

Ludwig Institute for Shared Economic Prosperity Technical Documentation

Written by Philip Cornell on behalf of LISEP

Chapter 1

Methodology for LISEP's True Rate of Unemployment
(TRU)

Chapter 2

Methodology for LISEP's True Usual Weekly Earnings

Chapter 1

Methodology for the LISEP True Rate of Unemployment

Written by Research Assistant Philip Cornell on behalf of the
Ludwig Institute for Shared Economic Prosperity

This document details the calculation and verification of the True Rate of Unemployment (TRU).

I. True Rate of Unemployment

Our data comes from the [IPUMS CPS database](#) (see endnotes for formal citation of the IPUMS database). According to this website: “IPUMS CPS harmonizes microdata from the monthly U.S. labor force survey, the Current Population Survey (CPS), covering the period 1962 to the present. Data include demographic information, rich employment data, program participation and supplemental data on topics such as fertility, tobacco use, volunteer activities, voter registration, computer and internet use, food security, and more.”

The BLS seeks to answer the question of “What percent of the people that are currently seeking jobs are unsuccessful?” We are broader in our aim. We develop the TRU to answer the question of “What percent of the labor force is functionally unemployed?” Functional employment is employment in which one is able to survive and simultaneously has the potential to advance their welfare.

LISEP’s definition of “TRU” accepts the Bureau of Labor Statistics U-3 unemployment rate for comparison purposes but modifies it by adopting two important stipulations. The first stipulation deals with the workweek. To be classified as employed for LISEP’s true employment concept, an individual must either have a full-time job (35+ hours per week) or a part-time job and no desire for a full-time job (e.g., students). The second stipulation is that an individual must earn at least \$20,000 annually. This annual wage is adjusted for inflation, calculated in 2020 dollars. (\$20,000 was chosen based on the U.S. poverty guidelines put out by the Department of Health and Human Services, which considers a three-person household to be in poverty if it has an income of less than \$20,000 per year). Because we build upon the U-3 unemployment rate, we also use the BLS-defined labor force as our sample. That definition is found here:

<https://www.bls.gov/cps/definitions.htm>. To ensure that a person is in the labor force, we use the variable LABFORCE, which is defined by IPUMS as:

- “LABFORCE is a dichotomous variable indicating whether the respondent participated in the labor force during the preceding week. See EMPSTAT for a more detailed employment status variable. Those coded as "yes" in LABFORCE were either: were at work; held a job but were temporarily absent from work due to factors like vacation or illness; were seeking work; or were temporarily laid off from a job during the reference period.”

If the LABFORCE variable indicated that the respondent is in the labor force, then the respondent was eligible for the sample within TRU. This is analogous to the BLS U-3 unemployment rate. So, before any other calculations, we excluded those who are not in the labor force. We take an alternative approach in Section III.

To calculate TRU, we treat wage and salaried workers differently from the self-employed. We needed to use different methods to calculate the true employment of wage and salaried workers versus the self-employed because of the lack of availability of income data for the self-employed population. To separate the sample, we use the variable CLASSWKR, defined by IPUMS as:

- “CLASSWKR indicates whether a respondent was self-employed, was an employee in private industry or the public sector, was in the armed forces, or worked without pay in a family business or farm. Workers with multiple sources of employment were classified according to the job in which they worked the most hours. For persons employed at the time of the survey, CLASSWKR relates to the respondent's job during the previous week.”

The potential responses to this are:

CODE	DEFINITION
13	Self Employed, not incorporated
14	Self Employed, incorporated
22	Wage/salary—private, for profit
23	Wage/salary—private, nonprofit
25	Wage/Salary— government, federal government employee
26	Wage/Salary—government, armed forces

27	Wage/Salary—government, state government employee
28	Wage/Salary—government, local government employee
29	Unpaid family worker

Using these responses, we first excluded the unpaid family workers from the entire labor force. This is a small subset of the data (less than 0.0005 of the sample). The reason for this is because the wages these people report are not accurate. Although they report zero wages, they are most likely working for other tangible items (e.g., education, food, shelter) that are being provided to them by the family. But it would be impossible to accurately estimate the dollar amount of these items, so we dropped them from the set. This exclusion makes little difference given the small share of unpaid family workers.

Once the respondents were identified as self-employed or wage earners, the next step is to calculate the percent of the sample that is part of the self-employed category for each month. We needed to produce an aggregate number, which required weighing the true employment of wage workers versus the self-employed by their representative size when we combined the two separate types of workers.

Once these proportions were calculated, we dropped the self-employed respondents out of the dataset (to be considered later in this paper) and looked solely at wage earners.

A. True Employment Calculation for Wage Earners

For the wage earners, we use the Outgoing Rotational Group Survey from the CPS. More information about the ORG and its methods can be found here:

https://cps.ipums.org/cps/outgoing_rotation_notes.shtml.

We determined that this portion of the CPS respondents would give us the best estimate for TRU, as it is the only sample of the wage earners for which we have both the full CPS Questionnaire about the type of work that the respondent participates in, as well as information on their wage. The downside of using the ORG sample is that it only accounts for one-fourth of the respondents of the entire CPS in a given month. Therefore, the sample that we use (identical to the sample used by the Federal Reserve Bank of Atlanta to track wages) is a smaller portion of the sample that BLS uses to publish the unemployment rate each month.

1. Full-time Definition

There are multiple ways to satisfy the first stipulation, which is a full-time job. First, the job must be 35+ hours per week. Second, one could work a part-time job and prefer this part-time status. The exact variable from IPUMS used to make this determination is “WKSTAT.” IPUMS defines

WKSTAT as:

- “[A] recode from the Census Bureau that states the part-time or full-time employment status for the respondent, and reasons. It is derived from a number of labor force questions asked in the monthly questionnaire.”

WKSTAT takes on many values, but the relevant ones are shown below.

CODE	DEFINITION
10	Full-time schedules
11	Full-time hours, usually full-time
12	Part-time hours for non-economic reasons, usually full-time
13	Not at work, usually full-time

If this variable takes on the codes of 10, 11, 12, or 13, then we count this person as being fully employed. This is the first way to meet the full-time stipulation.

If a person has a part-time job and prefers this over having full-time employment, we use the variable “WHYPTLWK,” which is described as:

- “WHYPTLWK reports the reason why respondents worked part-time (a total of less than 35 hours combined for all jobs) during the previous week. Some of these individuals normally worked a part-time job; others usually worked full-time but worked less than 35 hours during the week in question.”

We accept that some people choose to be part-time for wholly noneconomic reasons, so we did not want to include them as part of TRU. We wanted to consider only the part-time workers who were involuntarily part-time within TRU. In the interest of developing a statistic that was transparent, we did not want to misreport a high TRU if a proportion of those whom we categorized as unemployed were part-time by choice. This is an important departure from both U-3 and U-6. The former does not consider part-time workers as unemployed at all, while the latter includes all workers who are part-time employed for economic reasons. This definition of “for non-economic reasons” is a smaller category than what we exclude. Most notably, we do not consider those who report full-time status but cannot get 35+ hours per week at this job as choosing to be part-time. U-6, on the other hand, reports that those whose full-time jobs are not actually full-time by hourly definitions as still employed.

So, to meet LISEP’s full-time stipulation, the respondent had to note that he/she worked part-time last week because of one of the following reasons:

CODE	DEFINITION
30	Weather Affected Job
50	Job Started/ended during work week
90	Holiday
100	Own illness
101	Health/Medical Limitation
111	Vacation/Personal Day
121	Child Care Problems
122	Other Family/personal Obligations
123	School/training
124	Civic/Military duty

Because these people are working part-time for the above noneconomic factors, and most of these respondents would still choose part-time if given the choice, they should be classified as truly employed.

2. \$20,000 Wage Stipulation

The second stipulation that separates the TRU from the U-3 unemployment rate relates to the respondent's wage. Specifically, the respondent's job must be paid a wage of at least \$20,000 (for January 2020 using the CPI-U) before taxes for the respondent to be eligible. The variables that we use for this calculation are "EARNWEEK" and "WKSWORKORG." These are defined by IPUMS to be:

- "EARNWEEK reports how much the respondent usually earned per week at their current job, before deductions. Interviewers asked directly about total weekly earnings and also collected information about the usual number of hours worked per week and the hourly rate of pay at the current job. The figure given in EARNWEEK is the higher of the values derived from these two sources: 1) the respondent's answer to the question, "How much

do you usually earn per week at this job before deductions?" or 2) for workers paid by the hour (and coded as "2" in PAIDHOUR), the reported number of hours the respondent usually worked at the job, multiplied by the hourly wage rate given in HOURWAGE."

- "WKSWORKORG reports the number of weeks, in single weeks, that the respondent usually works per year. WKSWORKORG is one of the earner study questions, which is also known as the outgoing rotation group questions. Private wage or salaried workers are asked the periodicity for which they are paid. Workers who are paid annually, excluding those who are self-employed, are asked how many weeks a year they are paid for."

Using these two variables, we calculate the amount of yearly earnings by multiplying EARNWEEK with WKSWORKORG. If there is a missing value in the WKSWORKORG, we take the EARNWEEK value and multiply it by 50, assuming 50 weeks worked per year. Because the WKSWORKORG value is only recorded for salaried workers, this assumption applies to a large part of the workforce (77.4% of the sample). For this reason, we checked the robustness of this assumption (detailed later in the robustness check section). We then calculated the current income of the respondent by adjusting the response to the currency value of the 2020 January USD. We downloaded the CPI from Federal Reserve of St. Louis on the FRED database, which can be found [here](#).

B. True Employment Calculation for Self-Employed

The above approach was used for wage and salaried workers. We also had to account for the self-employed part of the workforce. For this, we used the Annual Social and Economic Supplement (ASEC) to calculate the TRU for the self-employed at the time that the ASEC data was collected (once per year).

Unfortunately, the self-employed are not included in the ORG survey that we use in the TRU calculation for the wage earners. Therefore, although we are aware of the self-employed full- or part-time status each month, there are no indicators of self-employed income levels on a monthly basis. So, after recording their proportion of the workforce in each month, we needed to use the ASEC to calculate their wages.

The ASEC asks the respondents about the previous year, so we adjusted for the fact that the 2019 ASEC is referring to the 2018 year and do this for all of the years in the sample. For this reason, we could not also apply the WKSTAT variable that we use for wage earners because this question asks about the previous week. The income data is for the previous year, and so we must also use the working status from the previous year in order to align the responses.

1. Full-Time Stipulation

To determine the full- or part-time status of the self-employed workers, we use the variable UHRSWORKLY, which is defined by IPUMS as:

- "UHRSWORKLY reports the number of hours per week that respondents usually worked if they worked during the previous calendar year. Individuals were asked this question if:

1) they reported working at a job or business at any time during the previous year or 2) they acknowledged doing "any temporary, part-time, or seasonal work even for a few days" during the previous year."

If UHRSWORKLY is 35 hours or greater, then we consider this person to be a full-time worker. We decided to also include the variable WHYPTLY in the code as it is analogous to the WHYPTLW variable that is utilized in the wage earners data. This is defined by IPUMS as:

- "WHYPTLY reports the reason why respondents worked part-time (less than 35 hours) for at least one week during the previous calendar year. Some of these individuals normally worked a part-time job; others usually worked full-time but worked less than 35 hours for some weeks (e.g., because of slack work or a shortage of materials). Paid time off due to vacations, holidays, or sick leave did not count."

The possible responses to this variable are:

CODE	DEFINITION
1	Could not find a full-time job
2	Wanted part-time
3	Slack work
4	Other

We considered respondents coded with the value 2 to be truly employed. This is because these respondents clearly indicated a preference for part-time rather than full-time employment.¹

2. \$20,000 Wage Stipulation

To calculate the wages of the self-employed, we combined four different variables from the ASEC: INCWAGE, INCBUS, INCFARM, and INCSS. They are defined here:

- "INCWAGE indicates each respondent's total pre-tax wage and salary income--that is, money received as an employee--for the previous calendar year."
- "INCBUS indicates each respondent's net pre-income-tax non-farm business and/or professional practice income for the previous calendar year. INCBUS is reported for self-employed persons; employees' earnings are given in INCWAGE."

¹ The "Other" category was roughly 3.7% of responses.

- “INCFARM indicates each respondent's net pre-income-tax earnings as a tenant farmer, sharecropper, or operator of his or her own farm during the previous calendar year. INCFARM collects income information for self-employed persons who had their own farms. Income earned as an employee on a farm is contained in the variable INCWAGE.”
- “INCSS indicates how much pre-tax income (if any) the respondent received from Social Security.”

The reason we include Social Security payments in the data is because of the relevance to elderly self-employed. Social Security income is a huge factor in the decision-making process of this segment of the population. They may have a small business or side job to pursue hobbies or supplement their income. This could be consulting, ride-share driving, or selling a homemade good. Their goal for this work might not be to sustain their living, so we are factoring in social Security income to get a better idea of their actual income for the period. Social Security, versus other types of income, is directly connected to one's labor in past years. Thus, it is more reasonably seen by an individual as a type of paycheck earned through past labor.

We combined all four of these variables to account for the income that a self-employed person might have. We then calculated the respondent's income by adjusting their response to the currency value of the 2020 January USD. We use the CPI from the Federal Reserve of St. Louis on the FRED database.

3. Linear Interpolation for Self-Employed

After calculating the TRU for the self-employed, we then took this TRU value and attributed it to June of the preceding year. The survey is in March of the current year, but all of the questions are about the preceding year, so we chose the sixth month of the year to spot the measure. We then linearly connected this TRU with the year before and the year after. For example, if we (entirely hypothetically) found that the TRU of 1980 is 12% and the TRU of 1981 is 24%, we would then have allocated these values such that in June 1980, the TRU for the self-employed is 12%; in July 1980, the TRU is 13%; August is 14%, and so on, so that this aligns with the fact that in June of 1981, the TRU would be calculated as 24%.

There is no perfect way to merge the annual self-employed data with the monthly wage earner data. Moreover, the problem is exacerbated in the present. This is because the annual data not only comes out once per year, but it also comes out at a lag. So, for the entire year until the new data comes out, we have to keep a linear trend of the self-employed. For example, we used the linear trend from 2018 to 2019 to calculate our 2020 values. We will not be able to come out with more accurate numbers until the 2020 values are released, and thus will revise the past numbers when the data becomes available. We recognize that this is not ideal, but we are attempting to be as precise as possible while also making the least number of assumptions given data availability.

C. Aggregation of True Rate of Unemployment (TRU)

We then took the TRU number calculated for the wage and salary earners and aggregated this with the self-employed. We used the proportions for each month calculated from the CLASSWKR variable outlined above. We took the share of the population that is self-employed and multiplied this by the TRU for the self-employed and then added this to the TRU for the wage and salary workers multiplied by their proportional representation in the sample. The final TRU is the weighted average of the wage and salaried worker TRU and self-employed TRU. We do this aggregation on a month-by-month basis to obtain a monthly TRU value.

II Analysis of Subgroups

We also stratified the data to create the TRU for different groups within the population. The different groups can be sorted into the broad categories of demographic groups and educational groups.

A. Demographic Groups—Race/Ethnicity and Gender

First, we calculated the TRU for three large racial/ethnic groups within the U.S. population. These groups are White (Non-Hispanic), Black, and Hispanic. To sort the respondents, we use RACE and HISPAN from the IPUMS website. They are defined as:

- “Racial categories in the CPS have been more consistent than racial categories in the census. Up through 2002, the number of race categories ranged from 3 (white, negro, and other) to 5 (white, black, American Indian/Eskimo/Aleut, Asian or Pacific Islander, and other). Beginning in 2003, respondents could report more than one race, and the number of codes rose to 21, and then up to 26 codes in 2013.”
- “HISPAN identifies and classifies persons of Hispanic/Spanish/Latino origin. Origin is ancestry, lineage, heritage, national group, or country of birth. Prior to 2003, information was collected by asking, "What is the origin or descent of each person in this household?" and asking the respondent to select the appropriate category from a limited number of choices on a flashcard (including "another group not listed.") The choices included five to eight choices that would be classified as Hispanic, "Negro" and "Black," and a small number of European ancestry groups such as "German." The primary intention of the question was to identify Hispanic respondents, rather than origin or descent for the general population. Beginning in 1976, the original CPS data preserved detail for only the Hispanic responses, with all other answers lumped together as "another group not listed" (relabeled "Not Hispanic" in IPUMS-CPS). In 2003 and later years, respondents were asked, "Are you Spanish, Hispanic, or Latino?" rather than the broad query about origin or descent. Detailed information about Hispanic ethnicity was collected only from those who answered "yes" to this initial question.”

RACE takes on several values, which are presented in this table:

CODE	DESCRIPTION
------	-------------

100	White
200	Black
300	American Indian/Aleut/Eskimo
651	Asian
652	Hawaiian/Pacific Islander
801	White-Black
802	White-American Indian
803	White-Asian
804	White-Hawaiian/Pacific Islander
805	Black-American Indian
806	Black-Asian
807	Black-Hawaiian/Pacific Islander
808	American Indian-Asian
809	Asian-Hawaiian/Pacific Islander
810	White-Black-American Indian
811	White-Black-Asian
812	White-American Indian-Asian
813	White-Asian-Hawaiian/Pacific Islander

814	White-Black-American Indian-Asian
815	American Indian-Hawaiian/Pacific Islander
816	White-Black-Hawaiian/Pacific Islander
817	White-American Indian-Hawaiian/Pacific Islander
818	Black American Indian-Asian
819	White-American Indian-Asian-Hawaiian/Pacific Islander
820, 830	Unspecified Mixed Races

HISPAN has the codes in the table below:

CODE	DESCRIPTION
0	Not Hispanic
100	Mexican
200	Puerto Rian
300	Cuban
400	Dominican
500	Salvadoran
600	Other Hispanic
611	Central American, (excluding Salvadoran)
612	South American

Specifically, we designated White Non-Hispanic respondents to be those that had a RACE value of 100 and a HISPAN value of 0. We defined the Black population to include those who are Black and also those who identified as mixed race in which one of those races is Black: the RACE values of 200, 801, 805, 806, 807, 810, 811, 814, 816, or 818. We conducted a robustness test that excludes mixed race, detailed later in the paper. Lastly, we designated the respondent as Hispanic if the HISPAN variable took on values that were 100, 200, 300, 400, 500, 600, 611, or 612.

Once these populations were sorted, we eliminated the respondents from the dataset that did not belong to these groups and calculated them in the same way laid out in Section I.

We also broke down the population by gender. For this we use the variable SEX, which took on the values:

CODE	DESCRIPTION
1	Male
2	Female

We calculated the TRU separately for each gender with the same process laid out in Section I. Furthermore, in a departure from the BLS's reporting of statistics, we did not exclude those in the 16-to-19-year-old age range when we reported the data by gender. The BLS gender breakdown is male, 20 or older, and female, 20 or older. We made no such age distinctions. To assure that the trends hold, we conducted a robustness check with this age exclusion.

B. Education

We sorted the responses by educational groups as well, calculating the TRU for each educational cohort. We chose five distinct groups based on their highest level of educational attainment: less than high-school degree, high-school diploma or equivalent, some college/associate's degree/vocational degree, bachelor's degree, advanced/professional degree. To assign people to these education groups, we use the EDUC variable. This is defined as:

- “EDUC indicates respondents' educational attainment, as measured by the highest year of school or degree completed. Note that completion differs from the highest year of school attendance; for example, respondents who *attended* 10th grade but did not finish were classified in EDUC as having *completed* 9th grade. EDUC is a combination of two other variables, HIGRADE and EDUC99, which measure educational attainment in different ways. HIGRADE is available for years prior to 1992 and gives the respondent's highest grade of school or year of college completed. EDUC99 is available beginning in 1992

and classifies high school graduates according to their highest degree or diploma attained.”

The potential codes for this variable are summarized in the table below:

CODE	DESCRIPTION
002 to 071	No schooling/preschool/kindergarten through Grade 12, no diploma
073	High school diploma or equivalent
081	Some college but no degree
091	Associates degree, occupational/vocational program
092	Associates degree, academic program
111	Bachelor’s degree
123	Master’s degree
124	Professional school degree
125	Doctorate degree

We sorted the groups by allocating the codes 002-071 to No HS Degree, 073 to HS Diploma, 081-092 to Some College, 111 to bachelor’s degree, and 123-125 to Advanced/Professional Degree. We then calculated the TRU for the population in the same way that is detailed in Section I.

III True Rate of Unemployment (TRU) Out of the Population

We designed TRU to be analogous to the BLS U-3 unemployment rate. The U-3 uses the labor force as its population and calculates the unemployment rate accordingly. Additionally, the BLS publishes their much-discussed labor force participation rate (LFPR), which shows the share of the working-age population that is employed or actively looking for work. However, although the LFPR provides some insights, it leaves out the important information conveyed by the unemployment rate, such as actually how many people are employed. Therefore, a combination

of the two measures is the employment-to- population ratio. This is defined as the number of people that are employed as a ratio of the entire adult population. Because of the crucial information that this statistic conveys, we produced our “True Rate of Unemployment Out of the Population.” This number conveys similar information as the employment-to-population ratio. Namely, by taking the entire civilian, noninstitutional population as the denominator, we considered those who are not in the labor force as well.

A. Definition of the Population

The BLS defines population as the civilian, non-institutional population, as per its website:

“Included are persons 16 years of age and older residing in the 50 states and the District of Columbia who do not live-in institutions (for example, correctional facilities, long-term care hospitals, and nursing homes) and who are not on active duty in the Armed Forces.”

We included members of the Armed Forces in our definition of the population.² Because those in the military make more than \$20,000 a year and have full-time employment, we deemed them truly employed. Other than this, we make no other alterations to the civilian, non-institutional population as defined by the CPS.

Unfortunately, the CPS states that: “Most group quarters are of an institutional nature (for example, prisons and nursing homes) or military (for example, barracks) and therefore have never been part of the CPS sample.” So, because of this, we are constrained by the CPS data collection procedures. Although we included those in the Armed Forces who are excluded by the BLS (those not in barracks, but still in the military), we have an imperfect measure of the adults living in the United States. We are unable to account for those in nursing homes, prisons, and psychiatric wards, among other such places. The most interesting of these variables is that of the incarcerated population. For this, we are currently working on implementing lagged data provided by the Bureau of Justice Statistics on the incarcerated population in both county prisons, as well as state and federal facilities. Until we have this methodology more robustly confirmed, we have restricted our population to the CPS sample.

B. Calculations

Our methodology matches that which we outlined in Section I, with the exception that we included those respondents who are not in the labor force. In Section I, we excluded those outside the labor force as indicated by the variable LABFORCE. Furthermore, as the BLS’s unemployment rate is a negative qualifier – it measures those who are NOT employed – we matched this with the TRU. But the BLS publishes the employment-to-population ratio as a positive qualifier, measuring those who are employed. However, we believed it would be more meaningful for our purposes to also use the negative qualifier for the TRU Out of the Population

² We do still use the BLS data and survey. Because of this, if BLS surveyors make exclusions to not go into military zones, or to under survey those in army barracks, we are also only privy to this data. More information can be found at the CPS survey methodology, linked [here](#).

so that the public could more clearly see who among the whole population is not functionally employed. We also stratified this sample with the same stratifications that we outlined in Section II.

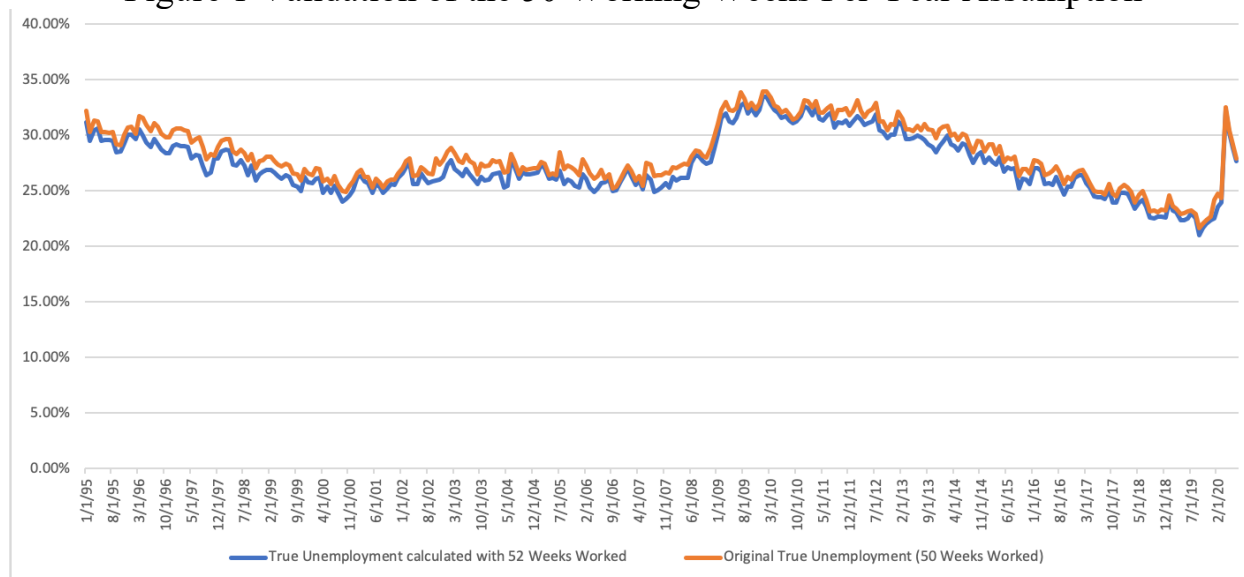
IV Robustness Checks

We ran several robustness checks on different aspects of the data throughout the process. We have listed them here and will explain them subsequently. All of the robustness checks are done on non-seasonally adjusted data.

1. Amount of Work Weeks Per Year

In the collection of the data, not every respondent reported the average number of weeks that he/she worked per year. Therefore, when calculating the yearly wages for the respondents, we had to assume a number of weeks worked. We assumed 50 working weeks, which gives two weeks of personal/vacation time. To confirm the robustness of this result, we calculated both 52 and 50 weeks worked per year. These two results are shown in Figure 1 below. Although the expected result of assuming more work and thus a higher wage lowers the number, the trends remain the same and the result is not significantly different. This check validates the appropriateness of the 50-week worked assumption.

Figure 1 Validation of the 50 Working Weeks Per Year Assumption

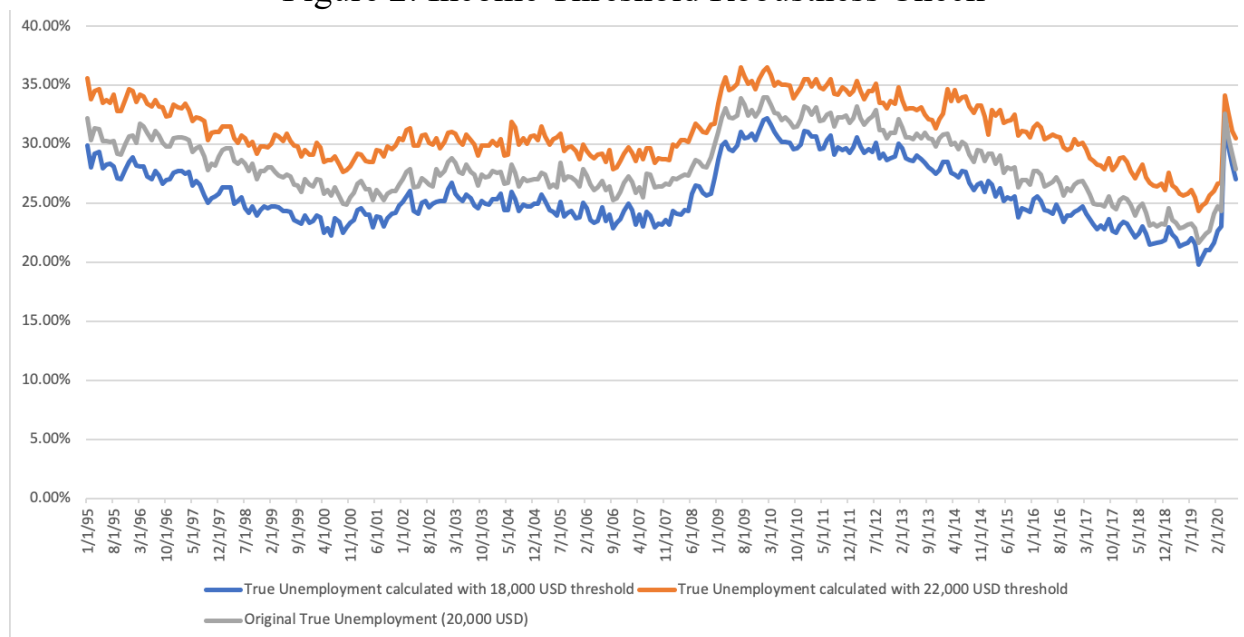


2. Income Threshold

We also conducted a robustness check on the income threshold. This is to show that the choice of \$20,000 as the cutoff is not critical. We modeled both the rate if we define the wage stipulation by yearly earnings of \$18,000, and also if we define the yearly earnings by \$22,000. We wanted the \$20,000 threshold to be a general measure of poverty. If this is true, then we would expect

the rate chosen by \$20,000 to be roughly symmetrically positioned in between the rate at \$22,000 and the rate at \$18,000. This is shown in Figure 2 below. All of the numbers both move together and are close together, as we would expect with similar income thresholds. Thus, we can conclude that the \$20,000 is not a special number, and that the TRU is not falsely inflated by this threshold. In other words, it is clear that the TRU is much higher than the BLS-reported unemployment rate, regardless of the exact number chosen for the income cutoff.

Figure 2: Income Threshold Robustness Check



3. Validation of the Income Calculation of the Self-Employed

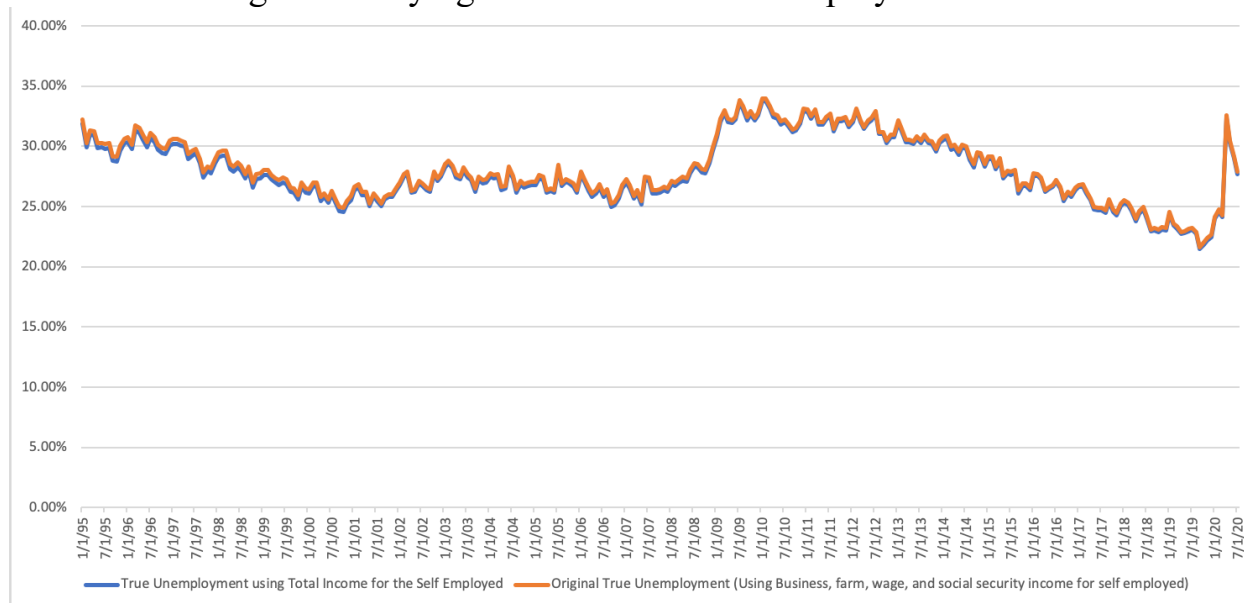
We sought to validate that someone who is self-employed relies on the income that is aggregated from their business, wages, and Social Security. These are the variables that we use to calculate the self-employed wage. One might critique this by suggesting that the self-employed are financially secure with other assets, that they choose to be self-employed with low hours/low wages because of other sources of financial security. To assess if this significantly drives the results, we changed the definition of income for the self-employed. In the robustness check, we took the self-employed income to be the INCTOT variable, which is defined by the IPUMS as:

- “INCTOT indicates each respondent's total pre-tax personal income or losses from all sources for the previous calendar year. Amounts are expressed as they were reported to the interviewer.”

This is an extremely conservative estimate, and in all cases includes the income that is generated by the four variables that we used to define self-employed labor income. In this case though, the income total includes other forms of income, primarily capital gains income or transfer payments.

We set the INCTOT as the wage variable, and then allowed eligibility for true employment if this INCTOT is greater than \$20,000 (adjusted to January 2020). This effort is to measure the statistical importance of those who have stable sources of other income and a “side” business that is not profitable. We still excluded those who don’t satisfy the full-time constraint (as defined in Section I). The results are shown in Figure 3. The results suggest that this portion of the population is insignificant. The original TRU and the TRU taken with total income are almost identical. Thus, we can verify that the four sources of income that we used in the construction of the self-employed income are representative of their financial situation.

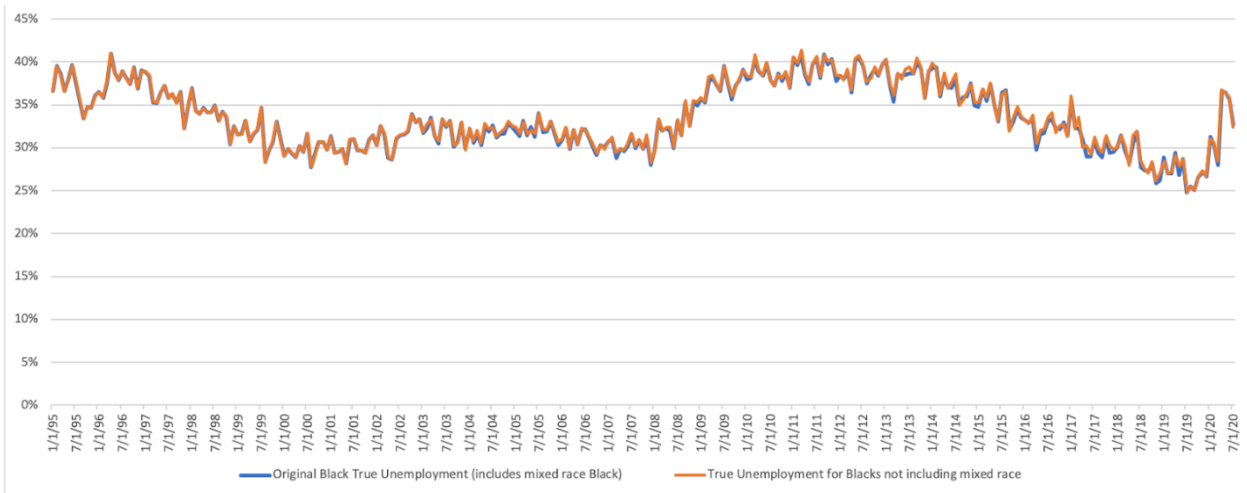
Figure 3 Varying Definitions of Self-Employed Income



4. Black Population Excluding Mixed Race

In determining the original TRU for Blacks, we classified respondents who identified as partially black as being Black or African American. We did not do this with the White Non-Hispanic population. To confirm the robustness of our number – that it is representative of those who identify as Black – we did a second calculation that excluded mixed-race people. Specifically, we only included the population coded as 200 for RACE (page 11 above). As depicted in Figure 4 below, the two TRU are virtually identical, and therefore including respondents who indicated mixed-race Black as well as just those who responded Black shows similar trends for these populations.

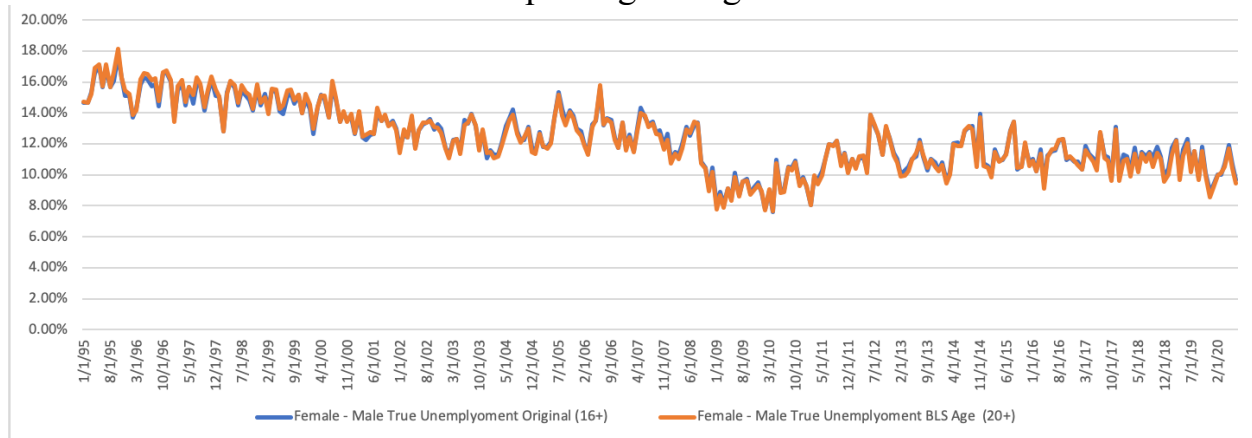
Figure 4 Verification of Trends with More Precise Versus More Inclusive Definitions of Black/African American



5. Gender Alignment to Compare to BLS

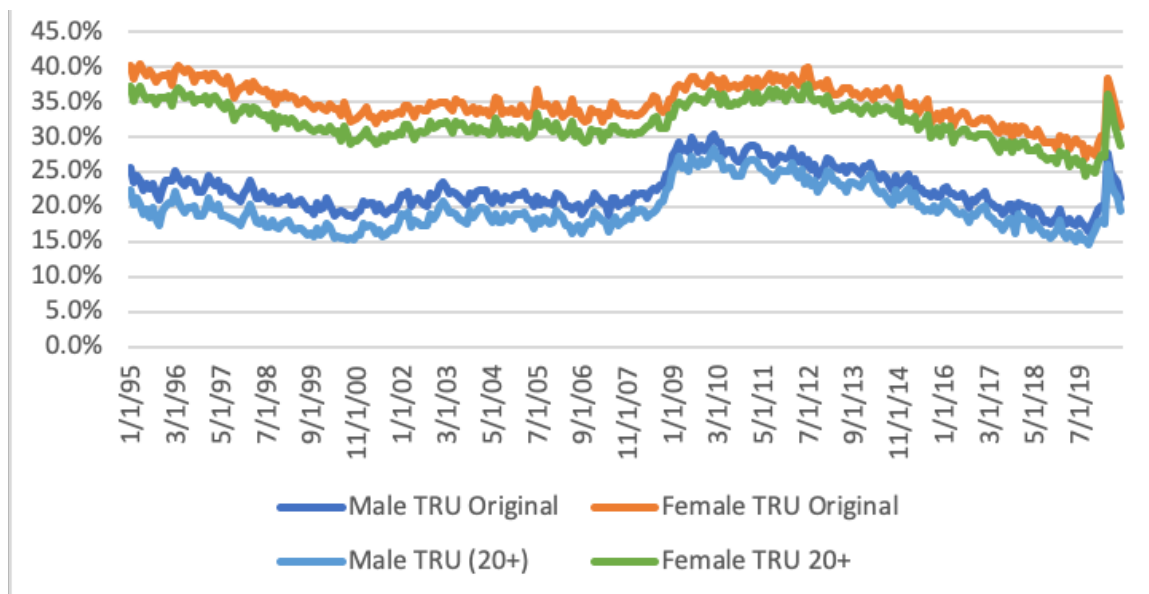
We also calculated the TRU by gender only including the 20+ population for each gender. The BLS publishes the U-3 rate by the two gender categories: men 20+ years and women 20+ years. Because the TRU is meant to be easily compared to the BLS unemployment rate, we needed to confirm that the gender trend for our number is very similar to the aged 20+ in each gender. So, although the absolute numbers will undoubtedly be different, the gender breakdown is used to compare the employment situation of males versus females. We calculated our measure (which includes age 16+) and determined the difference between the male and female rates. We then took the TRU by gender using only age 20+ and determined the difference between genders using that sample. We then compared the results to see if the gender differentials persisted even in the 16-20 age group in order to evaluate if our age range was a satisfactory comparison to the BLS. As is shown in Figure 5.1 below, the differences in the calculations are almost identical.

Figure 5.1 Verification that Gender Differences Persist Despite Age Range Selection



There is also a reasonable question as to why the BLS reports only on the 20+ age for the gender unemployment rate instead of the usual 16+ civilian population. We do not attempt to answer this question, but we note the difference in the two rates of TRU shown below. We can see from the difference in these lines the higher rate of youth TRU compared to overall TRU for both genders. This is unsurprising, as those under the age of 20 seldom have jobs that meet LISEP's stipulations to be counted as truly employed. This again proves that the 20+ age range and 16+ age range track each other. For women, the average difference for 16+ is 2.73% higher, and in men, the average difference for 16+ is 2.66% higher.

Figure 5.2 True Rate of Unemployment by Sex Using Different Age Ranges for the Labor Force

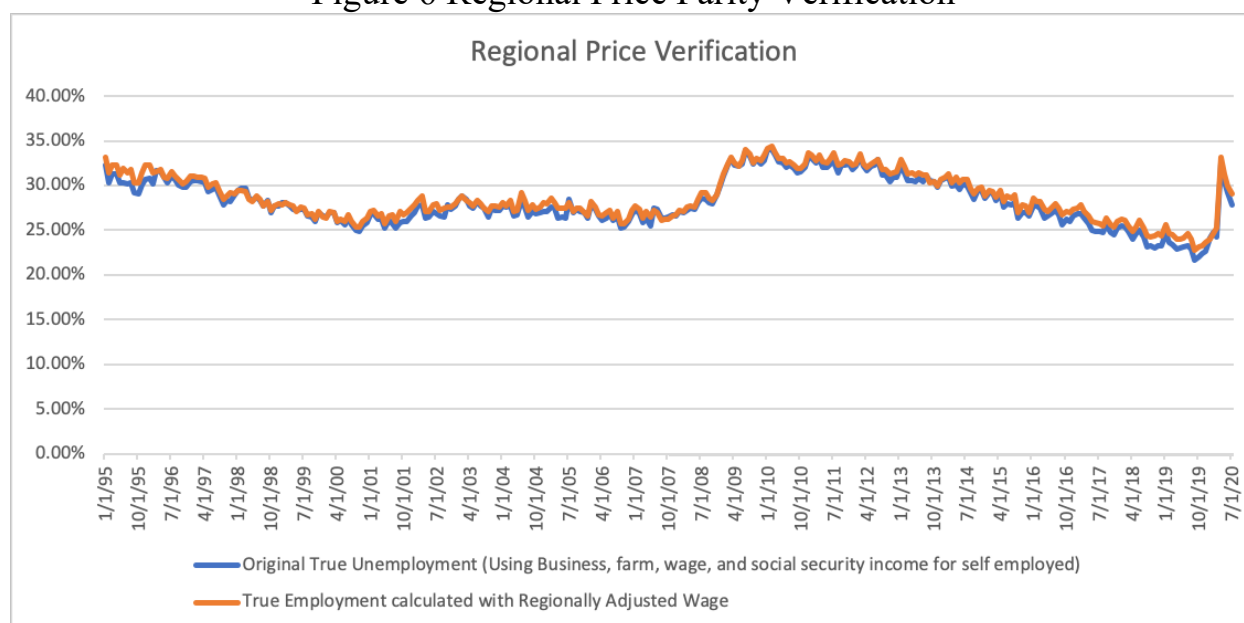


6. Regional Price Variation

This robustness check is based upon the fact that the wages in each area might have different real values. We used the regional price parities published for each state by the U.S. Bureau of Economic Analysis. It “allows comparisons of buying power across the 50 states and the District of Columbia, or from one metro area to another, for a given year. Price levels are expressed as a percentage of the overall national level.”³

However, this is a relatively new area of research, and so the state regional price parities only have values from 2008 to 2018. Because we wanted to confirm the values from 1995 to 2020, the data is adjusted using the same linear trends approach we used for the self-employment calculation. These rates are stable across the 10-year time period so it is not an unreasonable assumption that they would continue to be constant in the years slightly outside the sample. We based the \$20,000 on the entire U.S. January 2020 poverty level (It was not the exact poverty rate, rather it was based off the number, and rounded to be easily understood), and then adjusted the \$20,000 for the equivalent for each state and time period. We then used the same procedure described in Section I. The results are pictured in Figure 6. This comparison shows that the TRU calculated with regional price parities is very similar to that of the main TRU. Thus, it cannot be said that the \$20,000 threshold is invalid because regional discrepancies were not accounted for.

Figure 6 Regional Price Parity Verification



7. Imputed Wage Data from Outgoing Rotation Groups

This robustness check deals with a sampling non-response issue that must be addressed in all large surveys. Namely, when the households are selected, the initial selection is random. So, if

³ <https://www.bea.gov/data/prices-inflation/regional-price-parities-state-and-metro-area>

there is not a response from a randomly selected household, then one cannot ask another household without violating the random selection of the first. Thus, in extreme cases, when the respondent refuses to answer a question or does not know an answer, the CPS imputes the data. Luckily, this is a small part of the data, both currently and historically. The much larger imputation rates of income that are discussed in the literature are primarily about total income, not wage or earner income (Burkhauser, Feng, Jenkins 2012).⁴ See Figure 7 below, taken from the Design and Methodology of the CPS.^[1] This provides a suggestion of the size of the sample that is imputed.

Figure 7 Imputation Rates of the CPS

Table 16-4. **CPS Items With Missing Data (Allocation Rates, %), National Estimates**

Month/year	Household		Demographic		Labor force		I&O		Earnings	
	Range	Mean	Range	Mean	Range	Mean	Range	Mean	Range	Mean
Jan. 1997.....	0.16-5.84	0.41	0.08-9.2	1.35	0.15-12.0	1.57	2.12-5.0	4.09	0.99-32.0	11.02
Jan. 2004.....	0.01-3.42	0.56	0.12-19.0	2.41	0.21-15.0	2.18	5.43-10.0	8.74	0.01-36.0	17.68

Current Population Survey TP66

Quality Indicators of Nonsampling Errors 16-5

U.S. Bureau of Labor Statistics and U.S. Census Bureau

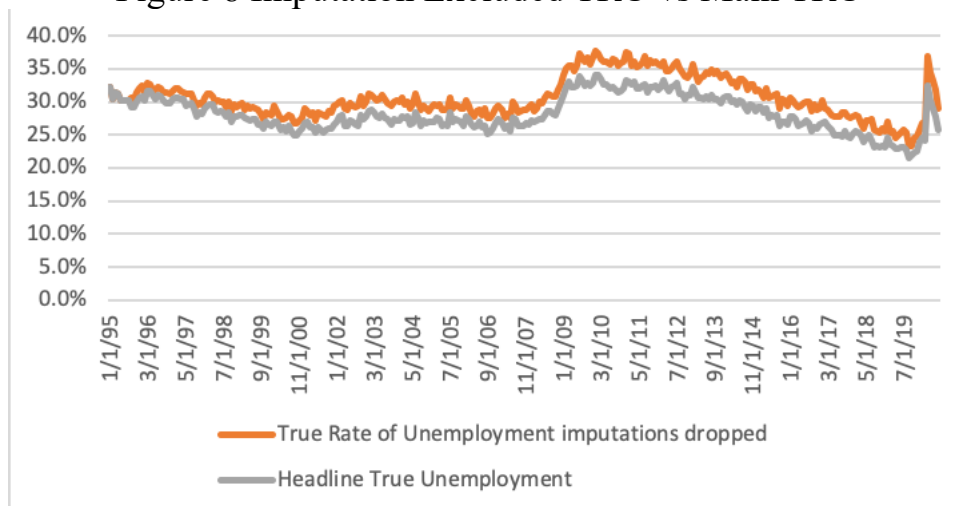
The variable used for earnings allocations by IPUMS is QEARNWEE, which takes on the following values:

CODE	DESCRIPTION
0	Not Allocated
4	Allocated

To check that the imputations did not overestimate TRU, we deleted all of the observations from the dataset that were imputed. Thus, we excluded the observations that took on 4 for QEARNWEE. If the imputations overstated TRU, one would expect TRU without the imputed data to be below LISEP's main TRU. Essentially, this takes the average TRU for each time period of the non-allocated values and then applies this TRU to the allocated values in the same proportion. Figure 8 below shows the results, and it is clear the opposite is true. If anything, the imputations from the CPS understate the main TRU. Thus, we can be sure the imputations do not overstate TRU's magnitude.

⁴ Burkhauser, Richard V., et al. "Recent trends in top income shares in the United States: reconciling estimates from March CPS and IRS tax return data." Review of Economics and Statistics 94.2 (2012): 371-388.

Figure 8 Imputation Excluded TRU vs Main TRU



Furthermore, Bollinger et. al. (2019) suggest that the imputation has a stronger effect with different subpopulations because there is a higher prevalence of allocated values in the tails of the data. Because of this, we calculated the TRU for each of the subgroups using only the non-allocated values, and then applied this rate to their respective racial/ethnic group. This assumes that the allocated group had the same TRU as the non-allocated. The results below show that the largest discrepancies are in the Black labor force, but when you include the imputations in all of the data, the TRU is lower. This confirms that if anything, there is a downward bias in our number, and the actual TRU might be higher if we were privy to no imputations.

Figure 8.1 Imputation for the Black Labor Force

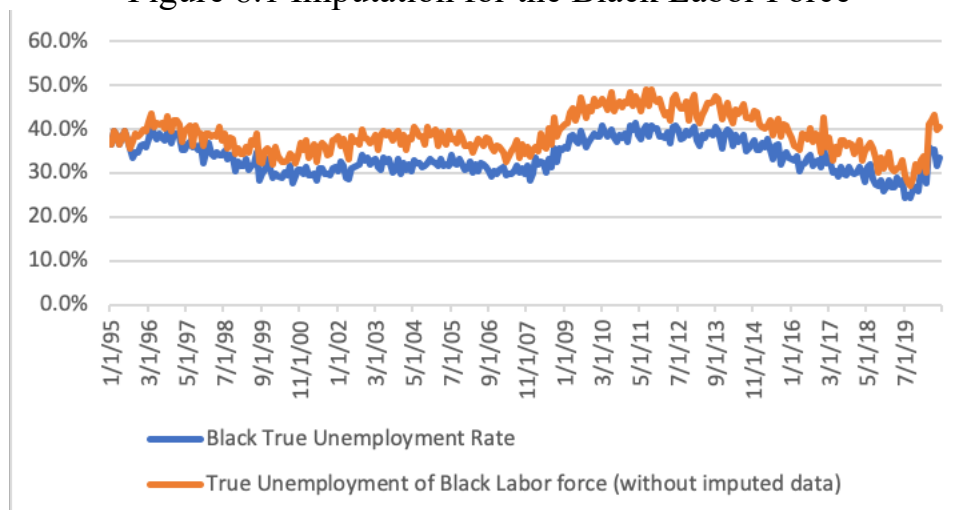


Figure 8.2 Imputation for the White Labor Force

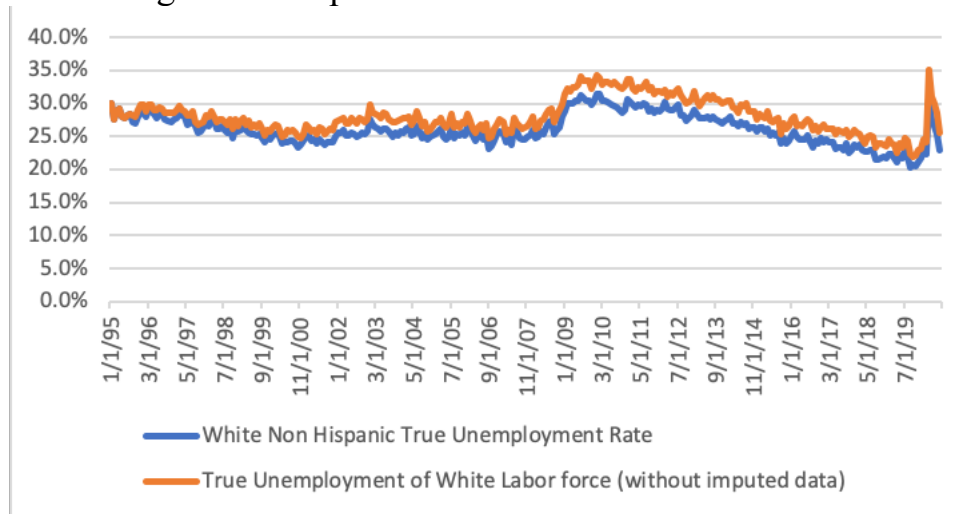
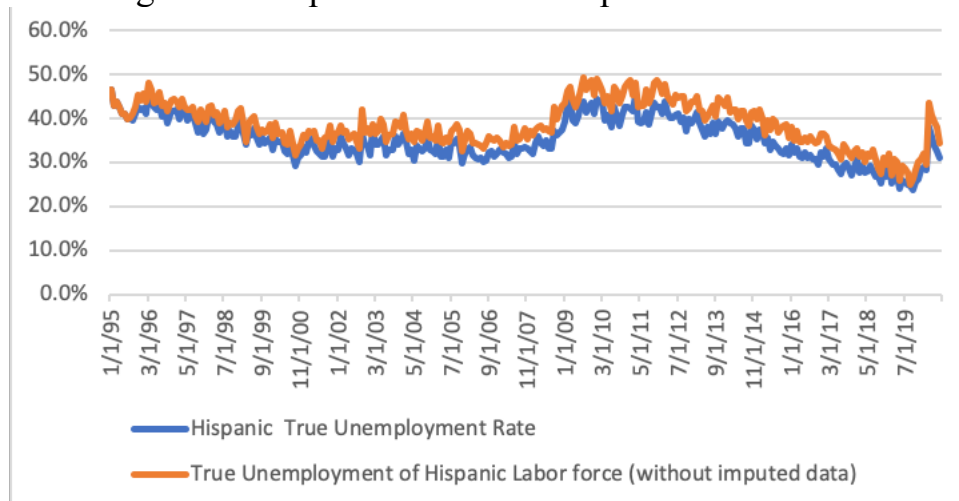


Figure 8.3 Imputation for the Hispanic Labor Force

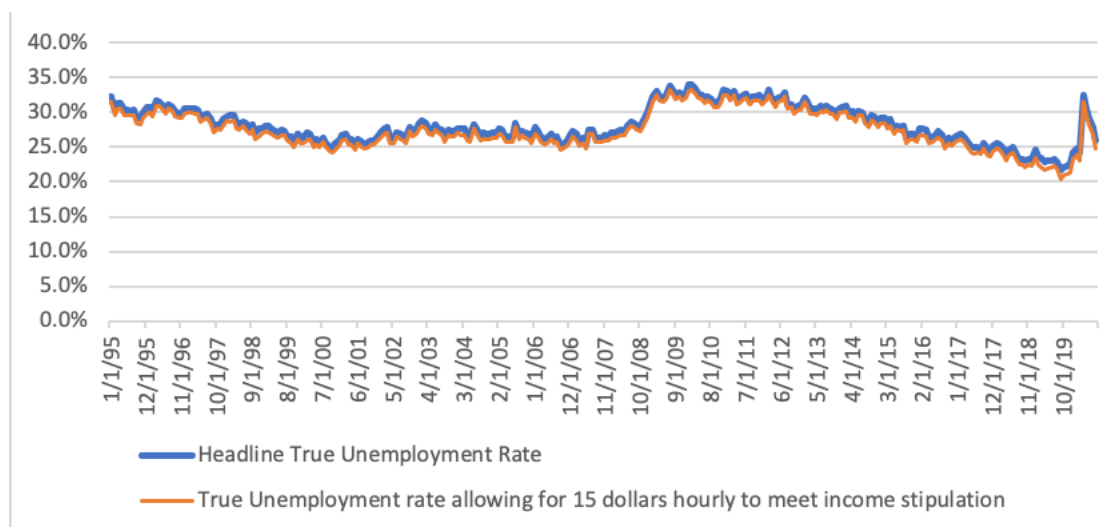


8. Hourly Wages to Substitute for Quality Jobs for Part-time Workers

This robustness check is to ensure that those in part-time jobs by their choice are not making a such a high wage that it allows them to work a limited number of hours. This is referred to as the income effect and would discourage the worker from working more because his/her hourly rate is higher. With the \$20,000 baseline, we know this is probably not the case, yet we test whether those working part-time by choice have higher wages but just do not work enough to allow them to make the \$20,000 income level to be counted as truly employed. In this check, we set the hourly wage bar to a mere \$15 an hour, and thus took the most conservative estimate as to where the income effect might decrease willingness to work. We recalculated the TRU two ways to meet the income stipulation to: 1) have a wage of \$20,000 per year (the original way) and 2)

have an hourly wage of \$15 per hour (in January 2020 dollars). The new rate trend is shown below. The differences between the rate throughout time averages at 0.74%. Thus, even with the low wage of \$15 per hour, the population that makes this wage and chooses to work less hours do not account for a significant part of the TRU.

Figure 9: TRU Comparison for Different Methods of Achieving Income Stipulation



V Seasonal Adjustment

For seasonal adjustment methods, we used the BLS's X13 ARIMA SEATS model to adjust each rate. We downloaded this program from the U.S. Census website and used the X-13 ARIMA adjustment, allowing it to select its own ARIMA model for each of the individual data series. The only exception to this was the population numbers for those with a bachelor's degree. The series covariance matrix produced by ARMA was singular, and thus the t-statistics for the parameters could not be computed. For this, we used the X11 adjustment method. We further specified that this is a stock series in the adjustment.

Chapter 2

Methodology for the LISEP True Usual Earnings

Written by Research Assistant Philip Cornell on behalf of the
Ludwig Institute for Shared Economic Prosperity

This document details the calculation and verification of the True Earnings.

I Definition and Comparability of LISEP Usual Weekly Earnings Measures

LISEP's definition of True Median Earnings adheres to the methodology used to calculate the median earnings measure in the Usual Weekly Earnings of Wage and Salary Workers quarterly report published by the Bureau of Labor Statistics (BLS).⁵ The measure of LISEP's True Median Earnings, however, differs in the sample used to calculate the statistic. The BLS includes only full-time wage earners whereas LISEP includes full-time and part-time wage earners as well as those actively searching for a job. The LISEP measure thus provides a more comprehensive measure of the U.S. workforce's "usual earnings."

In this we seek to answer the question of "What are the median earnings of those who are working and actively seeking work?" The answer to this question will fluctuate based upon two things. First, what are the earnings of those who are employed? Second, what is the fate for those seeking jobs? To indicate a healthy labor market, the former question must be answered with "high wages", and the latter question must be answered with "employment". The BLS fails to address the second question at all in their earnings metric.

First, we will outline the BLS sample and method and then differentiate the LISEP metric from that of the BLS.

A. BLS Metric of Median Weekly Earnings

⁵ See Appendix for other measures of earnings published by the BLS and explanations of why we don't compare our metric to those.

Although there are many different earnings measurements that government provides, we compare ours to the one produced from the Current Population Survey (CPS). We detail the other measures and reasons for our choice in the Appendix. The BLS publishes a quarterly report entitled the “USUAL WEEKLY EARNINGS OF WAGE AND SALARY WORKERS” that details a variety of earnings statistics calculated from the CPS. BLS uses three different types of calculations in the report.

The first, and most dominant statistic is the Median Weekly Earnings of full-time wage and salary workers, and the following are key definitions used:

- “Median Earnings- The median earnings level represents the midpoint in an earnings distribution, with half of workers having earnings above the median and the other half having earnings below the median.”⁶
- “Full time workers- For the purpose of producing estimates of earnings, workers who usually work 35 hours or more per week at their sole or principal job are defined as working full time.”⁶
- “Wage and salary workers- These are workers who receive wages, salaries, commissions, tips, payment in kind, or piece rates. The group includes employees in both the private and public sectors but, for the purposes of the earnings series, it excludes all self-employed persons, both those with incorporated businesses and those with unincorporated businesses.”⁶

The BLS reports the Usual Weekly Earnings statistic for full-time and salary workers for different strata of workers: by race/ethnicity, sex, age, educational attainment level, and in different occupations.

The second BLS statistic better illustrates the earnings distribution – it reports the earnings of full-time wage and salaried workers in the first decile, first quartile, third quartile, and ninth decile. These statistics are also reported by sex, race/ethnicity, and educational attainment.

The third statistic reported is the median earnings of part-time workers, with full-time workers excluded. It is reported by sex, age, and race/ethnicity.

Despite the three types of statistics, the “USUAL WEEKLY EARNINGS OF WAGE AND SALARY WORKERS” report is heavily skewed toward the first: median earnings of full-time wage and salary workers. Four out of six tables are about this metric.

B. LISEP True Earnings Measurement

⁶ <https://www.bls.gov/news.release/pdf/wkyeng.pdf>. The 2020 third-quarter release of the “USUAL WEEKLY EARNINGS OF WAGE AND SALARY WORKERS.” More information about these definitions can be found at <https://www.bls.gov/cps/definitions.htm#earningswagesalary>.

LISEP calculates our own form of the first two types of BLS statistics: median earnings of wage and salary workers and the selected income deciles and quartiles of earnings. But, in our version of these variables, we strive to present a better measure of the workforce’s usual weekly earnings by presenting the median earnings of all workers and job seekers. Thus, we included in our sample all members of the labor force that are potential wage and salary workers: full- and part-time wage and salary workers, as well as the unemployed. We did not include in our sample those outside the labor force, which is defined by the BLS this way:

- “The labor force includes all people age 16 and older who are classified as either employed and unemployed, as defined below. Conceptually, the labor force level is the number of people who are either working or actively looking for work.”²

We excluded those who are currently self-employed from the sample. This is for two reasons. First, the BLS does not collect monthly earnings data for those who are primarily self-employed. Second, the self-employed have an earnings structure that allows them to shift company profits to earnings or to reinvestment. Thus, it is hard to measure actual earnings per month if they are able to shift this income around.

We used the BLS definitions listed in Section A above for full-time wage and salary workers (with part-time being less than 35 hours per week). We also used the BLS classification for unemployment, described below.

- “In the Current Population Survey, people are classified as unemployed if they meet **all** of the following criteria:
 - They were not employed during the survey reference week.
 - They were available for work during the survey reference week, except for temporary illness.
 - They made at least one specific, active effort to find a job during the 4-week period ending with the survey reference week (see active job search methods) OR they were temporarily laid off and expecting to be recalled to their job.”⁷

Our first statistic is the median earnings of the entire salary and wage labor force. We also stratified this measurement by race/ethnicity, age, sex, and educational attainment. See Section II C for more information of the sample’s breakdown.

Our second statistic mirrors the BLS in that we reported the change in earnings from different points of the wage and salary earnings distribution. Namely, we reported the first decile, the first quartile, the third quartile and the ninth decile. We then reported this metric by gender and race/ethnicity.

Table 1 presents a summary of the BLS and LISEP samples.

Table 1 Characteristics of Usual Median Earnings of the Workforce Samples

⁷ <https://www.bls.gov/cps/definitions.htm#wagesalary>

	Bureau of Labor Statistics Sample of Headline Number	LISEP Sample of Headline Number
Full-time Wage and Salary Workers	Included	Included
Part-time Wage and Salary Workers	Not Included	Included
Unemployed	Not Included	Included
Incorporated Self-Employed	Not Included	Not Included
Non-Incorporated Self Employed	Not Included	Not Included
Unpaid Family Workers	Not Included	Not Included
Out of the Labor Force	Not Included	Not Included
Not Members of the Civilian Non-Institutional Population	Not Included	Not Included

II. Calculation of the Statistic

In this section, we first detail the procedure used to generate the sample. We then detail an example of the mathematical procedure to calculate the median of the sample.

A. Sorting Procedure for the Sample

To filter the data to only include those in the workforce, we used the variable LABFORCE.⁸ The other condition was that the individual could not be self-employed, and we used the CLASSWKR variable to make this distinction.⁹ We also used CLASSWKR to exclude unpaid family workers from the sample set. The BLS also does this for similar reasons as for the self-employed. Unpaid family workers often contribute to a small business, and their wages might come in the form of being able to live in the business or eat the food. If they don't make direct wages, this might not be an indication that they have no monetary earnings.

Lastly, we included those who were categorized as self-employed, but currently were unemployed. Because of a coding in the data, if a respondent's last job was self-employment, he/she might respond that they are in the self-employed class of worker despite being currently unemployed. So, because employment is not a necessary condition for self-employment, we also wanted to keep those who were unemployed, even if they categorized themselves in the class of the self-employed because it was their last status when actively working. To do this, we used a combination of the CLASSWKR and WKSTAT¹⁰ variables.

B. Mathematical Calculation of the Statistics

⁸ See thorough definition on page 1

⁹ See thorough definition and coding on page 2

¹⁰ See thorough definition and coding on page 3

To calculate median earnings given a sample, we used the same binning method used by the BLS in its calculation of median wages. Specifically, we grouped each weekly earnings measure into \$50 bins. We then determined which bin contains the median weighted wage. Then within this bin, we linearly interpolated across the endpoints of the bin, based on the weights of the rest of the sample.

These weights are used for the Outgoing Rotational Group (ORG) survey, which are four times the weights used for the regular CPS. This is because the ORG sample is just the fourth and the eighth month of the eight-month sampling done by the CPS and is one fourth of the households. This variable on IPUMS is EARNWT and signifies the number of persons in the civilian non institutional adult population that are represented by that specific entry.¹¹ The variable used to define the earnings of an individual is EARNWEEK, which is defined by IPUMS as:

- “EARNWEEK reports how much the respondent usually earned per week at their current job, before deductions. Interviewers asked directly about total weekly earnings and also collected information about the usual number of hours worked per week and the hourly rate of pay at the current job.”¹²

EARNWEEK is top coded to ensure anonymity of the data. Because of this, we cannot calculate exact values for some of the top earners, and thus finding the mean is impossible with the micro-survey data that is available to the public. But the top code does not affect the median because all of the top coded values do not lose their place above or below the median.

We used the BLS’s bin conversion method equation. Taken from its methodology, “For Usual Weekly Earnings (UWE) the starting point is \$24.50, and the bin size is \$50. The starting point is algebraically transformed to 32 with a starting bin size of one half.” The bin is then rounded down to the 0.5 interval.¹³ The reason for this binning of the data and then unbinning is to prevent modal stickiness in the measure of the median. Many people report round numbers in their measures, and thus although the entire distribution might shift up or down, these clusters prevent the median from moving. The binning method helps show these smaller changes.

In Example 1, we illustrate, from start to finish, finding the median of a sample with the binning method.

Example 1

Part A) Find the bin containing the median value

- 1) Sort the earnings from smallest to largest. Notate the observations such that

$$e_1 \leq e_2 \leq \dots \leq e_n$$

- 2) Calculate the running sum of weights of the sorted variables g such that

¹¹ https://cps.ipums.org/cps-action/variables/EARNWT#description_section

¹² https://cps.ipums.org/cps-action/variables/EARNWEEK#description_section

¹³ This was not found on the BLS website, and rather was obtained by a BLS Economist in the Earnings section.

$$g_i = \sum_{i=0}^i w_i$$

- 3) Find the weighted median of earnings m by finding the earnings observation whose running sum of weights is half of the total sum. Thus:

$$g_n / 2 = \sum_{i=0}^m w_i = g_m$$

- 4) Find the bin b_m of the median earnings value:

$$b_m = \text{int}(((g_m/100) + 31.755) * 2) / 2$$

Part B) Linearly interpolate over the bin and convert back to dollars to find the reported median

- 5) Find the relative location of the m within b_m . First, find the sum of weights of all of the bins up to b_m . Then find the difference between this sum and the sum of weights all the way to the median:

$$\begin{aligned} \delta &= g_m - g_h \\ &\text{such that} \\ h &= \max(n) \text{ while } h \in b_{b_m-1} \end{aligned}$$

- 6) Proportion this difference to the size of b_m

$$\begin{aligned} \phi &= \delta / (g_f - g_{h+1}) \\ &\text{s.t.} \\ f &= \max(n) \text{ while } f \in b_m \end{aligned}$$

- 7) Halve this to get the median weight, and then add it to the starting point of the bin (note that bins start on integers and half integers):

$$\text{Final binned median} = \phi * 1/2 + \min(b_m)$$

- 8) Convert this binned value back to dollars to get the final reported median:

$$\text{Final reported median} = (\text{Final binned median} - 31.775) * 100$$

This is the exact procedure the BLS conducts, and in an effort at a direct comparison, we follow the same method.

We used the same method for the calculation of the selected deciles and quartiles of the earnings distribution, save that we did not calculate the median; instead, we calculated the specific percentile. Thus, in step three, we found the percentile by dividing the sum of weights by the appropriate percentile that we are aiming to find. For example, if we were to try to find the first quartile, we would multiply the sum by 1/4 and find the 25th percentile at that place in the sorted weights. To find the 90th percentile, we multiplied the sum by 9/10ths and found the value at that place in the sorted weights, etc. In addition, in step seven, the final binned percentile would be found by taking $\phi * (\text{percentile}/100)$ and then adding that to the 2nd term.

C. Stratification

We calculated the median earnings by categories of racial and ethnic groups, sex, and educational attainment.

For the breakdown by race/ethnicity, we publish the Usual Weekly Earnings of the Black Workforce, the Usual Weekly Earnings of the White Workforce, and the Usual Weekly Earnings of the Hispanic Workforce. To sort the observations, we used the RACE variable and the HISPAN variable from the IPUMS-CPS.¹⁴ It is important to note that we categorized observations as Black if they indicated being part Black. This is not true for observations categorized as White. Thus, if an observation indicated that they were partly white as one of a combination of races, we did not categorize them as White. Furthermore, we categorized the White population to specifically be White non-Hispanic. Thus, we used the HISPAN variable both to identify those in the population that responded that they belong to Hispanic ethnicity, and also to exclude those who are Hispanic from what we refer to as the White workforce. We used the same sorting procedure as detailed on pages 9-12.

For the breakdown by sex, we publish the Usual Weekly Earnings of the Male Workforce as well as the Usual Weekly Earnings of the Female Workforce. For this we used the variable SEX. We used the same classification procedure that we used in TRU.

For the breakdown by education, we used five exclusive groups that are characterized by the respondent's furthest level of educational attainment. These groups are less than high-school degree, high-school diploma or equivalent, some college/associate's degree/vocational degree, bachelor's degree, advanced/professional degree. We used the variable EDU¹⁵ to determine which group to sort the respondents into. This is the same sorting procedure that we used to classify the educational groups on page 13 for the TRU statistic. Unfortunately, the granularity provided by the IPUMS-CPS did not allow us, prior to 1992, to generate all five of the different educational classifications. Thus, for this stratification, we report the data only from 1992 to present.

III. Robustness of Results

In this section, we first detail the Standard Error Calculation of the results. We provide statistics and illustrations of the standard errors that suggest validation of the accuracy and consistency of the LISEP measurements. We then conduct several robustness checks of the results. In these checks, we both vary the sample and the method slightly to suggest that our results depict an accurate measure of the Usual Weekly Earnings.

A. Standard Error Calculation

We wanted to assure that our measure was reasonably precise, and so we calculated standard errors for each measure. The method for this was calculating the bootstrap standard error at 95% confidence for 100 repetitions. Reported in Table 2 are the average standard errors, the

¹⁴ Detailed code and description found on page 9

¹⁵ Detailed code and description found on page 13

maximum 95% standard error for a single quarter for each of the statistics that we compiled, and the standard deviations of the standard errors.

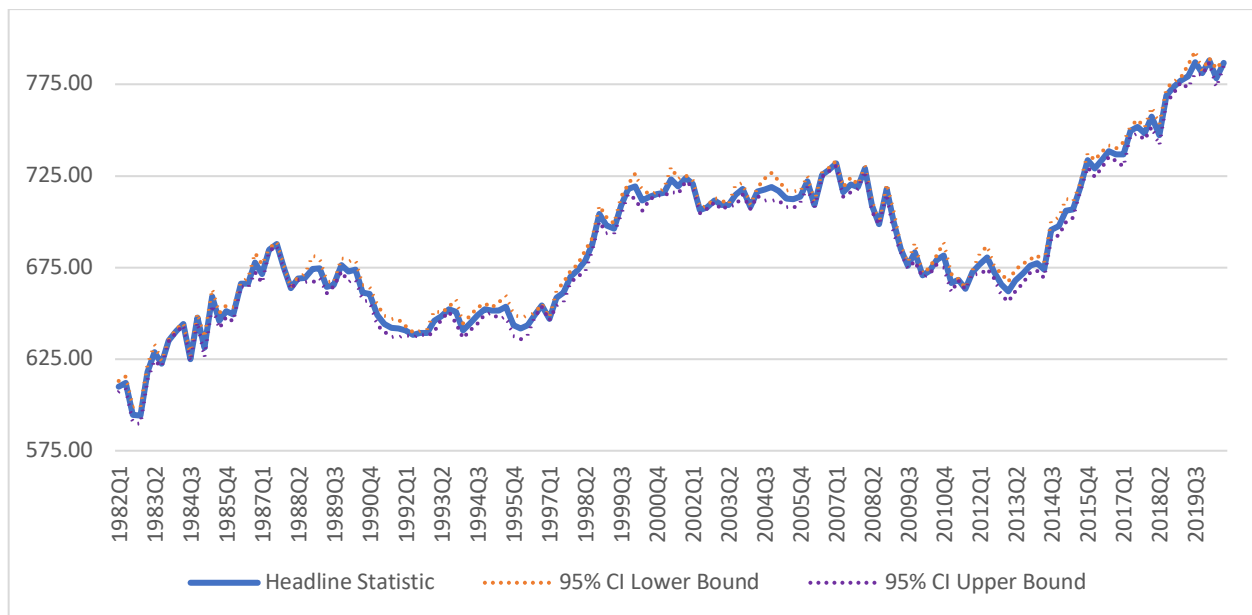
Table 2 Standard Errors and Related Metrics of LISEP Usual Earnings

	Average SE (95%)	Maximum Standard Error (95%)	Year and Quarter of Maximum	Standard Deviation of Standard Error
Headline LISEP True Earnings	3.43	7.57	2004 Q4	1.90
True Earnings for Black Workforce	8.15	13.87	2002 Q4	2.66
True Earnings for White Workforce	4.55	8.73	1986 Q1	1.97
True Earnings for Hispanic Workforce	7.83	16.81	1985 Q2	2.91
True Earnings for Male Workforce	5.62	12.82	1983 Q3	2.91
True Earnings for Female Workforce	3.99	7.54	1985 Q2	1.87
True Earnings for Workforce with No HS Diploma	6.47	12.86	2015 Q4	2.54
True Earnings for Workforce with Max Education of HS Diploma	4.20	7.92	2000 Q2	1.88
True Earnings for Workforce with Max Education of some College	10.11	17.04	1993 Q3	2.35
True Earnings for Workforce with	9.78	16.34	1993 Q2	3.70

Max Education of bachelor's degree				
True Earnings for Workforce with Advanced Degree	17.21	27.44	1996 Q3	5.22

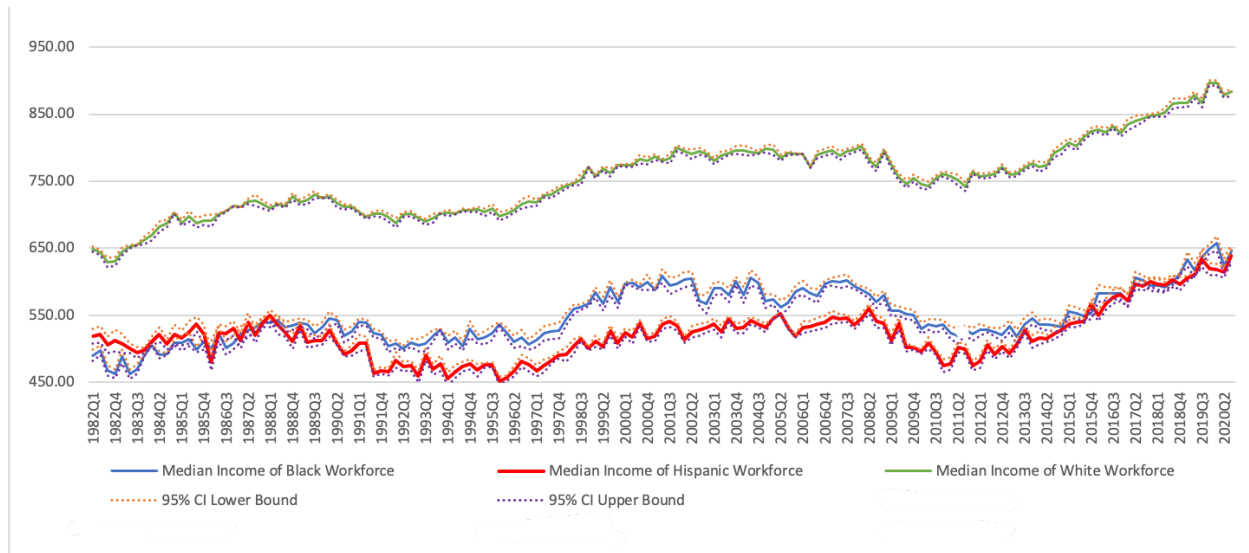
In figures 1.1-1.4, we graphed median wages as well as the 95% confidence interval for the headline statistic, median earnings by race/ethnicity, median earnings by sex, and median earnings by educational attainment.

Figure 1.1 LISEP Median Earnings Overall with Margin of Error



We see above the accuracy of the LISEP median earnings. The confidence interval closely tracks the reported point estimate.

Figure 1.2 LISEP Median Earnings by Race/Ethnicity with Margin of Error



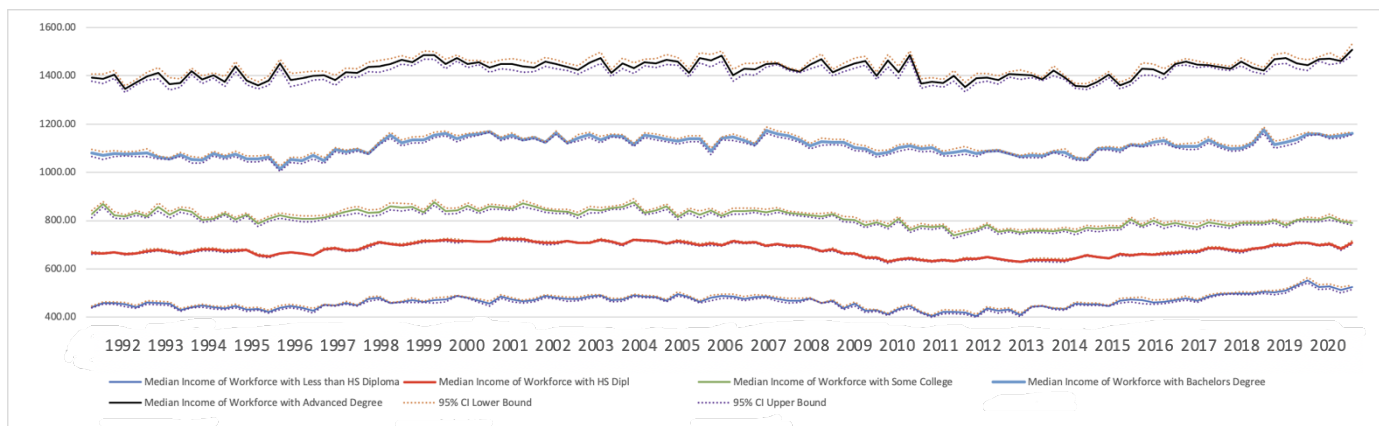
We see above that the White labor force has a statistically significant different median weekly earnings level than that of the Hispanic or the Black labor force.

Figure 1.3 LISEP Median Earnings by Sex with Margin of Error



At all points above, we can see that the male median earnings statistic is significantly different than the females because neither confidence interval intersects.

Figure 1.4 LISEP Median Earnings by Educational Attainment with Margin of Error



Each education level above is significantly different, as neither confidence interval intersects with another educational level's confidence interval.

B. Robustness Checks of Sample

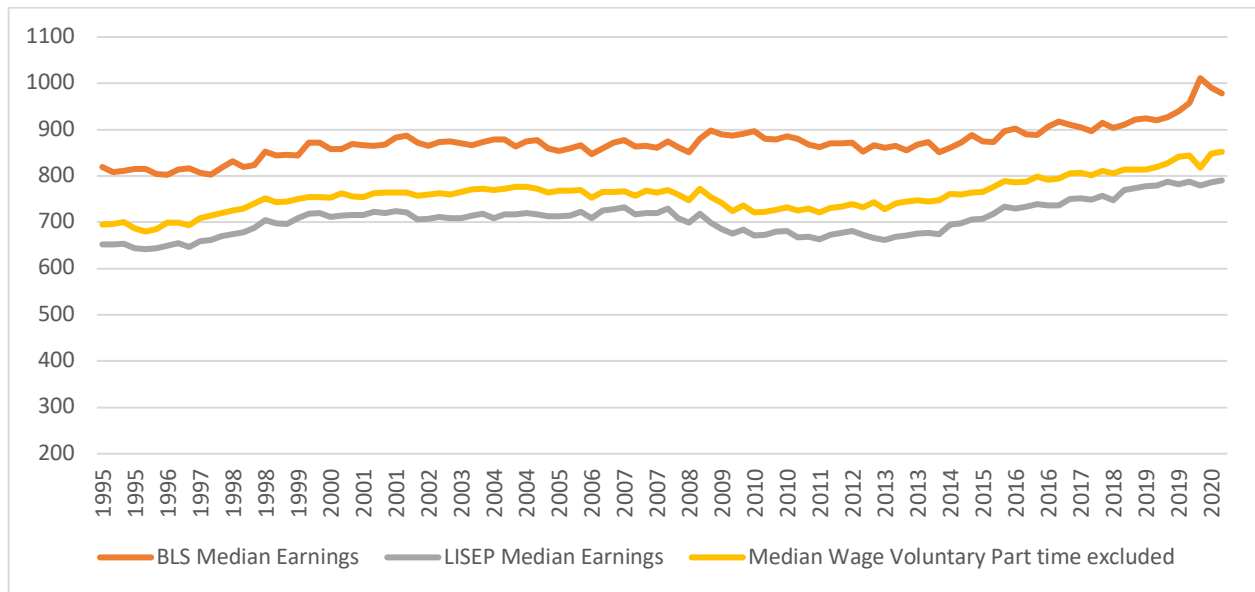
In this section, we detail several robustness checks to test LISEP's accuracy regarding the median worker's earnings in the labor force.

1. Voluntary Part-time Robustness

First, we gauged the importance of those who are part-time employed voluntarily and may be satisfied with a lower wage. If this is a substantial part of the labor force, then we might expect that the potential workforce earnings are much higher. Furthermore, we do not want to inaccurately portray the workforce as suffering if lower earnings are by choice. We tested this in two different ways. First, we excluded part-time voluntary workers from the workforce. To do this, we used the WHYPTLWK¹⁶ variable and use the same categorization as detailed in pages 4-5 to determine the voluntary aspect of this part-time status. In Figure 2.1, we show the results of this measure compared to the results of the LISEP earnings statistic.

Figure 2.1 Comparison of LISEP Statistic with LISEP Modified Part-time Excluded Sample

¹⁶ For detailed coding and description see pages 4-5



Two things are apparent in the graph above. The first is that, throughout time, the BLS statistic is greater than the LISEP statistic with the modified LISEP sample. Furthermore, the modified LISEP sample moves similarly to the LISEP sample, not the BLS sample. This suggests that the indicator change of the LISEP earnings measurement portrays similar information to the modified labor force, both of which differ from the BLS’s headline statistic.

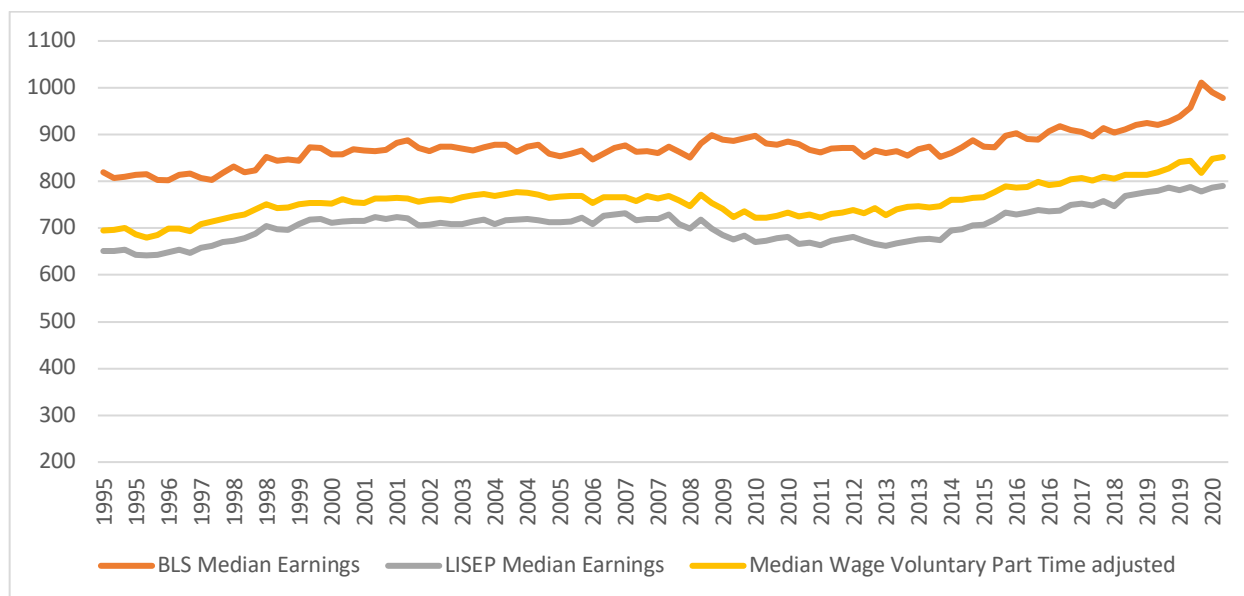
The second robustness check modified part-time earnings so that they are earning the weekly equivalent of a full-time worker’s hours, but at their current hourly rate. Thus, for those who are voluntarily employed part-time, we adjusted their wages upward. For example, take someone who voluntarily works part-time for 20 hours a week with weekly earnings of \$400. We adjusted this person’s earnings to be \$20 an hour for 35 hours a week (full-time status), and their adjusted earnings equal \$700 a week. To determine the number of hours that a person works in a given week, we used the UHRSWORKORG variable from the IPUMS-CPS, which is described as:

“UHRSWORKORG reports the total number of hours the respondent who is paid hourly usually works per week at their main job”¹⁷

We then took their earnings per week from EARNWEEK and adjusted it to the proportional time that the part-time worker works of the 35-hour, full-time week. For the full-time workers, we took their given EARNWEEK values. We also took the given values for those who are involuntarily part time (I do not adjust them). Figure 2.2 shows this compared to the LISEP main statistic.

Figure 2.2 Comparison of LISEP Main Statistic with
LISEP Modified Part-time Adjusted Sample

¹⁷ https://cps.ipums.org/cps-action/variables/UHRSWORKORG#description_section



This graph shows the same findings as figure 2.1 – the adjusted number moves similarly to the LISEP statistic and strays away from the BLS statistic, especially in times of crisis. This further validates the use of the LISEP statistic to accurately depict trends in the usual earnings of the workforce.¹⁸

Finally, we present an argument for the reason we did not use these modified samples as our main statistic. First, there is no clear indication that the employers who are hiring part-time would be able to hire these workers full-time at the same rate. The employer also may be hiring someone part-time and then hiring someone else at higher weekly earnings full-time. Thus, the full-time employee might be above the median earnings level and the voluntary part-time below that. If both of them chose to be part time, this could push them to equally split the aggregate pay and both end up below. Thus, it is overly optimistic to claim that the employers of voluntary part-time workers would be able to hire them at the same rate, but full time. Second, although there is merit to both measures, the main measure makes the least assumptions about the data. We don't cherry-pick the selected sample; we are including all of the labor force that data permits (again, self-employed are excluded because we do not have accurate monthly earnings data).

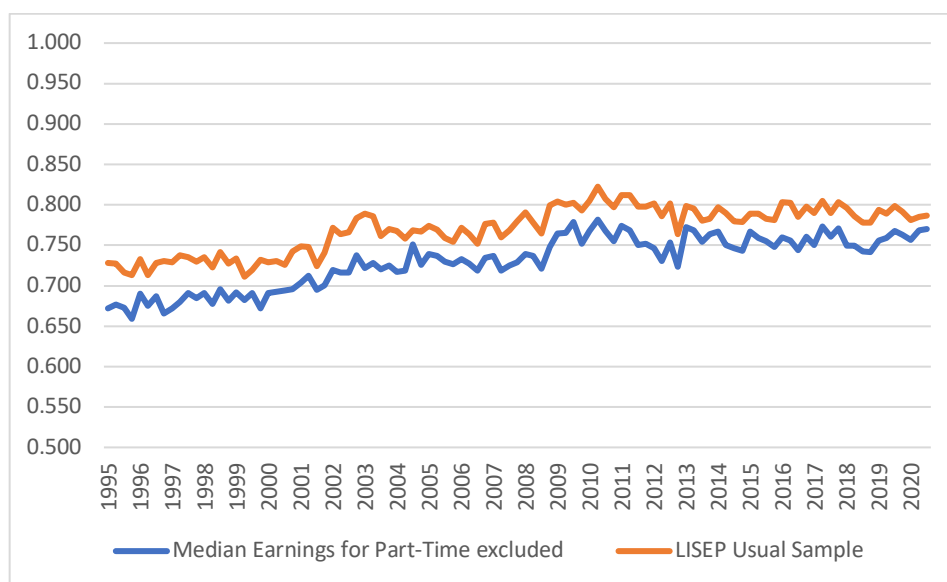
2. Robustness Check for Female to Male Earnings Ratio

The third robustness check shows there is negligible difference in the part-time voluntary inclusion in the sample between sexes. One might make the argument that including part-time voluntary workers in the sample lowers the female median earnings disproportionately more than

¹⁸ Figure 2.1 and 2.2 show data from 1995 onwards instead of 1982 onwards because the CPS before 1995 did not ask about reasons for part time employment. Thus, prior to 1995, it was impossible to differentiate between voluntary and involuntary part time employment.

it lowers the male earnings.¹⁹ This argument would say we falsely portray the true ratio of female to male earnings. So, to check whether the true difference in female-to-male earnings is due to voluntary part-time employment, we used the modified sample from the previous robustness check where we excluded all of those who were voluntarily part-time and calculated the ratio of male-to-female median earnings in this fashion. We graphed this ratio in Figure 3 in comparison to the ratio between the LISEP female-to-male true earnings ratio.

Figure 3 Ratio of Female-to-Male Earnings of Full LISEP Sample Compared to Modified LISEP Sample



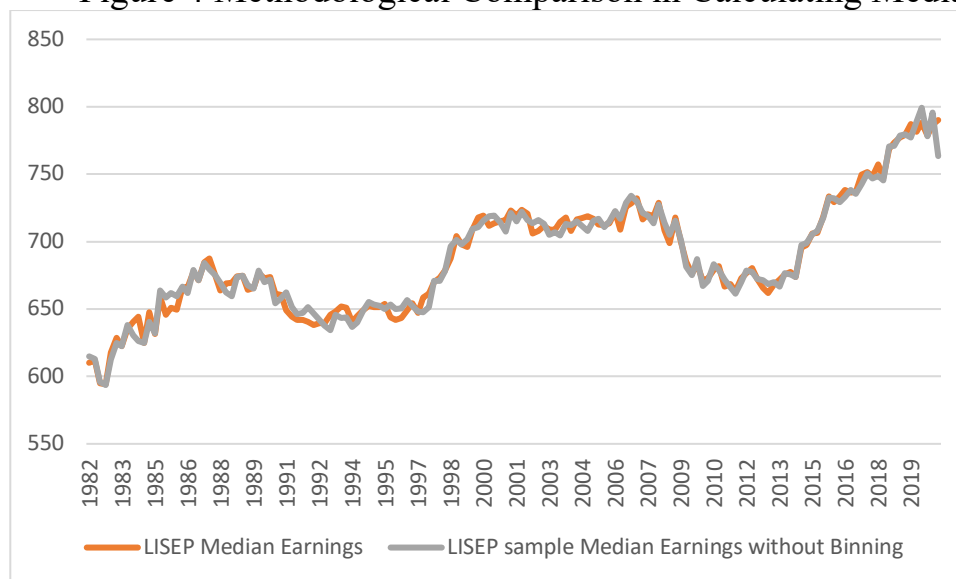
¹⁹ This is because many believe that females predominantly work part-time voluntarily because they might have childcare or other similar issues. We do not suggest this is the case but provide the argument for context. And this criticism does not affect our statistic.

The picture above shows that even with the exclusion of voluntary, part-time workers, we still can see that there is a large gap between male and female earnings. Furthermore, in 25 years, this gap shows only very small improvements. To clarify, if women and men earned the same earnings, this ratio would be at 1. Furthermore, we validated that the LISEP full sample median earnings and the LISEP modified sample median earnings are very correlated because these two ratios have 0.951 correlation since 1995.

C. Robustness Checks of Method

In the fourth robustness check, we used a different methodological approach in finding the median than the BLS binning method. In this approach, we took the weighted median for each quarter and year. We then plot this median versus the BLS median to compare below.²⁰

Figure 4 Methodological Comparison in Calculating Medians



As we can see from the graph above, there is very little difference in the BLS methodology of calculating the median (binning) and a regular weighted median methodology. We employed the BLS method in order to make a direct “apples-to-apples” comparison but do not to weigh in on the validity of the BLS method. Rather, we show this to prove that the BLS method of binning does not show a completely different result than an easier-to-conceptualize median of the data.

IV Inflationary and Seasonal Adjustment

²⁰ We used the same method as detailed in the II.B. Mathematical Calculation Section, but only do the first three steps and stop after we found the median weighted value.

In this section, we provide a brief summary of the inflationary and seasonal adjustment that we apply to the earnings series generated by LISEP.

A. Inflation Adjustment

In order to more easily compare numbers across time, we inflation adjusted all of them. The BLS produces an inflation-adjusted measure as well. But our inflation measure is based off of the 2020 Q1 dollar whereas the BLS number is based off of the 1982-1984 constant dollar. When looking at time trends and changes from quarter to quarter, the specific dollar that is chosen is unimportant. We used the 2020 dollar so that a common viewer is better able to understand the actual spending power of the listed wage.

We used the St. Louis Federal Reserve Economic Data time series of Consumer Price Index for All Urban Consumers: All Items in U.S. City Average (CPIAUCSL)²¹ to adjust the earnings.

Like the binning method of calculating the median, we did not use the CPI out of endorsement. Rather, we used it so that we can have a more straightforward comparison to the statistics BLS produces. We recognize that for parts of the population, the CPI does not accurately reflect their change in spending power. Furthermore, the bundle of goods contained in the CPI is gathered from real spending by the Personal Consumption Expenditures, and thus reflects what people spend their money on currently. This is not the same as what people would spend their money on if given the resources. Thus, CPI might go up marginally because the goods that people used to buy are outside their budget (e.g., education and medical care), not because their wanted bundle's price only goes up marginally. Thus, we entertain the possibility that the spending power portrayed by the CPI-adjusted earnings we present actually overstates the growth in spending power of the population.

B. Seasonal Adjustment

We used the same seasonal-adjustment software that we used for the TRU calculation. This is the official software used by the U.S. Census Bureau: X-13 ARIMA SEATS. In an effort at transparency, we publish both the seasonally adjusted as well as non-seasonally adjusted wages. Furthermore, we seasonally adjusted the wages after inflation adjusting the wages. The reported data on the website, www.lisep.org, are not seasonally adjusted but the seasonally adjusted measure is included in the accompanying raw data spreadsheet.

V Appendix

In this appendix, we provide a brief summary of the other measures of wages and earnings the BLS presents apart from those published using the CPS. We also provide brief reasoning as to why we chose not to calculate our statistic in comparison to these other measures.

²¹ <https://fred.stlouisfed.org/series/CPIAUCSL>

A. Current Employment Statistics Survey

The CES is a survey of establishments that is conducted by the BLS on a monthly basis. This is the only other measure of employment (besides the CPS, which we use) that BLS publishes monthly. Unlike the CPS, the CES is an establishment survey, meaning that instead of asking households the questions, the surveyors go straight to the establishments. The CES uses “a representative sample of businesses” to give an accurate measure of nonfarm employment and payroll statistics.²² The CES sample is about twice the size of the CPS sample (145,000 workplaces compared to 60,000 households) and includes establishments representative of businesses and government agencies (not including farms).

A strength of this data is that it is more exact because it is asking establishments to look exactly at their payroll. Moreover, in a statistical sense, the point estimates provided by the data are more precise because of the larger sample size. Additionally, the establishment survey includes part- and full-time work in its estimate. Thus, it has a wider sample than the headline statistic used in the CPS usual weekly earnings.

The first weakness is that the data excludes farm businesses. Second, the establishment survey “cannot include new firms immediately; they are incorporated with a lag. Similarly, the permanent closure of a firm is not always captured immediately.”²³ For this reason, the ability for this survey to act as a monthly indicator is largely diminished. Third, the CES estimates are averages and not medians. So, this is easily skewed by much higher earnings. For these three reasons, the measure is not a revealing enough tool to gauge workforce earnings. Finally, like the CPS data, it does not include self-employed.

For the reasons above, we chose not to make our measure as a comparison to the CES data. We would have had to use a similar methodology, so we would have had to take the wage average. Averages are not a good way to represent the “usual” worker because they are too easily skewed by the high earners. In addition, the CES microdata is not available to the public. Thus, the available refinement for the statistic is limited.

B. Metrics from the National Compensation Survey

The Employment Cost index is based off of data collected by the National Compensation Survey (NCS). The NCS samples establishments once a quarter (March, June, September, and December). The sample is across the nation and across industries and occupations. But left out of the sample are federal government agencies, only state and local governments are included.²⁴ The NCS then publishes estimates of employment and employer cost for employee compensation. The first metric is used to measure the change in compensation over time while the latter is a measure of the combined change in compensation and employment seen from the

²² https://www.bls.gov/web/empsit/ces_cps_trends.htm

²³ https://www.bls.gov/web/empsit/ces_cps_trends.htm

²⁴ <https://www.bls.gov/opub/hom/ncs/concepts.htm>

employer perspective²⁵ (productivity is a useful measure published by the NCS). The NCS also provides modeled wage estimates of hourly pay by occupation.

The first advantage to this data is that it measures full compensation. This is useful because it shows the different benefits that workers receive from employers beyond just wages. It also shows a wide variety of granular data, making sure to sample all of the occupations and regions of interest. Lastly, the data shows all private industries (including farms).

Yet, there are many disadvantages to the data in using them as earnings estimates. The first is that compensation and earnings are not entirely the same thing. If healthcare costs go up, compensation also goes up, but the worker is not really seeing a beneficial change in earnings if he/she is still receiving the same level of medical care and same wages. Second, the compensation data provided are quarterly estimates, not a grouping of monthly estimates. Thus, if there is a decline in earnings for January and February and then a recovery in March, there will be no indication of this change, and the microdata of the survey is inadequate to explain month-to-month variation. Next, the modeled wage estimates are by occupation but do not provide a modeled weekly estimate, so it is hard to estimate how many hours each person employed in that certain occupation works. Furthermore, it is hard to know how many people worked in that occupation to try to gain a sense of the aggregate usual earnings.

We did not use this data for the reasons detailed above, and also because the data is presented as an index. We find that presenting understandable data regarding earnings that people can relate to is more meaningful and easier to use for the general population. We also did not use this data because again the microdata is not available.

C. Quarterly Census of Employment and Wages

The QCEW takes data from all establishments, both farm and nonfarm private, as well as all three levels of government. The survey takes most of its information from state unemployment insurance programs. The holes in this data are then filled in by the Annual Refiling Survey and the Multiple Worksite Report. The data is then reviewed and compiled and published at about a half of a year lag.²⁶

The main advantage to this survey is its comprehensiveness. It uses data from over 10 million establishments to provide estimates. This lowers the potential sampling error. The survey also covers all types of workers, including all levels of government and private firms. In addition, it provides estimates at the county, MSA, state, and national levels at the industry level, so the data published is very granular.

The disadvantage of this data is that it publishes the mean, so it would be hard to compare without being skewed by high-wage earners. Moreover, it presents the data at least five months late. Thus, while it is useful for historical research, by the time that the indicator is ready to be presented, it is not valuable for current policymaking.

²⁵ <https://www.bls.gov/ncs/ect/ectfaq.htm#faqE1>

²⁶ <https://www.bls.gov/opub/hom/cew/data.htm>

We did not use this data for the reasons above, and being an establishment survey, the microdata is protected and not available to the public.

^[1] <https://www.census.gov/prod/2006pubs/tp-66.pdf>