample improve by about 49%, from 0.83 to 0.42, getting close to the baseline of 0.4 obtained with the autoregressive model. In the last three rows we report the results of the models that use the psychological variables extracted from the Twitter timelines of our three population samples. In column 1 we use the subset of variables obtained from the output of the stepwise regression using the job losers (column 1 of Table 4); in column 2 we use the subset obtained using the job gainers, and in column 3 the output obtained using the random users. We observe that the RMSEs reported for these models are lower than both the autoregressive and raw (un)employment count models. The worst of these models (unemployment dataset and specification 2) outperforms the autoregressive model by 7.5%. Further, the employment sample outperforms the other samples in all three specifications, with the lowest RMSE being 0.21 in specification 2 (a 47.5% improvement over the autoregressive model and 50% over the best raw counts model). In Figure 3 we show the official unemployment rate and our best prediction.

A final observation is in regards to the results obtained with the random user sample. Contrary to our expectation, we observed fairly low prediction errors from models using this sample. They always outperform the autoregressive model, and they are comparable to those obtained with the unemployment sample. This suggests that while data from a targeted selection of users (e.g. job gainers) generally leads to better performance, the average population

sentiment has enough signal to produce worthwhile results. Note that this has also been discovered in [38], where call detail record (CDR) covering approximately 10M randomly selected subscribers in an European country have been used to augment the prediction of general levels of unemployment.

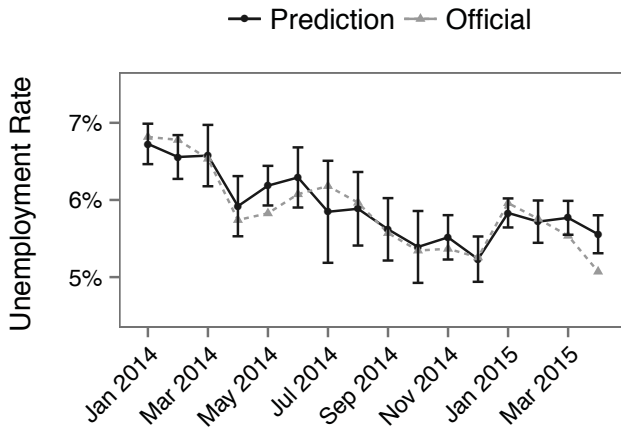


Figure 3: Official unemployment rate and its prediction with 95% confidence intervals obtained with the employed dataset and specification 2.

| | Spec. 1 | Spec. 2 | Spec. 3 |
|---------------------------|---------|---------|---------|
| Autoregressive | 0.40 | 0.40 | 0.40 |
| <i>Raw counts</i> | | | |
| Employed | 0.86 | 0.86 | 0.86 |
| Unemployed | 0.83 | 0.83 | 0.83 |
| Employed (year > 2012) | 0.64 | 0.64 | 0.64 |
| Unemployed (year > 2012) | 0.42 | 0.42 | 0.42 |
| <i>Behavioral</i> | | | |
| Unemployment | 0.33 | 0.37 | 0.35 |
| Employment | 0.27 | 0.21 | 0.29 |
| Random | 0.33 | 0.36 | 0.36 |

Table 5: Predictions error (RMSE) using different combinations of datasets and metrics: 1) Specification 1: psychological variables obtained by running a stepwise regression on the the measures of users that lost a job 2) Specification 2: psychological variables obtained by running a stepwise regression on the the measures of users that gained a new job 3) Specification 3: psychological variables obtained by running a stepwise regression on the the measures of the random users

6. DISCUSSION

6.1 Impact on and of psychological states

The primary goal of this work is to show that incorporating psychological state changes in individuals resulting from a job loss or gain event can lead to better prediction of future macro-scale unemployment. Although it is intuitive that psychological state variables like anxiety would impact an unemployed person’s ability to obtain a job, quantifying this relationship has been difficult. This is, at least in part, because the psychological state variables are at individual scale, while unemployment is at population scale. With our data, we obtain not only a large sample of people, but also a very

detailed, long-running record of their psychological state. Models using all three of our population samples outperformed baseline models for predicting future unemployment rates. This shows that these rich psychological state data streams carry information about future unemployment over and above recent histories of the number of people unemployed. Thus we hope these results help demonstrate that we can now measure a “micro-foundation” of something as macroeconomic as the unemployment rate.

This interplay between individual psychology and population scale unemployment starts with a first step, the initial impact of job loss or gain. This has been documented but not at this scale. Our difference in differences results in particular (Table 3) help us understand how people react to unemployment. Some results are expected, like increases in tentativeness, which is logical given the sudden removal of a job, which provides consistency in daily routine as well as financial stability. Others such as the decrease in focus on leisure after losing a job are a bit paradoxical, but sensible in retrospect. In this case, people who have lost a job may not feel they can afford to focus on leisure when they lack the resources to do so that a job provides. Generally these findings align with existing literature (e.g., [30]) and are able to be collected at a much greater study population scale and finer temporal granularity.

6.2 Population effects and the random sample

As mentioned in Section 5.5, the data from randomly selected users shows some predictive power for unemployment. We attribute this effect to the *level* of employment being an economy wide phenomenon that affects the general population. As a quick exploration, though certainly this warrants much deeper analysis, we relate the population average time series of our tentativeness signal from the random sample of users to a statistic called “Jobs Hard to Get”, measured by the Conference Board in their popular consumer confidence survey. We selected tentativeness based on the notion that if the prospect of finding employment is uncertain, a general feeling of tentativeness may permeate other areas of life. We find that the correlation of tentativeness to “Jobs Hard to Get” is 0.23 on a level basis and 0.19 for the first differences, suggesting that even our randomly selected users exhibit psychological effects related to the macroeconomy.

We also note that despite the bias of selecting from the Twitter user population, who are not a truly random sample of the U.S. population (e.g., Twitter users are known to skew toward urban geographies [20]), at 250,000 our set of randomly selected people is large and thus has a better chance of capturing large scale population effects. Again, deeper exploration of the effect of the economy on the population “psyche” is an area for future work. For instance, these psychological impacts of the economy may have secondary, mediating effect on related domains of life such as engagement in risky behavior or the divorce rate.

6.3 Better measurement for monetary policy

An implication of this paper is that the Federal Reserve may benefit from taking into account these psychological effects when setting monetary policy. One simple way to describe U.S. monetary policy is to treat it as a strict output gap or unemployment versus inflation trade-off [41], commonly called the Taylor rule, specified in Equations 5 and 6:

$$R(t) = 2 + i(t) + 0.5 * [I(t) - 2\%] + 0.5 * Y(t) \quad (5)$$

$$Y(t) = 2.3 * [5.6\% - U(t)] \quad (6)$$

Here $R(t)$ is the federal funds rate, $I(t)$ is the percent change in the headline Personal Consumption Expenditure (PCE) price index

from four quarters earlier, and $Y(t)$ is the output gap. The output gap is computed using Okun's law specified in Equation 5.

We recognize that the relationships between the specified macroeconomic variables is complex and that the Federal Reserve does not only rely on this simple rule, but it serves as a useful benchmark. In this framework, regardless of the starting point, everything else being equal, for a 2% decrease in unemployment rate, the policy response would be to increase the interest rate by roughly 2.3%. However, the effect of unemployment changing from say 12% to 10% may be different than unemployment changing from 5% to 3% because of the psychological effects described here. That is, given the impact of individual psychology on future unemployment that would affect a larger section of the workforce when it is 12%, as well as the higher overall population effect as suggested by the random sample results, optimal monetary policy responses may be quite different than when the unemployment rate is 5% and affects a smaller section of the workforce. This could happen if there were a nonlinearity or a "tipping point" in the psychological impact of unemployment such that rates above a certain percent have far greater consequences for the economy going forward than do lower rates.

6.4 Limitations

Generally, a lexicon-driven approach for emotion detection has limitations. First, our methodology utilizes written words that may not truly reflect the psychological state of the individual if, for example, the person is intentionally presenting themselves to a public audience. Second, the approach does not take into account negation that could be used in conjunction affective words (e.g., "not happy"). Finally, the data stream for any given person only contains snapshots of their lives they chose to post to Twitter. In our context, we argue that while these limitations may add noise to the data, they do not invalidate the findings because we consider posts of a particular person over a long time period (two years), and given the large numbers of tweets, we thus observe reasonably accurate psychological reflections of the users. We again note that a similar approach has been used a number of prior works (e.g., [15]).

In terms of implementation of our model, while we were fortunate to have an agreement with Twitter, we acknowledge that this is not the general case. Further there are privacy concerns that must be addressed when using sensitive information such as social media data. With the potential to use such data at very fine granularity, both in space and in time, privacy and security problems can arise. Anonymization techniques and data aggregation can be used to protect the users' privacy, but many researchers have shown that re-identification attacks are a potential threat ([31], [37]). We stress that we used only data from public Twitter timelines and that we always aggregated over large numbers of people and thus never identified or used data or content from any individual user.

Finally, the data collection is conditioned on who accesses and uses Twitter, introducing potential biases due to self-selection or sorting. If policy decisions are based solely on data derived from social media, the segment of the population that cannot access these applications may be underserved.

6.5 Advantages of using social media over traditional techniques

Despite these limitations, the use of social media data to predict economic effects and metrics have many advantages over traditional techniques. First the cost of collecting such data can be lower than the cost of running population surveys. Second, social media data can be collected at much finer granularity, both spatial (e.g., at ZIP code or neighborhood level) and in time (e.g., daily

instead of monthly). Third, the ability to predict economic measures weeks or months faster than traditional methods is extremely valuable for policy and decision makers both in private and public institutions. Finally, the fact that we can observe and measure the effect of economic shocks at the individual level and at an unprecedented scale, can enable economists to develop new empirical approaches and statistical methods to study the macroeconomy and the microeconomy and how they are connected.

7. CONCLUSIONS

In this paper we showed how large scale social media data can help predict important economic outcomes such as the unemployment rate. We started by showing the connection between users' psychological well-being and the economic shock of job loss or job gain. We then demonstrated that social media can capture and track changes in these variables over time. Finally, we leveraged the relationship between changes in users' psychological state and macroeconomic shocks to propose a model capable of predicting levels of U.S. unemployment better than baseline models.

Broadly our hope is that web-scale data, such as from social media, will help improve estimates of the current state of the economy. Given that these data are available in real time and at a more granular level, one improvement is faster and finer estimates. The quantity and richness of these data can change the way economists analyze and study the macro economy, and the statistical tools they employ. We do not argue that web and social media data should replace economic theories, but that they can augment standard methods, and we look forward to seeing how the integration of web data and economic analysis will be realized.

8. REFERENCES

- [1] G. A. Akerlof and J. L. Yellen. The fair wage-effort hypothesis and unemployment. *The Quarterly Journal of Economics*, pages 255–283, 1990.
- [2] D. Antenucci, M. Cafarella, M. Levenstein, C. Ré, and M. D. Shapiro. Using social media to measure labor market flows. Technical report, National Bureau of Economic Research, 2014.
- [3] B. E. Ashforth. Role transitions in organizational life: An identity-based perspective. mahwah, nj: Lawrence erlbaum associates. 2001.
- [4] R. J. Barro and J.-W. Lee. Losers and winners in economic growth. Technical report, National Bureau of Economic Research, 1993.
- [5] D. Belle and H. Bullock. The psychological consequences of unemployment. *SPSSI-Society for the Psychological Study of Social Issues*, 2011.
- [6] M. Bertrand, E. Dufflo, and S. Mullainathan. How much should we trust differences-in-differences estimates? Technical report, National Bureau of Economic Research, 2002.
- [7] N. Branden. The psychology of self-esteem: a new concept of man's psychology. 1969.
- [8] J. W. Brehm. A theory of psychological reactance. 1966.
- [9] C. C. Brown, C. Gilroy, and A. I. Kohen. The effect of the minimum wage on employment and unemployment: a survey, 1982.
- [10] M. C. Burda, D. S. Hamermesh, and J. Stewart. Cyclical variation in labor hours and productivity using the atus. Technical report, National Bureau of Economic Research, 2012.

- [11] D. Card. Origins of the unemployment rate: the lasting legacy of measurement without theory. *The American Economic Review*, 101(3):552–557, 2011.
- [12] C. Colin and L. George. *Behavioral economics: Past, present, future*. Princeton: Princeton University Press, 2004.
- [13] W. Darity and A. H. Goldsmith. Social psychology, unemployment and macroeconomics. *The Journal of Economic Perspectives*, pages 121–140, 1996.
- [14] W. Darity Jr and A. H. Goldsmith. Mr. Keynes, the new Keynesians, and the concept of full employment. In *Post-Keynesian Economic Theory*, pages 73–93. Springer, 1995.
- [15] M. De Choudhury, M. Gamon, S. Counts, and E. Horvitz. Predicting depression via social media. In *ICWSM*, 2013.
- [16] S. G. Donald and K. Lang. Inference with difference-in-differences and other panel data. *The review of Economics and Statistics*, 89(2):221–233, 2007.
- [17] P. Eisenberg and P. F. Lazarsfeld. The psychological effects of unemployment. *Psychological bulletin*, 35(6):358, 1938.
- [18] N. T. Feather. *The psychological impact of unemployment*. Springer Science & Business Media, 2012.
- [19] J. Hausman and E. Leibtag. Cpi bias from supercenters: does the bls know that wal-mart exists? In *Price Index Concepts and Measurement*, pages 203–231. University of Chicago Press, 2009.
- [20] B. Hecht and M. Stephens. A tale of cities: Urban biases in volunteered geographic information. In *Proceedings of the AAAI conference on weblogs and social media*, 2014.
- [21] J. J. Heckman, H. Ichimura, and P. E. Todd. Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme. *The review of economic studies*, 64(4):605–654, 1997.
- [22] J. J. Heckman and R. Robb. Alternative methods for evaluating the impact of interventions: An overview. *Journal of econometrics*, 30(1):239–267, 1985.
- [23] P. J. Huber. The behavior of maximum likelihood estimates under nonstandard conditions. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, volume 1, pages 221–233, 1967.
- [24] M. Jahoda. Work, employment, and unemployment: Values, theories, and approaches in social research. *American psychologist*, 36(2):184, 1981.
- [25] M. Jahoda. Economic recession and mental health: Some conceptual issues. *Journal of Social Issues*, 44(4):13–23, 1988.
- [26] A. Krueger, A. Mas, and X. Niu. The evolution of rotation group bias: Will the real unemployment rate please stand up? Technical report, National Bureau of Economic Research, 2014.
- [27] M. W. Linn, R. Sandifer, and S. Stein. Effects of unemployment on mental and physical health. *American Journal of Public Health*, 75(5):502–506, 1985.
- [28] A. Llorente, M. Cebrian, E. Moro, et al. Social media fingerprints of unemployment. *arXiv preprint arXiv:1411.3140*, 2014.
- [29] R. D. Marcuss and R. E. Kane. Us national income and product statistics. *Survey of Current Business*, 87:2–32, 2007.
- [30] F. M. McKee-Ryan and J. Harvey. “i have a job, but...”: A review of underemployment. *Journal of Management*, 37(4):962–996, 2011.
- [31] A. Narayanan and V. Shmatikov. De-anonymizing social networks. In *Security and Privacy, 2009 30th IEEE Symposium on*, pages 173–187. IEEE, 2009.
- [32] G. E. O’Brien. *Psychology of work and unemployment*. John Wiley & Sons, 1986.
- [33] K. I. Paul and K. Moser. Unemployment impairs mental health: Meta-analyses. *Journal of Vocational behavior*, 74(3):264–282, 2009.
- [34] J. W. Pennebaker, C. K. Chung, M. Ireland, A. Gonzales, and R. J. Booth. The development and psychometric properties of liwc2007, 2007.
- [35] R. H. Price, D. S. Friedland, and A. D. Vinokur. Job loss: Hard times and eroded identity. *Perspectives on loss: A sourcebook*, pages 303–316, 1998.
- [36] T. Sakaki, M. Okazaki, and Y. Matsuo. Earthquake shakes twitter users: real-time event detection by social sensors. In *Proceedings of the 19th international conference on World wide web*, pages 851–860. ACM, 2010.
- [37] L. Sweeney. Simple demographics often identify people uniquely. *Health (San Francisco)*, 671:1–34, 2000.
- [38] J. L. Toole, Y.-R. Lin, E. Muehlegger, D. Shoag, M. C. González, and D. Lazer. Tracking employment shocks using mobile phone data. *Journal of The Royal Society Interface*, 12(107):20150185, 2015.
- [39] H. White. Maximum likelihood estimation of misspecified models. *Econometrica: Journal of the Econometric Society*, pages 1–25, 1982.
- [40] A. H. Winefield, H. R. Winefield, M. Tiggemann, and R. D. Goldney. A longitudinal study of the psychological effects of unemployment and unsatisfactory employment on young adults. *Journal of Applied Psychology*, 76(3):424, 1991.
- [41] J. L. Yellen. The economic outlook and monetary policy. Technical report, Board of Governors of the Federal Reserve System (U.S.), 2012.
- [42] J. L. Yellen. What the federal reserve is doing to promote a stronger job market, 2014.
- [43] A. X. Zhang and S. Counts. Modeling ideology and predicting policy change with social media: Case of same-sex marriage. In *Proceedings of the SIGCHI conference on human factors in computing systems*, 2015.