Extended Abstract: The Relationship Between the "Norm to Work" and Labor Market Outcomes

PRELIMINARY RESULTS

Jennifer Kades

Last updated: December 20, 2024

Abstract

This project explores how work norms may shape the labor market. Media narratives and existing research suggest such norms are important drivers of worker behavior and, ultimately, macroeconomic outcomes like labor supply and wages, but this idea has not been tested empirically due to the difficulty of measuring norms. I construct a novel measure of work norms using unstructured text data from Twitter, combined with a custom-built machine learning model to label the text. I then use this measure to examine the relationship between work norms and labor market outcomes such as labor force participation, employment, and wages.

1 Introduction

The past decade has seen a surge of interest among economists in the extent to which social norms may shape labor market outcomes.[1] Although much of this work is focused on gender norms, recent labor market upheavals in the wake of Covid-19 have also highlighted the potential importance of broader social norms and attitudes regarding work, with many commentators claiming that a pandemic-induced re-evaluation of the importance of work has led to more quits, higher wages, and a resurgence in worker power (e.g., Hirsch 2021; Krugman 2021; Sull, Sull, and Zweig 2022). While there is some evidence of decreased willingness to work and increased worker power in the form of unionization, this narrative has proven difficult to test empirically (Faberman, Mueller, and Şahin 2022; Kinder and Stateler 2022), in large part because social norms are so challenging to quantify. The big-picture question remains: to what extent does the "norm to work"-—i.e., the social norm that working-age adults should be employed-—influence worker behavior and, ultimately, macro-level outcomes such as employment and wages?

Existing research has established that social norms are powerful drivers of human behavior (e.g., Coleman 2018; Elster 1989; McAdams 1997), and this appears to be especially true in the labor market. Unemployment carries heavy psychological costs even after controlling for financial duress (Blanchflower and Oswald 2004; Brand 2015; Jahoda 1988). Research has also shown that these psychosocial costs are determined in part by unemployment rates or the prevailing social norms in one's community (often treated as equivalent in earlier work) (Powdthavee 2007; Stavrova, Schlösser, and Fetchenhauer 2011). Unemployed people tend to be unhappier in areas with stronger work norms, and there is some evidence suggesting that the unemployed expend more effort on job search when

1. Recent examples of papers in this area include, but are certainly not limited to: Bertrand et al. 2016; Binzel and Carvalho 2017; Breza, Kaur, and Krishnaswamy 2019; Bursztyn, Fujiwara, and Pallais 2017; Bursztyn, González, and Yanagizawa-Drott 2020; Dube, Naidu, and Reich 2022, Fortin 2015; Jayachandran 2021.

employment norms are stronger (Chadi 2014; Clark 2003; Eugster et al. 2017; Stutzer and Lalive 2004). It also seems plausible that employed people take greater pains to avoid unemployment in locations with a stronger NTW, though there is currently no empirical evidence of this. Taken together, then, the NTW has a potentially large impact on job search and labor supply behavior.

While the literature has established a clear link between the NTW and the well-being and behavior of unemployed people, no existing research has moved beyond these individual-level, direct effects to study the potential aggregate-level impacts of work norms. Additionally, prior work has been limited by the difficulty of measuring the NTW. The papers referenced above largely rely on survey data, which tends to be irregularly intervalled and geographically coarse, or else simply use measures of local or network-level unemployment as a proxy for employment norms (so-called "descriptive" norms). Creative exceptions include Stutzer and Lalive 2004 and Eugster et al. 2017, which exploit region-level vote shares from a Swiss referendum on unemployment benefits and a language border that separates cultures (but not labor markets), respectively. However, such solutions are unique and highly context-specific. This paper aims to address these gaps.

This paper offers preliminary evidence of a relationship between the social norm to work ("NTW") and labor supply, using a novel, social media-based measure of norms. The measure is based on a large sample of tweets, which are categorized based on the norms or attitudes toward work that they convey. The labeled tweets are then aggregated by time and region to create a dynamic and geographically specific measure of norms in the United States over roughly the past decade. Whereas this type of unstructured text classification would historically have required human coders, I instead leverage new machine learning technology to label the tweets, making the process of parsing large amounts of text data much quicker and cheaper.

I use the resulting measure to provide, to my knowledge, the first empirical evidence documenting an effect of work norms on aggregate labor market outcomes. Initial results are noisy, but suggest a positive relationship between the NTW and labor supply, as measured by labor force participation and employment. This aligns with theoretical predictions (based on, e.g., a reference-dependent labor supply framework). Additional work aiming to distinguish the direction of the causality is ongoing. These results have important implications for issues such as unemployment protections and take-up in the context of psycho-social costs of unemployment.

The measure and measure construction methodology are also key contributions of this paper. The NTW measure developed for this project is unique in its scope and geographic granularity, and thus may prove useful to economists studying a broad array of questions exploring the impact of norms in the labor market. More broadly, the methodology employed to create the measure may be easily generalized to other norms, providing researchers with a new tool for quantifying social norms in different contexts.

2 Measuring the Norm to Work

The norms measures used in this paper are based on data from Twitter, labeled using a custom-trained text classification model. Figure 1 presents an outline of the measure construction workflow, which is described in more detail below.

Figure 1: Outline of the measure construction workflow

Pull **20 million tweets** using relevant keywords and phrases

↓

Restrict the sample to U.S. tweets that can be identified at the state or city level based on geo-tags and user-reported data (**4.3 million tweets**)

↓

Pre-process tweet text (e.g., remove user tags, links, etc.)

↓

Pilot labels/labeling performance across OpenAI large language models, M-Turk, and Prolific to determine whether to use human- or AI-generated labels

↓

Use OpenAI's GPT-4 API to label a training dataset of **100,000 tweets**

↓

Train a machine learning model to label the remainder of the tweets in the sample

↓

Aggregate labeled tweets by time and geography to construct the final measures

Tweet sample

The NTW measures in this paper are based on a sample of roughly 20 million tweets collected through Twitter's Academic API.[2] The sample was selected using keyword searches intended to filter for tweets relating to the norm to work and attitudes towards employment and unemployment[3]. The sample begins in 2011—the year when Twitter reached 100 million users—and ends in November 2022, when the tweets were collected (Meyer 2019). After dropping tweets that were from outside the U.S. or missing geographic data (based on user-reported locations or geo-tags), the final sample consists of about 4.3 million tweets[4].

Tweet labeling

I employ a supervised machine learning approach to perform multi-class label classification on each tweet in the sample. As there is little concrete guidance regarding the exact language that best characterizes the NTW, I include several different (exclusive) category labels capturing different facets of work norms, as follows:

1. Anti-work
2. Norm to work
3. Pro-workers' rights

2. The Academic API, along with all other affordable tiers of Twitter API access, was discontinued in 2022 following Elon Musk's acquisition of the platform.

3. While some searches were based on single keywords (e.g., for "unemployment") many topic words (e.g., "work") were too vague to yield a reasonable proportion of relevant tweets. Therefore, I employed many multi-word searches intended to triangulate beliefs and attitudes towards work and employment (e.g., a keyphrase search for tweets containing ("work" OR "job") AND "lazy", meant to roughly capture sentiments along the lines of "people who don't work (or have a job) are lazy").

4. The Twitter API does not permit filtering based on users' self-reported locations (the primary source of geographical data), so many of the collected tweets had to be dropped from the sample ex-post.

4. Unemployment stigma
5. Not applicable

The final label captures relevance, as many of the tweets in the sample are simply off-topic. The labels were chosen to distinguish between positive versus negative framings of the social norm to work, as well as positive versus negative framings of the opposing (anti-work) attitudes. In practice, the labels appear to capture meaningfully different concepts. Upon inspection of the labels chosen by workers on MTurk or Prolific, “norm to work” and “anti-work” seem to be interpreted more broadly or abstractly, while “pro-workers’ rights” or “unemployment stigma” tend to be chosen for tweets containing more concrete and specific language around, e.g., advocating for better working conditions or criticizing unemployment benefits. See Table 1 for an example of tweets classified into each label in the training data (as discussed in the following paragraphs).

Table 1: Example Labeled Tweets

Tweet text	**Label***
thrive global: why we need to talk about elizabeth holmes’s destructive work ethic. (link) via @User	anti-work
people get burned out because they’re working so hard not chasing their dream @User	anti-work
looking forward to getting back to work at the ballpark early in the morning and being a productive designer. #earlymorningdesign #natitude	norm to work
being on time means 15 minutes early. employers will notice. followed this for 35 years	norm to work
@User i wonder how many people commenting pro wwe sentiment think it’s smart and clever when their boss rips them off.	pro-workers’ rights
happy work/life balance being addressed, but y anyone wld ever consider it a “women’s issue” is beyond me: (link) @User	pro-workers’ rights
@User what the heck is wrong with you? get a job!	unemployment stigma
yeah, people who are lazy and don’t work for their money while taking handouts from the government usually are “chill and laid back”	unemployment stigma
another hard working @User family safeguarding their cash and time by not overpaying ridiculously high list fees (link)	not applicable
framingham jobs: server: amc dine-in theatres - framingham, ma - purpose deliver superior service. (link) #jobs #framingham	not applicable

*Reflects labels generated by OpenAI’s GPT-4 model and used to train the current version of the machine learning model.

Supervised machine learning requires a pre-existing set of correctly-labeled data to train a model—typically a dataset that has been manually labeled by humans. After conducting a pilot study comparing labels for 700 tweets generated by workers on M-Turk and Prolific to labels chosen by OpenAI’s GPT-3.5 and GPT-4 models, I instead opted to use GPT-4 to label the training dataset, as the data quality appeared to be similar across these sources (see Figure 2 for detail). I used GPT-4 to label a training dataset of 100,000 tweets, with each tweet classified into one of the categories listed above. I then used this set of 100,000 labeled tweets to train a machine learning model[5] for this multi-class labeling task, which generated predicted labels for the remainder of the tweets in the sample.

5. The machine learning model used for labeling the tweets in this sample was built using the Keras and Tensorflow

The details and performance of the preliminary tweet classification model are summarized in Tables 2 and 3 below. The overall accuracy of the multi-class classification on hold-out data is $\sim$ 62%, though accuracy considering each label individually (binary classification) is higher. This level of accuracy is in line with the accuracy of text classification models used in similar contexts[6], but introduces a significant amount of noise into the final NTW measures.

This process resulted in a sample of about 4.3 million tweets, each classified into one of the categories listed above.

Figure 2: Comparison of Human and Large Language Model Labels from Pilot*

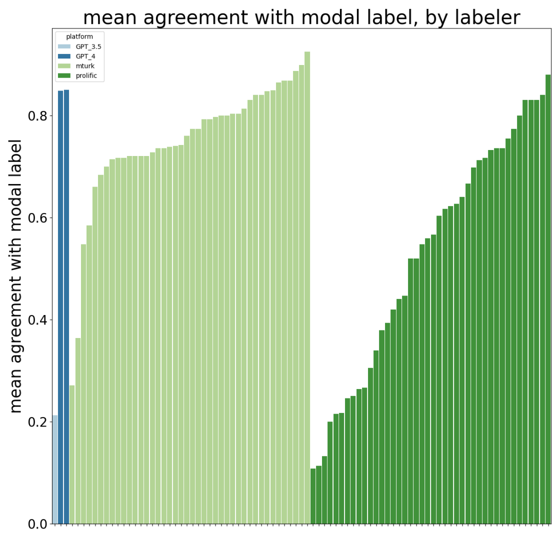

*Each tweet in the pilot sample was labeled by nine separate labelers: GPT-3.5, GPT-4 at two different "temperature" settings, and three different Prolific and M-Turk workers, respectively. The figure shows the extent to which each labeler's labels agreed with the modal label for each tweet. Distributions for M-Turk and Prolific reflect different labelers as captured by worker ID. As shown by the dark blue bars, the GPT-4 labels aligned with the modal label over 80% of the time—higher than most human labelers from either platform.

packages in Python. It is a convolutional neural network model trained to conduct multi-class text classification. Convolutional neural networks are generally considered the gold standard for classification tasks with high-dimensional, unstructured data (such as tweets), as their multi-layer and flexible structure allows for highly complex data representations (Chollet 2021). Specifically, the models used in this paper are recurrent neural network (RNN) models that include long short-term memory ("LSTM") layers. RNNs are ideal for sequential data (such as text data), as they allow the model to have "memory" of data from earlier in a sequence of elements (essentially, state-dependence).

6. E.g., text classification models achieved accuracy of about 70% in Adams-Prassl, Barbanchon, and Marcato 2022 (identifying training offers in job ad text), 74%-85% in Lassébie et al. 2021 (identifying skills in job ad text), and 54%-97.5% in Osnabrügge, Ash, and Morelli 2023 (identifying topics in parliamentary speeches).

Table 2: Model summary

Layer (type)	Output shape	Parameters
embedding 1 (Embedding)	(None, 67, 200)	10000000
lstm 1 (LSTM)	(None, 64)	67840
dense 1 (Dense)	(None, 64)	4160
dense 2 (Dense)	(None, 5)	325

Table 3: Model performance on hold-out data

Model	Accuracy	Precision	Recall
Multi-class classification (all labels simultaneously) (used for preliminary measure)	62%	71%	51%
Binary label classification: “not applicable”	78%	78%	67%
Binary label classification: “anti-work”	85%	60%	38%
Binary label classification: “norm to work”	80%	70%	59%
Binary label classification: “pro workers’ rights”	93%	63%	33%
Binary label classification: “unemployment stigma”	95%	63%	42%

Preliminary NTW measures

The labeled tweets are aggregated by time and geography to create four NTW measures (corresponding to categories 1-4 listed in the previous section). Each measure is adjusted for the volume of *relevant* tweets in a given area at time t. The measure for each label is constructed as follows:

$$NTW_{irt} = \frac{\sum\limits_{\text{tweet} \in \text{tweets}_{rt}} \text{label i}}{\sum\limits_{\text{tweet} \in \text{tweets}_{rt}} \text{relevance label}}$$

Where r denotes region (with measures constructed at both the U.S. state and metro-area level), t denotes time (year), and i denotes the label being used (ranging from 1-4). Measures were constructed for all 50 states and the top 50 largest cities by population. Recall that all labels are binary, so this is equivalent to dividing the count of positive (label=1) tweets by the count of relevant (i.e., not categorized as "Not applicable") tweets for each label.

While I examined results using alternative denominators (population and gross tweet volume), adjusting for relevant tweet volume is theoretically the most robust approach. Area population does not account for variation in the amount people tweet in different areas at different times if this is not directly proportional to population. Adjusting for gross tweet volume solves this issue, but is sensitive to changes in popular interest in work-related discussions. For example, if people become more interested in discussing unemployment policy from one year to the next, but gross tweet volume remains unchanged, the NTW could look artificially higher simply because there are more relevant tweets in the latter sample (both supporting and opposing the NTW). Indeed, the gross volume-adjusted and relevant volume-adjusted measures are highly correlated (typically >95%) within years, but produce different time trends. For the reasons stated above, this paper focuses on results using the relevant volume-adjusted measures.

Figure 3: Label shares by year

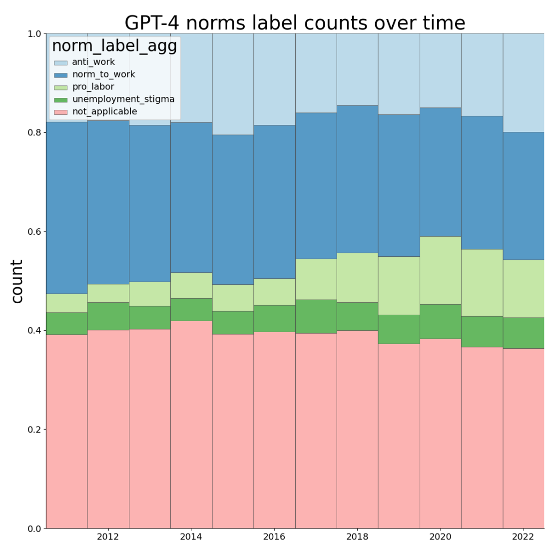

As shown in Figure 3, the share of tweets expressing a "norm to work" attitude has been slowly declining since 2011, while the share of tweets expressing "pro-workers' rights" has grown over the sample period. Interestingly, there was not a dramatic shift in label shares post-Covid, suggesting that worker dissatisfaction may have been steadily building leading up to the pandemic, with Covid serving as a catalyst for discussion and change.

Figure 4: Preliminary norm levels by state (2011-2022 average)

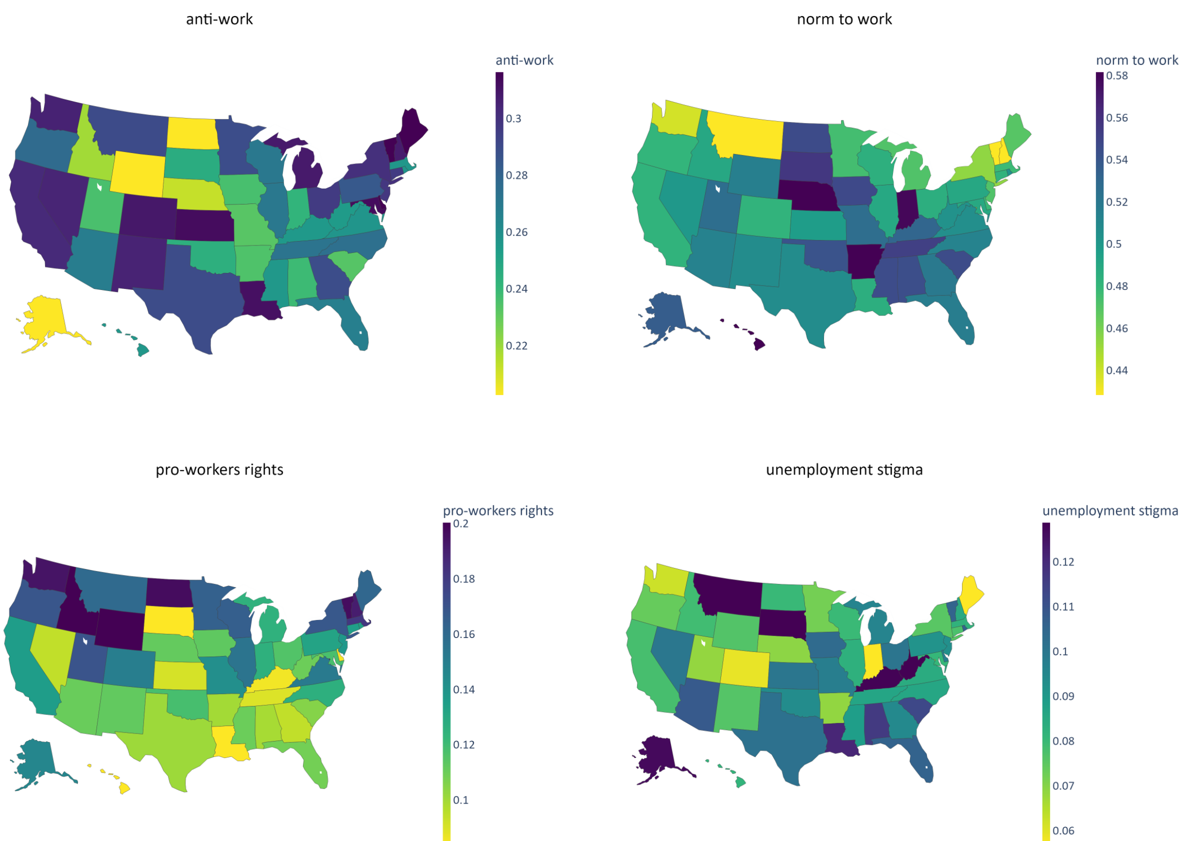

Levels of each preliminary norm measure by state are shown in Figure 4, aggregated over all of the sample years. The level of "anti-work" sentiment by state appears somewhat noisy, but is generally consistent with what one might expect based on media narratives (e.g., stronger anti-work sentiment on the coasts, weaker in the South). The level of "norm to work" sentiment is roughly inverse to the level of "anti-work" attitudes (e.g., stronger in the South and midwest), as we should expect conceptually. Note that this is not by construction, as each tweet could be sorted into one of four possible relevant categories (in other words, while the exclusivity of the labels means each label necessarily takes shares from all the others, the correlation between levels of any two particular labels could be null or even positive). Northeastern and upper midwestern states have the highest levels of "pro-workers' rights" attitudes, consistent with historical union strongholds, while "unemployment stigma" levels approximately invert this pattern.

3 Results

Using the preliminary norms measures described above, I examine the empirical relationship between work norms and labor supply (as measured by labor force participation and employment) and wages. Data for the individual-level labor market outcomes used in the analysis—labor force status, employment, and wage data—come from the Census Bureau's Annual Social and Economic Supplement (ASEC) of the Current Population Survey for years 2011-2019. The sample is restricted to civilian (non-military) adults aged 25-65.

Using the framework of reference-dependent labor supply from Farber 2008, I predict that stronger work norms (in terms of "norm to work" and "unemployment stigma") will be associated with higher labor force participation. In particular, to the extent that workers share a common understanding of the minimum number of hours they need to work to be considered employed, thus complying with the norm to work, we should see a positive relationship between work norms and binary labor force participation and employment, as well as hours up to the reference point. This also means we may observe a negative relationship between work norms and wages, operating indirectly through labor supply. Preliminary results are consistent with these predictions.

These initial results are based on the time period from 2011-2019. While the Covid-19 pandemic presents a potentially interesting shock to work norms, it also impacted the labor market in a number of other ways, both directly and indirectly, making any interpretation of post-Covid results inherently challenging. I hope to explicitly address the impact of Covid-19 on the NTW, and possible implications for labor market dynamics post-Covid, in future iterations of this project.

All analyses employ two-way fixed effects to control for unobserved heterogeneity due to year and time-invariant geographical characteristics. However, this specification is not robust to simultaneity, an obvious concern in this setting. Therefore, the current results should be interpreted with caution—in particular, it is not currently possible to say anything definitive about the direction of causality. Teasing out causality, for example through policy or labor events (e.g., large strikes) that could offer plausibly exogenous variation in labor supply or work norms, is a focus of ongoing work. The estimating equation is given by:

$$Y_{irt} = \alpha + \beta_1 norm_{rt} + \mathbf{X}_i \beta_X + \mathbf{Z}_{rt} \beta_Z + \phi_r + \tau_t + \epsilon_{irt} \tag{1}$$

where β_1 is the coefficient of interest, Y_{irt} is individual i's binary employment status [unemployment status, wage] at time t in region r, $norm_{rt}$ is the norm measure at time t in region r, $\mathbf{X}_i$ is a vector of individual-level controls (e.g., demographics, union membership), $\mathbf{Z}_{rt}$ is a vector of region-by-time controls (e.g., lagged regional unemployment, unemployment benefit generosity), and ϕ_r and τ_t are region and time fixed effects.

The observed relationships between norms and labor supply are broadly consistent with theoretical predictions, noting however that not all results are significant at conventional levels. In particular, as shown in Tables 4 and 5, there is a strong and significant negative relationship between "anti-work" attitudes and labor force participation and employment, and a strong positive relationship between NTW and these outcomes. Labor supply results for the "pro-workers' rights" and "unemployment stigma" measures are directionally consistent but imprecisely estimated. This is not surprising, as these labels were relatively rare in the training data, and ML models tend to struggle with low-incidence labels, meaning that these measures are likely noisier.

Wage results are imprecisely estimated (indistinguishable from zero) for all norms measures. This may be due to simultaneity, as theoretical frameworks predict opposing effects of norms (higher wages would presumably cause workers to view work more favorably, but stronger NTW could indirectly depress wages by increasing labor supply).

Table 4: Effect of state-level norms on Labor Force Participation

	Effect of state-level norms on LFP			
VARIABLES	(1)	(2)	(3)	(4)
anti-work (relevant vol-adj.)	-0.00574***			
	(0.00186)			
norm to work (relevant vol-adj.)		0.00277**		
		(0.00108)		
pro-workers rights (relevant vol-adj.)			0.00397*	
			(0.00225)	
unemp. stigma (relevant vol-adj.)				-0.000529
				(0.00133)
Controls				
tweet count (state)	-2.07e-07	-4.91e-07	-2.01e-07	-4.03e-07
	(4.43e-07)	(4.43e-07)	(4.29e-07)	(4.16e-07)
lagged unemp rate (state)	0.0858	0.0698	0.104	0.0899
	(0.101)	(0.102)	(0.101)	(0.100)
state population	-7.31e-09***	-6.94e-09***	-6.69e-09***	-6.57e-09***
	(1.66e-09)	(1.60e-09)	(1.56e-09)	(1.53e-09)
state UI replacement ratio	0.000949	-0.00171	-0.00728	-0.00790
	(0.0315)	(0.0317)	(0.0313)	(0.0314)
Constant	0.319***	0.319***	0.312***	0.313***
	(0.0305)	(0.0299)	(0.0294)	(0.0292)
Mean LFP	**0.782**	**0.782**	**0.782**	**0.782**
Observations	785,432	785,432	785,432	785,432
R-squared	0.113	0.113	0.113	0.113
Controls for demographics	✓	✓	✓	✓
Year and region FE	✓	✓	✓	✓

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

This table shows that there is a meaningful relationship between state-level norms and labor force participation (LFP). Columns show the relationship between individual LFP (binary) by state and year with the level of different norms in the same region and year, estimated according to equation (1). Controls include those shown above as well as individual-level demographic controls (e.g., age, gender) and state and year fixed effects. Standard errors clustered at the region-year level are shown in parentheses.

Table 5: Effect of state-level norms on Employment

Effect of state-level norm to work on employment				
	(1)	(2)	(3)	(4)
VARIABLES	anti_work	ntw	pro_labor	unemp_stigma
anti-work (relevant vol-adj.)	-0.00602***			
	(0.00174)			
norm to work (relevant vol-adj.)		0.00326***		
		(0.00108)		
pro-workers rights (relevant vol-adj.)			0.00311	
			(0.00238)	
unemp. stigma (relevant vol-adj.)				-0.00115
				(0.00128)
Controls				
tweet count (state)	1.07e-07	-2.05e-07	6.58e-08	-1.28e-07
	(5.43e-07)	(5.15e-07)	(5.59e-07)	(5.23e-07)
lagged unemp rate (state)	-0.490***	-0.510***	-0.474***	-0.489***
	(0.101)	(0.102)	(0.103)	(0.102)
state population	-7.91e-09***	-7.56e-09***	-7.26e-09***	-7.01e-09***
	(1.72e-09)	(1.66e-09)	(1.68e-09)	(1.66e-09)
state UI replacement ratio	-0.0123	-0.0143	-0.0212	-0.0212
	(0.0321)	(0.0322)	(0.0329)	(0.0327)
Constant	0.254***	0.255***	0.247***	0.247***
	(0.0337)	(0.0330)	(0.0330)	(0.0329)
Mean Employment	**0.742**	**0.742**	**0.742**	**0.742**
Observations	785,432	785,432	785,432	785,432
R-squared	0.108	0.108	0.108	0.108
Controls for demographics	✓	✓	✓	✓
Year and region FE	✓	✓	✓	✓

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

This table shows that there is a meaningful relationship between state-level norms and employment. Columns show the relationship between individual employment (binary) by state and year with the level of different norms in the same region and year, estimated according to equation (1). Controls include those shown above as well as individual-level demographic controls (e.g., age, gender) and state and year fixed effects. Standard errors clustered at the region-year level are shown in parentheses.

Table 6: Effect of state-level norms on Wages

Effect of state-level norms on wages				
VARIABLES	(1)	(2)	(3)	(4)
anti-work (relevant vol-adj.)	0.00137			
	(0.00256)			
norm to work (relevant vol-adj.)		-0.000602		
		(0.00186)		
pro-workers rights (relevant vol-adj.)			-0.00365	
			(0.00418)	
unemp. stigma (relevant vol-adj.)				0.00318
				(0.00236)
Controls				
tweet count (state)	-2.03e-06*	-1.97e-06*	-2.13e-06*	-1.76e-06
	(1.16e-06)	(1.16e-06)	(1.18e-06)	(1.17e-06)
lagged unemp rate (state)	-0.185	-0.181	-0.199	-0.174
	(0.223)	(0.224)	(0.222)	(0.224)
state population	4.02e-09	3.93e-09	3.81e-09	2.99e-09
	(4.34e-09)	(4.34e-09)	(4.28e-09)	(4.41e-09)
state UI replacement ratio	0.149**	0.149**	0.151**	0.150**
	(0.0610)	(0.0613)	(0.0614)	(0.0612)
state union coverage	-0.104	-0.104	-0.102	-0.0976
	(0.132)	(0.132)	(0.132)	(0.132)
state min. wage	0.00697***	0.00697***	0.00685***	0.00661***
	(0.00220)	(0.00220)	(0.00220)	(0.00218)
Constant	9.121***	9.121***	9.126***	9.134***
	(0.0756)	(0.0758)	(0.0749)	(0.0755)
Mean log wage	**10.66**	**10.66**	**10.66**	**10.66**
Observations	503,621	503,621	503,621	503,621
R-squared	0.378	0.378	0.378	0.378
Controls for demographics	✓	✓	✓	✓
Year and region FE	✓	✓	✓	✓

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

This table shows the relationship between state-level norms and wages, which appears to be small and imprecisely estimated. Columns show the relationship between log wages by state and year with the level of different norms in the same region and year, estimated according to equation (1). Controls include those shown above as well as individual-level demographic controls (e.g., age, gender) and state and year fixed effects. Standard errors clustered at the region-year level are shown in parentheses.

References

Adams-Prassl, Abi, Thomas Le Barbanchon, and Alberto Marcato. 2022. "On The Job Training and Labor Competition." Presented at the 5th IDSC of IZA Workshop: Matching Workers and Jobs Online - New Developments and Opportunities for Social Science and Practice, May.

Bertrand, Marianne, Patricia Cortés, Claudia Olivetti, and Jessica Pan. 2016. *Social norms, labor market opportunities, and the marriage gap for skilled women.* Technical report. National Bureau of Economic Research.

Binzel, Christine, and Jean-Paul Carvalho. 2017. "Education, social mobility and religious movements: The Islamic revival in Egypt." *The Economic Journal* 127 (607): 2553–2580.

Blanchflower, David G, and Andrew J Oswald. 2004. "Well-being over time in Britain and the USA." *Journal of public economics* 88 (7-8): 1359–1386.

Brand, Jennie E. 2015. "The far-reaching impact of job loss and unemployment." *Annual review of sociology* 41:359–375.

Breza, Emily, Supreet Kaur, and Nandita Krishnaswamy. 2019. *Scabs: The social suppression of labor supply.* National Bureau of Economic Research.

Bursztyn, Leonardo, Thomas Fujiwara, and Amanda Pallais. 2017. "'Acting wife': Marriage market incentives and labor market investments." *American Economic Review* 107 (11): 3288–3319.

Bursztyn, Leonardo, Alessandra L González, and David Yanagizawa-Drott. 2020. "Misperceived social norms: Women working outside the home in Saudi Arabia." *American economic review* 110 (10): 2997–3029.

Chadi, Adrian. 2014. "Regional unemployment and norm-induced effects on life satisfaction." *Empirical Economics* 46 (3): 1111–1141.

Chollet, Francois. 2021. *Deep learning with Python.* Simon / Schuster.

Clark, Andrew E. 2003. "Unemployment as a social norm: Psychological evidence from panel data." *Journal of labor economics* 21 (2): 323–351.

Coleman, James S. 2018. "The emergence of norms." In *Social institutions,* 35–60. Routledge.

Dube, Arindrajit, Suresh Naidu, and Adam D Reich. 2022. *Power and Dignity in the Low-Wage Labor Market: Theory and Evidence from Wal-Mart Workers.* Technical report. National Bureau of Economic Research.

Elster, Jon. 1989. "Social norms and economic theory." *Journal of economic perspectives* 3 (4): 99–117.

Eugster, Beatrix, Rafael Lalive, Andreas Steinhauer, and Josef Zweimüller. 2017. "Culture, work attitudes, and job search: Evidence from the Swiss language border." *Journal of the European Economic Association* 15 (5): 1056–1100.

Faberman, R Jason, Andreas I Mueller, and Ayşegül Şahin. 2022. "Has the willingness to work fallen during the Covid pandemic?" *Labour Economics* 79:102275.

Farber, Henry S. 2008. "Reference-dependent preferences and labor supply: The case of New York City taxi drivers." *American Economic Review* 98 (3): 1069–82.

Fortin, Nicole M. 2015. "Gender role attitudes and women's labor market participation: Opting-out, aids, and the persistent appeal of housewifery." *Annals of Economics and Statistics/Annales d'Économie et de Statistique,* nos. 117/118, 379–401.

Hirsch, Peter Buell. 2021. "The great discontent." *Journal of Business Strategy.*

Jahoda, Marie. 1988. "Economic recession and mental health: Some conceptual issues." *Journal of social Issues* 44 (4): 13–23.

Jayachandran, Seema. 2021. "Social norms as a barrier to women's employment in developing countries." *IMF Economic Review* 69 (3): 576–595.

Kinder, Molly, and Laura Stateler. 2022. "Frontline workers were excluded from companies' pandemic windfalls. No wonder so many are forming unions." *The Avenue* (May 4, 2022). https://www.brookings.edu/blog/the-avenue/2022/05/04/frontline-workers-were-excluded-from-companies-pandemic-windfalls-no-wonder-so-many-are-forming-unions/.

Krugman, Paul. 2021. "Working Out: Is the Great Resignation a Great Rethink?" *New York Times* 29.

Lassébie, Julie, Luca Marcolin, Marieke Vandeweyer, and Benjamin Vignal. 2021. "Speaking the same language: A machine learning approach to classify skills in Burning Glass Technologies data."

McAdams, Richard H. 1997. "The origin, development, and regulation of norms." *Mich. L. Rev.* 96:338.

Meyer, Jack. 2019. "History of Twitter: Jack Dorsey and The Social Media Giant." TheStreet. Accessed March 28, 2023. https://www.thestreet.com/technology/history-of-twitter-facts-what-s-happening-in-2019-14995056.

Osnabrügge, Moritz, Elliott Ash, and Massimo Morelli. 2023. "Cross-domain topic classification for political texts." *Political Analysis* 31 (1): 59–80.

Powdthavee, Nattavudh. 2007. "Are there geographical variations in the psychological cost of unemployment in South Africa?" *Social indicators research,* 629–652.

Stavrova, Olga, Thomas Schlösser, and Detlef Fetchenhauer. 2011. "Are the unemployed equally unhappy all around the world? The role of the social norms to work and welfare state provision in 28 OECD countries." *Journal of Economic Psychology* 32 (1): 159–171.

Stutzer, Alois, and Rafael Lalive. 2004. "The role of social work norms in job searching and subjective well-being." *Journal of the European Economic Association* 2 (4): 696–719.

Sull, Donald, Charles Sull, and Ben Zweig. 2022. "Toxic culture is driving the great resignation." *MIT Sloan Management Review* 63 (2): 1–9.