

WORKING PAPER · NO. 2020-41

Labor Markets During the COVID-19 Crisis: A Preliminary View

Olivier Coibion, Yuriy Gorodnichenko, and Michael Weber

APRIL 2020

LABOR MARKETS DURING THE COVID-19 CRISIS: A PRELIMINARY VIEW

Olivier Coibion
UT Austin
and NBER

Yuriy Gorodnichenko
UC Berkeley
and NBER

Michael Weber
University of Chicago
And NBER

First Draft: April 11th, 2020

Abstract: We use a repeated large-scale survey of households in the Nielsen Homescan panel to characterize how labor markets are being affected by the covid-19 pandemic. We document several facts. First, job loss has been significantly larger than implied by new unemployment claims: we estimate 20 million lost jobs by April 8th, far more than jobs lost over the entire Great Recession. Second, many of those losing jobs are not actively looking to find new ones. As a result, we estimate the rise in the unemployment rate over the corresponding period to be surprisingly small, only about 2 percentage points. Third, participation in the labor force has declined by 7 percentage points, an unparalleled fall that dwarfs the three percentage point cumulative decline that occurred from 2008 to 2016. Early retirement almost fully explains the drop in labor force participation both for those survey participants previously employed and those previously looking for work.

JEL: E31, C83, D84, J21, J26

Keywords: Labor market, unemployment, employment, covid-19.

We thank the National Science Foundation for financial support in conducting the surveys. We also thank Shannon Hazlett and Victoria Stevens at Nielsen for their assistance with the collection of the PanelViews Survey. Results in this article are calculated based on data from The Nielsen Company (US), LLC and marketing databases provided by the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. Information on availability and access to the data is available at <http://research.chicagobooth.edu/nielsen>.

The arrival of the covid-19 virus and the policy responses have led to unprecedented numbers of initial claims for unemployment since early 2020: over 16.5 million by April 4th, 2020, with new claims arriving at a rate of 6-7 million per week. But concerns about state governments' inability to process so many claims in such a short period, combined with the fact that many workers are ineligible for unemployment benefits, has led to concerns that total job losses are being understated by these numbers. Furthermore, because official labor market indicators compiled by the Bureau of Labor Statistics (BLS) take time to be released, the current state of the U.S. labor market remains unclear.

Using new ongoing large-scale surveys of U.S. households much like the ones run by the BLS, we provide some preliminary evidence on the response of labor markets in the U.S. to the current crisis. We focus on three key variables typically measured by the BLS: the employment-to-population ratio, the unemployment rate, and the labor force participation rate. Historically, the employment-to-population ratio and the unemployment rate are near reverse images of one another during recessions as workers move out of employment and into unemployment (or workers in unemployment find it harder to move into employment). More severe recessions also sometimes lead to a phenomenon of "discouraged workers," in which some unemployed workers stop looking for work. This leads them to be reclassified as "out of the labor force" by the BLS definitions, so the unemployment rate can decline along with the labor force participation rate while the employment-to-population ratio shows little recovery, not because the unemployed are finding work but rather because they stop trying to find one. Jointly, these three metrics therefore provide a succinct and informative summary of the state of labor markets.

Using surveys prior to and at the height of the covid-19 crisis, we provide new estimates of how these variables have changed in the last two months. Our most recent numbers are from individuals surveyed April 2nd-8th, 2020, and therefore provide a sneak preview at what the equivalent BLS numbers will likely show when they are ultimately released. Our findings are striking. First, the employment-to-population ratio (the fraction of the adult population reporting that they had a paid job) has declined by about 7.5 percentage points. With an adult (civilian non-institutional) U.S. population of 260 million, this corresponds to nearly 20 million jobs lost as of April 8th. This estimate is in line with (albeit even higher than) new unemployment claims through this time period and confirms the widespread job loss.

Twenty million jobs lost relative to the pre-crisis labor force would correspond to an increase in the unemployment rate of 12.2 percentage points, so to a level of around 16%, were all the newly unemployed looking for work. This is, however, not what we find. When we construct an unemployment rate in an analogous manner as the BLS (i.e., define the unemployed as not working but looking for work), we find an increase in the unemployment rate of only 2 percentage points. This reflects the fact that most of the newly unemployed surveyed are not looking for new work at this time, so they are defined as out of the labor force rather than unemployed. Correspondingly, we document an extraordinary decline in the labor force participation rate of nearly 8 percentage points. In short, we find labor market changes that differ markedly from those of a typical recession.

We first mention the related literature, then describe the surveys and how the resulting labor market statistics before the crisis compare to those of the BLS prior to the crisis and finally turn to what our newer survey implies for labor markets during the covid-19 crisis.

I Related Literature

We relate to the fast-growing literature studying the economic consequences of the covid 19 pandemic. Binder (2020) shows that 30% - 40% of American are very concerned about the corona crisis, postponed travel and delayed purchases of larger ticket items as early as March 2020 but became more optimistic about the unemployment situation and revised downward their inflation expectations once being told about the cut in the federal funds target rate on March 3rd. Fetzner et al. (2020) show the arrival of the corona virus in a country leads to a large increase in internet searches around the world. In a survey experiment on a U.S. population, they find survey participants vastly overestimate the mortality rate and extent of contagion. Barrios and Hochberg (2020) and Allcott et al (2020) use internet searches, survey data, and travel data from smartphones to document that political partisanship determines the perception of risk associated with covid-19 and non-essential travel activity. Dingel and Neiman (2020) use data from responses to two Occupational Information Network surveys and estimate that about 37% of jobs can be performed from home, whereas Mongey (2020) documents that employees that are less likely to be able to work from home are mainly non-white and without a college degree. On the quantitative side, a growing literature jointly models the dynamics of the pandemic and the economy to quantify the economic costs and benefits of different economic policies (Atkeson (2020), Barro et al. (2020), Eichenbaum et al. (2020), Guerrieri et al. (2020), Alvarez et al. (2020), and Dietrich et al. (2020)). Finally, our Nielsen survey builds on previous work using the Nielsen panelists to study the formation and updating of economic expectations (Coibion et al. (2019, 2020) and D'Acunto et al (2020a, b)).

II Measuring Labor Markets using the Nielsen Survey

We start by describing the Nielsen Homescan survey and how it can be used to construct measures of labor market outcomes. The pre-crisis wave of the survey was run between January 6th and January 27th 2020 prior to much of the covid-19 outbreak. Potential participants were those households participating in the Nielsen Homescan, which is a panel of 80,000-90,000 households who track their purchases on a daily basis. Nielsen allows for surveys to be emailed to those households and respondents receive points and prizes for participating in Nielsen surveys. The panel of households used by Nielsen is meant to be representative of the U.S. population in age, size, income, etc. Possible imbalances are corrected using weights provided by Nielsen. We received 18,344 responses to the first wave of the survey.

In this and subsequent waves of the survey, we asked respondents several job-related questions. First, they were asked whether they have a paid job, with answers being either yes or no. Anyone answering yes we define as being employed. Note that this is slightly different from the BLS, which asks respondents whether they have

worked in the survey reference week, and those who “did any work at all for pay or profit” are classified as employed. This means some respondents who would be classified as employed by the BLS are classified as non-employed using our question. Consistent with this, we find somewhat lower employment rates (as a share of adult population) in our pre-crisis data than was the case in corresponding BLS surveys. Table 1, for example, shows that the BLS was reporting an employment to population ratio of 61.1% in February 2020 while our survey yielded an employment-to-population ratio of 57.7%.

Respondents who said they did not have a paid job were then asked if they were actively looking for a job, with possible answers being yes or no. We define those who answer yes as unemployed while those who answer no are classified as out of the labor force. Again, this is slightly different from the BLS questionnaire, which asks individuals to select ways in which they had looked for jobs during the prior 4 weeks and only classifies individuals as unemployed if they select answers which indicate a sufficiently active search such as posting resumes, contacting potential employers or filling out applications (i.e., not just scanning newspaper ads). Given that we allow individuals to specify themselves if they are “actively” searching, one might expect that this would also lead to a higher prevalence of unemployment in our Nielsen survey than in the corresponding BLS survey. Consistent with this, our estimated aggregate unemployment rate prior to the crisis is 8.6% while the corresponding BLS estimates for January and February of 2020 were 3.6% and 3.5% respectively. The labor force participation rate, however, is very similar to that estimated by the BLS: 63.1% vs 63.3% respectively.

While some difference in levels is to be expected given that questions are not identical across the two surveys, it is important to verify that they are still capturing broadly similar features. We do so by comparing state-level estimates of all variables from both the pre-crisis BLS and Nielsen surveys. These are illustrated in Figure 1, which shows state level unemployment rates from the two surveys in Panel A, state level labor force participation rates in Panel B, and state level employment to population ratios in Panel C. All three figures yield a strong positive correlation: the two surveys identify the same states as having relatively high or low unemployment, labor force participation, and employment to population ratios.

We can control for the initial difference in levels of unemployment that stems from the different questions asked in the surveys in the following way. Let $UE_{i,t}^{Nielsen}$ be the unemployment rate in state i in the Nielsen survey in year t and $UE_{i,t}^{BLS}$ be the unemployment rate in state i in the Current Population Survey (CPS) survey in year t . We run the following regression

$$UE_{i,t}^{Nielsen} = b_0 + b_1 UE_{i,t}^{BLS} + error \quad (1)$$

and then use the estimated coefficients to adjust Nielsen statistic for year s is

$$\widetilde{UE}_{i,s}^{Nielsen} = (UE_{i,s}^{Nielsen} - \hat{b}_0) / \hat{b}_1. \quad (2)$$

By construction, the average value of $\widetilde{UE}_{i,t}^{Nielsen}$ is now equal to the average value of $UE_{i,t}^{BLS}$ across states, which may still differ from the aggregate unemployment rate (which is a population-weighted average across states). We apply a similar procedure for employment-to-population ratios and labor force participation rates. The resulting

adjusted measures of aggregate unemployment, labor force participation, and employment-to-population ratios pre-crisis from the Nielsen survey are presented in Table 1, column 4. For easier comparison to BLS numbers, we focus primarily on these adjusted measurements (applied to both waves of our survey) from now on.

III Labor Markets since the Covid-19 Crisis

A second survey was run on households participating in the Nielsen Homescan panel between the afternoon of April 2nd-8th, 2020. The response rate was 18.6% with 9,445 responses. The same labor market questions were asked in both waves of the survey, so we can directly compare the two surveys to get a sense of how labor markets have evolved since the onset of the covid-19 virus and the associated policy responses. Table 1 summarizes results from the second wave of the survey. Our main results are as follows.

First, the employment-to-population ratio has declined sharply. Using the adjusted metrics described above, we find that the employment ratio fell from 60% of the population down to 52.2%, a nearly 8% point decline. As illustrated in Panel B of Figure 2, this decline in employment is enormous by historical standards and is larger than the entire decline in the employment to population ratio experienced during the Great Recession. Given that the U.S. civilian non-institutional population is approximately 260 million, this drop in employment to population ratio is equivalent to 20 million people losing their jobs. This drop is even larger than the 16.5 million new unemployment claims over this time period.

Second, we find a much smaller increase in the unemployment rate. As documented in Table 1, the adjusted unemployment rate rose from 4.2% to 6.3%. Panel A of Figure 2 plots this rise relative to previous changes in unemployment over the last 15 years. While this increase is the single biggest discrete jump in unemployment over the time period, this change in unemployment corresponds only to about one-third of the increase observed during the Great Recession. For comparison with the employment to population ratio, if all twenty million newly unemployed people as measured by the decline in the employment to population ratio were counted in the unemployment rate, we would have found an increase in the unemployment rate from 4.2% to 16.4%, the highest level since 1939.

The reason for the discrepancy between the two is that many of the newly non-employed people are reporting that they are not actively looking for work, so they do not count as unemployed but rather as exiting the labor force. Consistent with this, we find an extraordinary decline in the adjusted labor force participation rate, from 64.2% to 56.8%. For comparison, Panel C of Figure 2 plots the historical evolution of the labor force participation rate over the last 15 years which includes a historically large decline in participation between 2008 and 2016 of 3 percentage points. Even this cumulatively large decline in participation over an eight year period is dwarfed by the historic decline in participation that we document.

Figure 3 plots state-level changes in unemployment, labor force participation and employment to population ratios around the time of the covid-19 crisis. The forty-five degree lines indicate no change. When looking at labor force participation rates and employment-to-population ratios, we see that almost all states fall

below the 45 degree line, indicating the declines in each variable are widespread throughout the country and relatively homogenous in size. The change in unemployment across states is noisier, due in part to higher measurement error in measuring unemployment rates at local levels, but also indicates geographically dispersed increases in unemployment. Still, the pattern that comes out from state-level variation is one of broad-based declines in both employment and labor force participation across the country.

How unusual are these patterns? We have already seen that the size of the changes in each variable is exceptional, at least for employment and changes in labor force participation. What about their simultaneous changes? Figure 4 illustrates the historical comovement of the employment to population ratio with unemployment in Panel A and with labor force participation in Panel B. Panel A shows that, historically, unemployment and employment are very strongly negatively related. Within short periods, movements in one are reflected almost perfectly in the other as workers move from employment to unemployment and back. Slow-moving demographics cause the relationship to gradually change over time, as can be seen by decadal shifts in the curve, but short-run movements are close to linear. The change that we document since the covid-19 crisis jumps out: we see an enormous change in the employment-to-population ratio with a much smaller change in unemployment that would have typically been expected. This pattern is therefore qualitatively different from the historical experience of U.S. labor markets, even after taking into account the size of the changes.

We find a less unusual pattern relative to historical experience when looking at the change in labor force participation and employment to population ratios in Panel B. There, we see that the two tend to comove positively and closely on average: periods when employment growth is strong are also periods during which more people are entering the labor force. In that sense, the simultaneous decline in employment and labor force participation is mostly unusual because of the size of the changes. Still, the drop in labor force participation appears large relative to the historical experience given the size of the decline in employment, which is consistent with the smaller than normal increase in unemployment.

Why do so many unemployed choose not to look for work? Both surveys included a question asking those who said they were neither working nor looking for a job to select among possible answers why this was the case. The results for both waves are presented in Table 2. Prior to the crisis, most respondents out of the labor force claimed that it was because they were retired, disabled, homemakers, raising children, students, or did not need to work. Only 1.6% of those out of the labor force were claiming that they could not find a job as one of their reasons for not searching. At the height of the covid-19 crisis with a much larger number of people now out of the labor force, we see corresponding declines in the share of homemakers, those raising children and the disabled. However, we see a large increase in those who claim to be retired, going from 53% to 60%. This makes early retirement a major force in accounting for the decline in the labor-force participation. Given that the age distribution of the two surveys is comparable, this suggests that the onset of the covid-19 crisis led to a wave of earlier than planned retirements. With the high sensitivity of seniors to the covid-19 virus, this may reflect in part a decision to either leave employment earlier than planned due to higher risks of working or a choice to not look for new employment

and retire after losing their work in the crisis. To see this more clearly, we exploit the panel dimension of the survey and identify respondents who were out of the labor force in April 2020 but in the labor force in January 2020. Column (3) of Table 2 presents reasons for being out of the labor force reported by those who were *employed* in January. Of those, 28% report that they are now out of the labor force because of retirement. Similarly, of those who were *unemployed* in January and out of the labor force in April (column 4), 21% report that it is because they retired. And while 9% of those going from employment to out of the labor force between January and April report that they are on break from working (as might be the case for some affected by temporary work closures due to covid-19), the equivalent proportion is 8% for those going from unemployment to out of the labor force, suggesting that these breaks are voluntary ones, not ones due to temporary work closings. In short, these results point to an unusual rise in the share of retirements accounting for the exceptional decline in labor force participation during this time period.

IV Summary


It is still very early on in the covid-19 crisis, but preliminary indicators point toward catastrophic declines in employment. Our surveys provide additional evidence on this decline in employment, pointing to a 20 million decline in the number of employed workers. Most strikingly, we find a much less than proportional increase in unemployment, indicating that most of these newly unemployed workers are not looking for new work. Hopefully this reflects a transitory characteristic as these individuals face shelters-at-home and few work opportunities. But the wave of early retirements that we document suggests that more permanent changes may already be taking place.

References

- Allcott, Hunt, Levi Boxell, Jacob Conway, Matthew Gentzkow, Michael Thaler, and David Yang, 2020. "Polarization and Public Health: Partisan Differences in Social Distancing during the Coronavirus Pandemic," Working Paper.
- Alvarez, Fernando, David Argente, and Francesco Lippi (2020) "A Simple Planning Problem for COVID-19 Lockdown" University of Chicago, Becker Friedman Institute for Economic Working Paper No. 2020-34.
- Atkeson, Andrew (2020) "What Will Be the Economic Impact of COVID-19 in the US? Rough Estimates of Disease Scenarios" NBER Working Paper w26867.
- Barrios, John and Yael Hochberg, 2020. "Risk Perception Through the Lens of Politics in the Time of the COVID-19 Pandemic," Working Paper.


- Barro, Robert J., José F. Ursúa, and Joanna Wang (2020) “The Coronavirus and the Great Influenza Pandemic: Lessons from the ‘Spanish Flu’ for the Coronavirus’s Potential Effects on Mortality and Economic Activity” NBER Working Paper w26866.
- Binder, Carola, 2020. “Coronavirus Fears and Macroeconomic Expectations” Working Paper.
- Coibion, Olivier, Dimitris Georgarakos, Yuriy Gorodnichenko and Michael Weber, 2020a. “Forward Guidance and Household Expectations,” NBER Working Paper w26778.
- Coibion, Olivier, Yuriy Gorodnichenko and Michael Weber, 2019. “Monetary Policy Communications and their Effects on Household Inflation Expectations,” NBER Working Paper w25482.
- D’Acunto, Francesco, Ulrike Malmendier, Juan Ospina, and Michael Weber, 2020a. “Exposure to Daily Price Changes and Inflation Expectations,” NBER Working Paper w26237.
- D’Acunto, Francesco, Ulrike Malmendier, Juan Ospina, and Michael Weber, 2020b. “Gender Roles and the Gender Expectations Gap,” NBER Working Paper w26837.
- Dietrich, Alexander M., Keith Kuester, Gernot J. Müller, and Raphael S. Schoenle (2020) “News and Uncertainty about COVID-19: Survey Evidence and Short-Run Economic Impact,” Becker Friedman Institute for Economic White Paper.
- Mongey, Simon and Alex Weinberg, 2020. “Characteristics of Workers in Low Work-From-Home and High Personal-Proximity Occupations,” Becker Friedman Institute for Economic White Paper.
- Eichenbaum, Martin S., Sergio Rebelo, and Mathias Trabandt (2020) “The Macroeconomics of Epidemics” NBER Working Paper w26882.
- Fetzer, Thiemo, Lukas Hensel, Johannes Hermlle, and Christopher Roth, 2020. “Perceptions of Contagiousness and Mortality of COVID-19 Shape Economic Sentiment,” Working Paper.
- Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Ivan Werning, 2020. “Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?” NBER Working Paper w26918.

Figure 1: Pre-crisis employment statistics at the state level.




Notes: Bureau of Labor Statistics (BLS) data refer to the January-February 2020 period. Nielsen data are for the pre-crisis wave of the survey.

Figure 2. Time series of key employment statistics.




Notes: Each panel plots monthly time series of an employment statistic. The black, solid line shows data from the Bureau of Labor Statistics (BLS). The red, solid line with circles shows the corresponding values from the Nielsen survey. 2020M2 are the values from the Nielsen pre-crisis survey. 2020M3 are the values from the Nielsen crisis wave.

Figure 3. Adjusted employment statistics by state, pre-crisis vs. crisis levels, Nielsen survey.




Notes: Nielsen employment statistics are adjusted to match average pre-crisis levels observed in the official data compiled by the Bureau of Labor Statistics (BLS). The adjustment is described in equations (1)-(2).

Figure 4. Historical comovement of key employment statistics.
 Panel A: Employment-to-population ratio vs Unemployment rate.



Panel B: Employment-to-population ratio vs Labor force participation rate.



Notes: Each panel shows comovement (by decade) of key official employment statistics compiled by the Bureau of Labor Statistics (BLS) as well as employment statistics based on Nielsen surveys (red circles with dates). 2020M2 are the values from the Nielsen pre-crisis survey. 2020M3 are the values from the Nielsen crisis wave.

Table 1. Employment statistics at the state and aggregate levels, Nielsen survey.

	BLS pre-crisis	Nielsen			
		Raw moments		Adjusted moments	
		pre-crisis	crisis	pre-crisis	crisis
	(1)	(2)	(3)	(4)	(5)
Panel A. State-level data (average)					
Unemployment rate	3.52 (0.84)	7.84 (2.80)	10.93 (5.07)	3.50 (2.55)	6.30 (4.61)
Labor force participation rate	63.91 (3.94)	62.89 (7.19)	56.94 (8.58)	63.97 (8.17)	57.20 (9.75)
Employment to population ratio	61.67 (4.03)	58.00 (7.27)	50.77 (8.63)	61.74 (7.68)	54.11 (9.12)
Panel B. Aggregate level					
Unemployment rate		8.62 [0.25]	10.95 [0.43]	4.21 [0.23]	6.32 [0.39]
Labor force participation rate	63.4	63.13 [0.33]	56.57 [0.51]	64.23 [0.38]	56.79 [0.58]
Employment to population ratio	61.1	57.68 [0.34]	50.38 [0.52]	59.94 [0.36]	52.22 [0.54]

Notes: Panel A report statistics across states. Averages are in the top row of each subsection and standard deviations are in the bottom rows (in parentheses) of each subsection. Panel B reports statistics for the aggregate level. Point predictions are reported in the top row of each subsection and standard errors are reported in the bottom rows (is in square brackets) of each subsection. Column (1) reports pre-crisis data from the Bureau of Labor Statistics (BLS), average values for January-February 2020. These data are based on the Current Population Survey (CPS). Columns (2) and (3) report moments for the Nielsen survey unadjusted for possible differences in design between the CPS and Nielsen surveys. Columns (4) and (5) report moments for the Nielsen survey adjusted for possible differences in design between the CPS and Nielsen surveys. Adjustment is based on state-level regressions and is given by equations (1)-(2).

Table 2. Reasons for not looking for a job (for those who do not have a job).

Reason	Share of people choosing a reason			
	All people		Employed pre-crisis, out of labor force in crisis	Unemployed pre-crisis, out of labor force in crisis
	Pre-crisis	Crisis		
(1)	(2)	(3)	(4)	
Homemaker	0.204	0.155	0.158	0.244
Raising children	0.125	0.082	0.115	0.222
Student	0.023	0.018	0.050	0.019
Retiree	0.527	0.595	0.277	0.210
Disabled, health issues	0.297	0.261	0.142	0.144
Couldn't find a job	0.016	0.018	0.061	0.133
On break	0.012	0.011	0.092	0.080
No financial need	0.049	0.037	0.039	0.029

Notes: the Nielsen survey question is “Here are a number of possible reasons why people who are not working choose not to look for work. **Please select all that apply to you.**” This question is asked for people who do not have a job and are not looking for a job.