Introducing Demographic Labor Market Data into the U.S. National Accounts

Authors Jon Samuels, U.S. Bureau of Economic Analysis and
Jerin Varghese, U.S. Bureau of Labor Statistics

Contact jon.samuels@bea.gov

Date June 2023

Abstract The U.S. gross domestic product and its foundational National Income and Product Accounts contain some of the most widely used and followed economic statistics in the world yet contain limited information on the labor market and almost no information on demographic groups. We build a new dataset that includes labor market data cross-classified by sex, age, education, and industry and integrate this into the National and Industry Economic Accounts. To overcome small sample size issues for poorly measured demographic groups, we apply small area estimation to refine the estimates. We present examples of how this data can be used to better understand relationships between economic growth and labor market outcomes by demographic group.

Keywords National Accounting, Labor Market, Demographic Groups

JEL Code E01

Bureau of Economic Analysis
U.S. DEPARTMENT OF COMMERCE

1. Introduction

Labor market income has accounted for the majority of national income generated each year since the national accounts data begins in 1929.[1] Information within the national accounts on how this income is generated and dispersed is limited. The National Income and Product Accounts (NIPAs) contain information on total employee labor compensation and wage and salary income by industry and the U.S. Bureau of Economic Analysis (BEA)-U.S. Bureau of Labor Statistics (BLS) Integrated Industry-Level Production Account (ILPA) contains the same information but augments labor compensation to include measures of self-employed labor income and presents information on two demographic groups: workers with at least a bachelor's degree and other workers.

The NIPA and gross domestic product (GDP) accounts were conceptualized to measure production but the call for the national accounts to reflect economic outcomes other than top-line production is broad and deep. Executive Order 13985 Sec. 9 calls for economic datasets to be disaggregated by demographic variables in order to inform efforts around analyzing economic equity. As noted in an influential independent report on U.K. economic statistics (Bean 2016), "GDP is not a measure of welfare and does not reflect economic inequality or sustainability (environmental, financial, or other)." This call for action has been echoed across the institutional framework of the United Nations Statistical Division (UNSD), the International Monetary Fund, the World Bank, the Organisation for Economic Co-operation and Development, and national statistical offices. For example, the UNSD recently convened a network of economic statisticians to convene a "Beyond GDP" sprint.[2] Within the U.S. national accounts, these efforts have led to new statistics on the distribution of income.[3]

This paper integrates information on labor market outcomes by demographic group into the U.S. national accounts. The advantage of integrating this data into the national accounts is that it recognizes key interplays between the evolution of production and the labor market and how changes in production may impact demographic groups differentially. Furthermore, there may be changes in the labor market (either on the supply or demand side), such as whether to search for employment or the need to telework or have flexible hours, that impact industries in different ways. Presenting this information in an internally consistent national account allows for analysis of these types of questions that separate (and likely inconsistent) labor market and industry data cannot easily address.

1. Based on NIPA table 1.12, labor compensation as a share of national income has exceeded 0.50 in every year since 1929.

2. UNSD — Economic Statistics

3. Distribution of Personal Income | U.S. Bureau of Economic Analysis (BEA)

At the outset, it is important to emphasize that the value added of the new labor data presented in this paper is not the labor data itself, as much of this is available elsewhere. The value added is that having the data constructed to be consistent with and integrated into the national accounts is potentially useful for addressing important economic and policy questions. For example, recent work indicates that the COVID-19 pandemic may have impacted women differentially (Fairlie, Couch, and Xu 2021). Importantly, Goldin 2022 finds that education played an important role in the impact of COVID-19 on women, confirming that it is important to have cross-classified data to understand economic impacts. The premise of this paper is that having labor data integrated with national and industry accounting can help shed light on the interactions between the economy, industry, and labor market outcomes.

Two other trends that are gaining importance in understanding the evolution of economic well-being are the age distribution of the workforce and the capacity to work from home. Acemoglu, Mühlbach, and Scott 2022, for example, find large increases and employment gains in “age-friendly” occupations and Dingle and Neiman 2020 show that a large portion of jobs can be done from home. A major motivation for the current paper is that understanding the relationship between these changes and other changes in the structure of economic production at the industry level would be advanced by integrated production and labor market data.

The remainder of the paper proceeds as follows. In Section 2, we describe the methodology for constructing the labor data and integrating it into the national accounts. In Section 3, we provide examples of how the data can be used. In the first part of this section, we conduct an expanded industry growth accounting exercise that parses out workers by demographic group and in the second part, we use the data to explore the relationship between changes in the structure of industry production and the labor market. Section 4 concludes and provides potential next steps.

2. Measuring Labor Market Data by Demographic Group

This section provides details on how demographic information on labor market outcomes is introduced into the national accounts. As an accounting identity, gross domestic product equals gross domestic income, the largest source of which is labor compensation paid to employees. Information on the aggregate value of this labor compensation is included in the first section of the NIPA tables, specifically Section 1 Domestic Product and Income. In the Section 6 Tables on Income and Employment by Industry, number of workers, hours worked, wage and salary, and labor compensation are presented by industry. None of this data is available by demographic group, however. Within BEA’s industry-level accounts, the ILPA includes information on labor compensation and labor input that is broken down by workers with a bachelor’s degree and above along and workers below this level of educational attainment. This is the only BEA data with a demographic classification that is constructed to be consistent with the national accounts.

The starting point for introducing demographic data into the national accounts is the underlying data of the ILPA because this data is already used to construct official statistics. To measure the growth in total factor productivity (TFP) by sector, BEA and BLS distinguish between growth in labor input and labor hours, where the labor input measure captures substitution to or away from different types of workers by constructing price and quantity measures for the heterogenous workforce. To have a tractable approach over a long time series, workers are cross-classified by educational attainment, age, sex, and class of worker. The groups are listed in the table below.

Age group	Education group	Sex	Class of Worker
<16	0-8th Grade	Male	Employee
16-17	9th-12th grade	Female	Proprietors and unpaid family workers
18-24	High school diploma		
25-34	Some college		
35-44	Bachelor's degree		
45-54	More than a bachelor's degree		
55-64			
65+			

To account for industry heterogeneity, these groups are tallied across 63 sectors that span the economy. In total, this amounts to 8x6x2x2x63 or over 12,000 different worker types every year. Labor measures that are tracked include the number of workers, average weekly hours, and average labor compensation per hour. Together we refer to this set of data as the labor matrices. In addition to providing the basic building blocks for the measure of labor input by industry that is included in the ILPA, this labor data is useful for understanding the evolution of the labor market and broader economic outcomes across demographic groups. For example, the distribution of workers across demographic characteristics is important for understanding how changes in the industry impact workers differentially across demographic groups. Differences in average weekly hours are important for assessing whether changes in the distribution of hours is due to changes at the extensive or intensive margin. Influential research (Goldin 2014) has demonstrated that flexibility in hours worked is related to pay differences between male and female workers.

Because many of the cross-tabulated groups are likely small cells with high standard errors based on the raw data (e.g., 18–24-year-old female workers with a master's degree in the water transportation sector) we have implemented a new approach to estimating the labor matrices based on a combination of data. The new method incorporates information from the 1- and 5-year American Community surveys (ACS) and the outgoing rotation group from the basic monthly Current Population Survey (CPS).

An important finding led us to consider the proposed method that combines raw data and estimates small cells instead of using the direct data-based measure. Due to the number of industries and the level of cross-classification in demographic groups, using a single data source (even one as large as the ACS) is not sufficient. This conclusion was based on an analysis of the labor composition measures that are derived from the labor matrices. Labor composition is defined as the difference in the growth rate of labor input and hours worked (Jorgenson, Ho, and Stiroh 2005) and is constructed to capture substitution towards workers with a relatively high marginal product. The BEA-BLS ILPA described in Garner, Harper, Russell, and Samuels 2021 uses the method described in Jorgenson, Ho, and Stiroh 2005. The Jorgenson, Ho, and Stiroh 2005 method makes use of the decennial census micro data sample to fill in the complete set of labor matrices in census years. In that work, they estimated the number of workers, hours worked, and labor compensation for about 7,400 demographic groups per year. In the BEA-BLS ILPA, estimates for about 12,000 different types of workers are needed because the ILPA provides estimates for 63 industries instead of the 44 industries in Jorgenson, Ho, and Stiroh 2005. Because the ACS replaced the decennial census, the first method that was considered for an updated set of labor matrices was to use the ACS as the sole data source, as the decennial census was the sole data source in benchmark years in Garner, Harper, Russell, and Samuels 2021. Unfortunately, labor composition measures produced solely with the ACS appeared to be too volatile for relatively small industries (even at the three-digit level), implying that even if data from a 1 percent census was available every year, it likely would not be suitable for industry level measures of labor composition at the current level of industry detail. Furthermore, this implies that any single data source that is currently publicly available would not be sufficient for measuring distribution of workers across detailed industries and demographic groups.

As an overview of the new method, we take the following approach to construct employment, hours, and compensation estimates cross-classified by demographic groups. First, we construct initial estimates for each year using the ACS. That is, we fill in each of the matrices directly with observations from the ACS. Second, we use small area estimation (SAE) to refine the direct estimates. Small area estimation constructs new estimates of the complete cross-classification of workers as a weighted average of the direct data-based estimate and a model-based estimate, where the weights depend on the uncertainty in direct data-based estimate relative to the model-based estimate. Third, we use information from the monthly current population survey to adjust workers across industries for multiple job holders. Fourth, we use iterative proportional fitting to fit the industry-based labor matrices to labor matrices at the aggregate level based on annual averages from the outgoing rotation group of the CPS. Fifth, we scale the labor matrices to control totals in the national accounts.

We describe each of these steps in more detail below.

Step 1: Construct "direct" estimates of the number of workers, average hours per week, weeks per year, and average hourly wages from the ACS. These estimates are cross-classified by industry, sex, age group, education group, and class of worker.[4] Based solely on information from the ACS, we construct four multidimensional matrices:

E_{jsaect}^{ACS} : number of workers cross-classified by industry, sex, age group, education group, class, and time period. We construct this matrix by allocating person weights to each cell in the matrix. That is, summing over all workers gives the ACS estimate of the total number of workers in year t. The armed forces and institutionalized population are excluded from this total.

h_{jsaect}^{ACS} : average weekly hours cross-classified by industry, sex, age group, education group, class, and time period. We construct this matrix by allocating person weights*average usual weekly hours reported in the ACS to each cell in the matrix and then dividing by the number of workers in the cell that reported working non-zero hours. This implicitly sets average hours worked for those that did not report hours to equal the average for other workers in the cell. Importantly, usual hours are adjusted to actual hours in a later step based on information in the CPS. While using actual hours worked is essential for total hours measures, the difference between actual and usual hours may not be as important in the labor composition adjustment.

w_{jsaect}^{ACS} : average weeks worked cross-classified by industry, sex, age group, education group, class, and time period. We construct this matrix by allocating person weights*average weeks reported in the ACS to each cell in the matrix and then dividing by the number of workers in the cell that reported working non-zero weeks. This implicitly sets average weeks worked for those that did not report hours to equal the average for other workers in the cell. Because the ACS reports only the range of weeks worked, we use the midpoint.[5] Importantly, we adjust weeks worked to actual reported weeks in a subsequent step using information in the CPS.

---

4. We classify all government workers into the government sector if they report government as their class of worker. For example, a State and Local teacher who reports Education as their industry, and State and Local government as their class will get classified into the State and Local government industry. This is consistent with the approach used in the national accounts.

5. This includes paid time off and includes weeks when the person only worked a few hours.

Using E^{ACS}_{jsaect}, h^{ACS}_{jsaect}, and w^{ACS}_{jsaect}, total hours worked for each cell is $E^{ACS}_{jsaect} * h^{ACS}_{jsaect} * w^{ACS}_{jsaect}$ and total hours worked in the industry can be attained by summing over all of the demographic groups within each industry.

c^{ACS}_{jsaect} : average hourly wages cross-classified by industry, sex, age group, education group, and time period. Note that we estimate this matrix for employee (class 1) workers only because wages for self-employed workers capture a mix of labor and other income. We construct this matrix by allocating person weights*annual wages reported in the ACS to each cell in the matrix and then dividing by the number of workers in the cell that reported wages. This implicitly sets annual wages worked for those that did not report wages to equal the average for other workers in in the cell. Total annual wages worked in the industry is then annual wages * E^{ACS}_{jsaect}. Total annual wages is divided by total hours worked ($E^{ACS}_{jsaect} * h^{ACS}_{jsaect} * w^{ACS}_{jsaect}$) to get c^{ACS}_{jsaect}. In a later step, we set the wages of class2 (unincorporated self-employed and unpaid family workers) workers equal to that of class1 (employee) workers by industry and demographic group.

Using E^{ACS}_{jsaect}, h^{ACS}_{jsaect}, w^{ACS}_{jsaect}, and c^{ACS}_{jsaect} total wages for each cell is $E^{ACS}_{jsaect} * h^{ACS}_{jsaect} * w^{ACS}_{jsaect} * c^{ACS}_{jsaect}$, and total wages worked in the industry can be attained by summing over all of the demographic groups within each industry. In a later step, wages are adjusted to total labor compensation.

A final note on this step is on the timing of the ACS. The ACS uses a rolling sample and asks people questions about their previous 12 months. Therefore, responses collected in January of a given year mostly reflect the previous year, while responses collected in December reflect most of the current year. We do not attempt to adjust for this timing issue directly; in later steps the labor matrices are adjusted to the annual CPS so this provides some adjustment into period-t timing.

Step 2: In this second step, we use SAE to refine estimates from the ACS based on each cell's sampling error. The basic objective in using SAE is to borrow information from similar cells with low sampling errors to inform estimates of similar demographic cells with relatively high sampling errors. We compute sampling errors for each cell for each year based on replication weights in the ACS. Instead of applying SAE directly to the employment matrix, we apply it to a matrix of employment shares by industry. The reason for this is twofold. First, our dataset fundamentally involves the shares of workers across demographic groups. Second, operating on employment levels themselves may provide odd results. Consider the case where the poorly estimated (small) industry has few observed highly educated, young workers, while the well-estimated (large) industry has many highly-educated young workers. Using the level of workers in the large industry to inform the level of workers in the small industry likely would lead to odd results yet using the share of highly-educated young workers would provide reasonable comparisons.

We now describe the SAE model, which we implement in two passes. Each SAE is an iteration of the Fay and Herriot 1979 model. To provide intuition for the approach, we start with the last step in constructing the final estimate for each cell of the labor matrices. Borrowing some of the notation in Mukhopadhyay and McDowell 2011, the final estimate for each cell i in each run of the model is constructed as:

$$\theta_i = \gamma_i \bar{y}_i + (1-\gamma_i)\hat{y}_i \tag{1}$$

Equation (1) states for each cell i in each labor matrix (employment (share), hours per week, weeks/yr, compensation/hr), the final estimate is constructed as a weighted average of the direct data-based estimate, $\bar{y}_i$, that is the tabulation of the raw data from the ACS, and a model-based estimate $\hat{y}_i$. Note here that i stands for industry, sex, age, education, class cell. The weight that goes on each, γ_i, reflects the uncertainty in the direct data-based estimate and the model-based estimate and is determined by comparing the standard errors in the direct survey data to the standard errors from the model. Thus, if a direct estimate using the ACS data is relatively imprecise (in terms of standard errors), the model-based estimate would receive a relatively higher weight, and vice versa.

The parameters necessary to implement the weighted average are estimated using an area-level small area model. The basic model is:

$$\bar{y}_i = x_i\beta + \mu_i + \bar{\varepsilon}_i \tag{2}$$

Where $\bar{y}_i$ is the cell average of the variable of interest (again industry, sex, age, education, class cell), x_i is a set of predictors, μ_i, is a cell-specific random effect, and $\bar{\varepsilon}_i$ is an (assumed to be known) error capturing the standard error of the measured cell average in the direct data-based estimate (tabulated from the replication weights). In our application, $\bar{y}_i$ is a scalar (which gets grouped with other scalars included in the model, i.e., other areas included in the model to make a vector of observations) and holds the average for a single sex, class, age, education, industry cell and $\bar{\varepsilon}_i$ is the standard error in that estimate (think due to sample size). Both μ_i and $\bar{\varepsilon}_i$ are assumed mean zero, so that the best linear unbiased predictor of $\bar{y}_i$ is $\hat{y}_i = x_i\beta$. This is not a standard ordinary least squares regression; μ_i is random effect but because $\bar{\varepsilon}_i$ is assumed to have a known variance; this specification corresponds to the Fay and Herriot 1979 model. The objective of the estimation is conditional on the measures of $\bar{y}_i$, x_i, and $\bar{\varepsilon}_i$ to estimate the unknown parameters β and the variance of μ_i ,which is labeled as σ_u^2 . This is estimated using the MIXED command in SAS. With parameter estimates in hand, the weight in equation (1) is:

$$\gamma_i = \frac{\sigma_u^2}{\sigma_u^2 + D_i} \tag{3}$$

Where D_i is the (measured and assumed known) sampling variance in the ACS-based estimate. Thus, if the model has a high variance relative to the estimate based on the raw data, the raw data (direct estimate) will get a higher weight. D_i is calculated using the replication weights from the ACS.[6]

While the model in (2) is broadly understood and applied, implementing the model requires specification of areas (what to include in each small area as observations in model) and the covariates to include in X. In our implementation, for each demographic group, we treat each industry as an observation, i.e., another area. That is for each demographic group, there are n observations in the model where n corresponds to the number of industries. We run the model separately for every demographic group in every year. That is, for each sex, age, education, class (except for the compensation matrix), and year we run a small area model. This is a total of up to 168 (7 age * 6 education * 2 sex * 2 class of worker) models per year. On the left-hand side are observations across each 4-digit industry (separate areas). The reason why the SAE is run twice is to account for two versions of the right-hand side variables (discussed below). The basic model is specified to test whether within each demographic group, differences across industries are significant relative to the sample mean, conditional on the right-hand side variables and the observed sampling error in the observations. Thus, if the model fits well (resulting in a low σ_u^2), its interpretation is that the differences across industries are not large given the uncertainty in the direct sample-based estimates. We note that we did consider alternative specifications, such as using a broader group of cell neighbors. For example, we experimented with specifying the model where the data was grouped into areas based on a nearest neighbor approach. That is, cells from the sex x class x age x education matrix were pooled into the same model if they were in a neighboring age or education group. The potential advantage of this approach is that allows for more differences across cells to be accounted for in the model in comparison to just industry, but in practice this approach produced model results with much higher variances. The intuition for this high variance is that trying to fit differences across demographic groups without other explanatory variables is often problematic, leaving large variances in the model. For example, using the share of workers with a master's degree or above (which is typically low) grouped with workers with Some College to predict the share of workers with a bachelor's degree (which is typically high) is problematic.

---

6. The model requires estimates of standard errors of each cell. For estimates of the number of workers the formula for these standard errors are a straight-forward application of formulas provided by the U.S. Census Bureau for using the replication weights. But the replication weights are only useful to tell us about uncertainty in the relative weights of each person in the sample and therefore these are not useful in deriving standard errors in certain cases. For example, consider two respondents both of whom reported working 40hrs per week. The replication weights can give us standard errors of these workers share in the workforce, but because both workers reported working 40hrs per week, the resulting standard error for hours per week would be zero. In cases like this, the proposed method uses an adjustment to modify standard errors for cells that have less than or equal to 5 observations.

We use SAE to fill in four matrices: 1) employment shares of each demographic group by industry, 2) hours worked per week for the same groups, 3) weeks worked per year for the same groups, and 4) wages per hour for employees cross-classified by demographic group. That is, we run SAE separately on each of these matrices for each year of the data.

In the first pass of the model, for each year and each demographic group and each of the four matrices, a version of the model is run with two explanatory variables: a constant and an estimate of the same cell from the 5-year sample of the ACS. For clarity, there would be a model for each demographic group and the “areas” would be each four-digit industry. The explanatory variables would be a constant and the same observation for each industry based on the 5-year ACS. The 5-year ACS is used to give the model power against the direct data-based estimate. If the 5-year ACS is a good predictor of the individual industry cells within a demographic group, then the model would have a relatively low standard error and the final estimate would put relatively high weight on the model-based estimate. This potentially has a smoothing effect on the labor composition estimate because it relies on ACS information over multiple years. With this version of the model, we estimate γ_i for each cell and construct θ_i to completely fill in the four labor matrices for each year.

In the second pass of the model, for each year and each demographic group and each of the four matrices, a version of the model is run with an additional explanatory variable: the estimate of the same cell based on the average of the outgoing rotation group from CPS over the same year; however, instead of using the equivalent four-digit cell from the CPS, we use the three-digit/sector level data from the CPS instead due to the limited record count in that dataset. So, for North American Industry Classification System (NAICS) 1111, we use 111CA (NAICS 111,112) data from the CPS. We include observations in this second-step version of the model only if the CPS estimate for that cell exists for every year in the sample. This has the effect of limiting the second-step model to being run only on groups of cells for which the CPS exists for every year in the sample. In practice, this (intentionally) restricts bringing in information from the CPS to cases when the observation is persistent enough to be in every sample, thus dropping observations from very small industry-demographic group pairs.

For example, for a given demographic group, the second pass of the model would include industries where the data from the CPS existed. With this second run of the model, the standard error of the data-based estimate is taken to be the standard error from the final estimate from the first pass. This second pass brings in contemporaneous information to counteract some of the potential smoothing from the first pass, but only for cases when the CPS consistently captures the cell over time. If the cell does not exist in the CPS, the final estimate is that from the first pass. Remember that this final model-based estimate after the second pass is weighted against the variability in the direct data-based estimate via equation (3).

In practice, the small area model results in a distribution of weights between the direct data-based estimate and the model-based estimate that reflects the distribution of sample sizes of each of the cells. The estimated distribution of weights is bimodal. Figure 1 shows the distribution of weights for the employment share estimate. The peak at the top end of the distribution reflects the many cells where there are sufficient observations such that the direct data-based estimate has a much lower standard error than the model-based estimate, resulting in a weight that is near one on the raw data. The peak at the low end of the distribution reflects that there are many small cells where the model-based estimate has a much lower standard error than the data, resulting in a weight close to zero. The distribution between the two peaks shows a shallow U-shape, indicating that the method has the tendency to push weights away from a simple 50-50 average. The distribution of weights suggests that the model works reasonably; a large share of cells is well-estimated in the raw data and receives little model weight, another large chunk are small cells with raw data that receives little weight, and the remaining group has weights that reflect a balance of the two.[7]

Figure 1: Weights on the Direct Estimate of Employment Share from the ACS

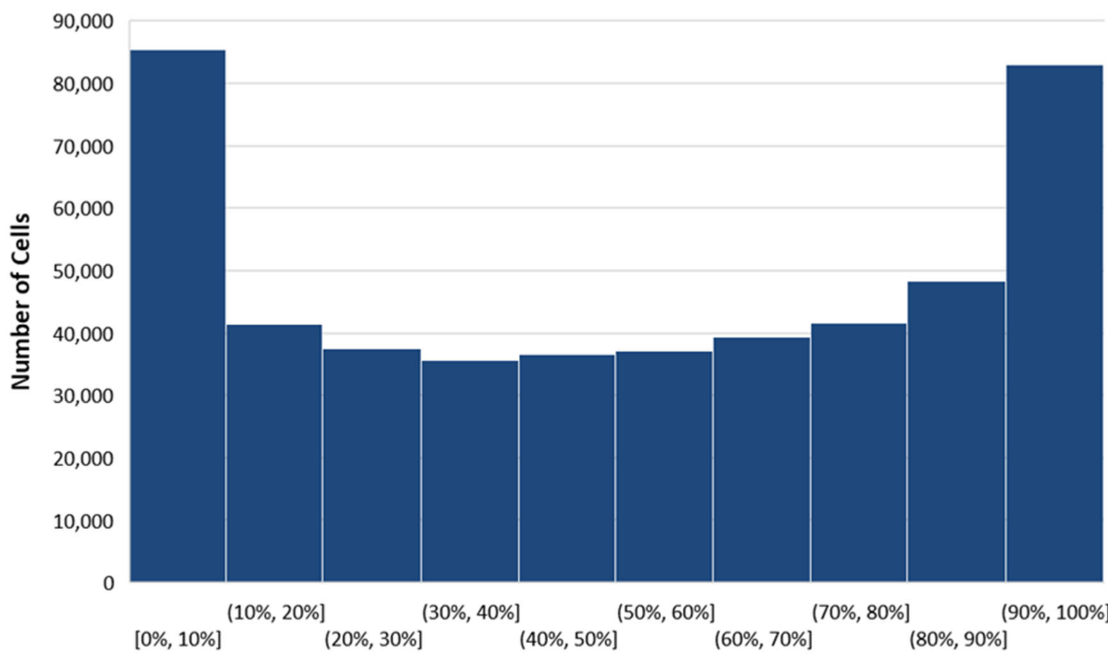

7. We note that not all small cells receive high model weights. Many small cells are precisely estimated in the data, that is precisely estimated to be small. If the standard error of the direct data-based estimate is small enough, the small cell receives a weight that reflects this low standard error.

Step 3: Adjust estimates for multiple-job holders. Many workers hold more than one job but the ACS tracks only their primary job. For example, a 9-5 white collar worker who drives a taxi after-hours would only be assigned to their primary industry in our tabulations of the ACS. If this worker's demographic group is atypical in comparison to the secondary industry, this could be important in measuring the demographic distribution in the second industry (for example, a lawyer that drives a taxi). This is a separable problem from constructing total hours worked by industry that must also account for secondary jobs; secondary jobs impact the composition of workers only if the primary industry has a different composition of workers than the secondary industry. To adjust workers for secondary jobs, we construct a concordance of primary and secondary jobs by detailed industry based on monthly information available in the CPS. This results in an estimate for each primary industry that shows what percentage of workers' hours take place in each secondary industry. We then reallocate this percentage of workers and hours from each demographic group within the ACS industry matrix to each secondary industry. For example, if our concordance told us that five percent of workers in the primary industry retail had a secondary job in transportation, we would reallocate five percent of the workers and their hours to the transportation sector.[8] This approach pushes the labor composition of the secondary industry toward that of the primary industry. Note, we do not reallocate wages. The assumption for wages is that a worker who has a second job earns the average wages by demographic group in that secondary industry. For example, a young lawyer with a J.D. degree who drives a taxi on the weekend to pay off her loans is assumed to earn the wage rate of the average young taxi driver in the transportation industry who has an advanced degree, not the average wage in the legal sector.

Step 4: Balance the matrices from step 3 to aggregate demographic composition from the annual average of outgoing rotation group of the CPS. The CPS on its own is too small of a survey to estimate the approximately 168 worker types by industry that are necessary to fill in the demographic labor matrices, yet the CPS has advantages as well. A major advantage is that the measures in the CPS more closely correspond to the concepts that appear in the model of labor input and TFP. First, it contains measures of actual hours and weeks worked for the same time period every year. Second, it is used broadly in the research community to analyze trends in the labor market over time. Finally, there is a long time series of available micro data from the CPS allowing for consistent control totals back through history for historical measures of the labor matrices.

---

8. We take the additional set of accounting for class of worker in this reallocation. We split the reallocated workers and hours by class of workers in the secondary industry. For example, if a lawyer grows vegetables as a secondary job, these hours would get reallocated to the farm sector proportionally to the distribution of hours between employees and self-employed in the farm sector.

To balance the industry-based matrices to the national totals we rely on iterative proportional fitting (RAS). The target for the RAS is the same four matrices on workers, hours per week, weeks per year, and wages per hour, but excluding the industry dimension: $E_{saect}^{CPS}, h_{saect}^{CPS}, w_{saect}^{CPS}, c_{saect}^{CPS}$. In practice, the RAS operates by collapsing over all dimensions except one, and scaling each of the rows in the industry-based matrix to match the corresponding distribution total based on the CPS. But doing this scaling changes the distribution across other dimensions. Therefore, the RAS procedure continues to iterate on each dimension until a convergence criterion is met. The large number of industries per each demographic group corresponds to many degrees of freedom for the RAS procedure, so it typically converges to the national totals very quickly. It is not possible to RAS directly to the $h_{saect}^{CPS}, w_{saect}^{CPS}, c_{saect}^{CPS}$ matrices because the way the RAS works is by summing up across dimensions to create a scaling factor. But it would be nonsensical to sum hours per week, for example. Thus, we construct temporary target matrices as the control total. For example, to RAS the hours per week matrix, we construct a target matrix $E_{saect}^{ACS,R} * h_{saect}^{CPS}$ (ie excluding the industry dimension) where $E_{saect}^{ACS,R}$ is the employment matrix after small area estimation and RAS have been applied and use this as a target for the $E_{jsaect}^{ACS,R} * h_{jsaect}^{ACS}$ matrix. After the RAS, we divide each cell in RASes matrix by $E_{jsaect}^{ACS,R}$ to retrieve the final estimate of hours worked per week, and similar for the weeks and wages per hour matrices. We apply steps 1)-4) at the NAICS four-digit level, and then aggregate up to the three-digit level to construct three-digit estimates after the RAS. This ensures consistency between the three- and four-digit results.

Step 5: Scale the labor matrices to control total in the national accounts. The matrices based on the ACS and CPS are based on household surveys. The GDP accounts are a production account from the perspective of a producer. To convert the household matrices to be consistent with the establishment-based constructs in the national accounts, a simple scaling is applied. Total workers by industry are scaled to workers in the NIPA accounts separately for employees and self-employed workers. The national accounts estimates are constructed using establishment surveys thus reflect workers on the payroll by industry; thus this aligns with the number of workers actually used in the production and aligns conceptually with the output and other inputs measured in the GDP accounts.[9] Similarly, total hours by industry and demographic group are scaled to total hours in the NIPA accounts and wages are scaled to total labor compensation in the NIPA accounts.[10] Finally, because labor compensation per hour is not available for self-employed workers, we set labor compensation per hour for self-employed

9. Isenberg, Landivar, and Mezey 2013 find that the match on industry between survey and administrative is not perfect. Thus our method can be interpreted as basing the labor composition measure on information reported in the household survey and hours with establishment-based data.

10. The same level of industry detail is not available for hours. Therefore, hours are first scaled to an unpublished tabulation from BLS on total hours worked by industry and then to the available industry detail in the national accounts.

workers to the be the same as the payroll employee cross-classified by demographic group. Therefore, the 55–64-year-old female lawyer with a private practice is assumed to earn the same compensation per hour as the 55–64-year-old female lawyer who works for the corporate sector.

3. Demographics and Industry Economic Growth

This paper does not aim to provide comprehensive coverage of the use cases for having labor data cross-classified by demographic group integrated into the national accounts. But a few examples of results that come from integrated growth accounting and demographic labor data are given in the following sections. The major benefit is that data of this nature allows users to perform an integrated examination of many issues related to industries, labor markets, and their interactions. In part I of this section, a few examples of the role of different demographic groups (based on the above classification) in industry economic growth are explored. A reason that this is important is that policies that change the tradeoffs faced by producers in hiring workers from different demographic groups may impact workers and industries differentially.

3.1. Part I: Decomposing the Sources of Economic Growth

Standard growth accounting decomposes industry growth into the contributions of capital, labor, intermediate input, and TFP growth (Garner, Harper, Russell, and Samuels 2021). In this section, the contribution of labor input is decomposed into workers from different demographic groups. Table 1 focuses on one cut of this tabulation, the contribution of older workers (65 and older) by sex. As in the literature on growth accounting, the contribution of each input is defined as its value share in nominal gross output times the growth rate of the input. It is important to note that the terminology "contribution" should not be taken to mean value in a normative or societal sense, it simply reflects the amount of workers and these workers' share in output as measured. For example, if female workers have a higher contribution to output in the education sector in comparison to male workers, this should not be taken to mean that each female worker is more valuable than each male worker; it means that the growth of female workers multiplied by these workers factor share exceeds the growth of male workers times their factor share. This is an internally consistent way (i.e., consistent with other inputs like capital and energy) to measure the way workers are employed in production.

Table 1: Contributions of Labor Input to Industry Output Growth 2005-2020

	Male 65+ Workers	Female 65+ Workers	Industry Output Growth
Farms	0.26	0.02	1.49
Forestry, fishing, and related activities	0.14	0.05	-1.07
Oil and gas extraction	0.01	0.03	5.15
Mining, except oil and gas	0.04	0.00	-2.20
Support activities for mining	0.01	0.00	-2.04
Utilities	0.05	0.01	0.55
Construction	0.19	0.02	-0.67
Wood products	0.09	0.01	-1.44
Nonmetallic mineral products	0.09	0.02	-1.45
Primary metals	0.05	0.01	-0.64
Fabricated metal products	0.15	0.02	-0.65
Machinery	0.08	0.02	-0.43
Computer and electronic products	0.14	0.04	2.12
Electrical equipment, appliances, and components	0.08	0.02	-0.99
Motor vehicles, bodies and trailers, and parts	0.02	0.00	0.35
Other transportation equipment	0.09	0.01	-0.49
Furniture and related products	0.15	0.03	-3.28
Miscellaneous manufacturing	0.08	0.02	-0.04
Food and beverage and tobacco products	0.04	0.01	0.27
Textile mills and textile product mills	0.05	0.03	-4.89
Apparel and leather and allied products	0.04	-0.03	-6.77
Paper products	0.05	0.00	-0.99
Printing and related support activities	0.08	0.02	-2.59
Petroleum and coal products	0.00	0.00	-0.23
Chemical products	0.04	0.01	-0.67
Plastics and rubber products	0.09	0.02	-0.80
Wholesale trade	0.12	0.03	1.99
Retail trade	0.09	0.04	1.58
Air transportation	0.05	0.03	-1.93
Rail transportation	0.03	0.00	-0.65
Water transportation	0.03	0.00	-0.30
Truck transportation	0.14	0.01	0.69
Transit and ground passenger transportation	0.21	0.03	1.53
Pipeline transportation	0.06	0.00	2.22
Other transportation and support activities	0.19	0.05	0.97
Warehousing and storage	0.10	0.06	6.35
Publishing industries, except internet (includes software)	0.08	0.07	2.49
Motion picture and sound recording industries	0.06	0.00	0.22
Broadcasting and telecommunications	0.02	0.01	3.05
Data processing, internet publishing, and other information services	0.03	0.01	11.30
Federal Reserve banks, credit intermediation, and related activities	0.05	0.04	0.10

Table continues on next page

	Male 65+ Workers	Female 65+ Workers	Industry Output Growth
Securities, commodity contracts, and investments	0.23	0.07	1.03
Insurance carriers and related activities	0.11	0.07	4.05
Funds, trusts, and other financial vehicles	0.01	0.00	1.16
Real estate	0.03	0.02	1.25
Rental and leasing services and lessors of intangible assets	0.02	0.00	1.06
Legal services	0.38	0.12	-1.06
Computer systems design and related services	0.18	0.04	5.83
Miscellaneous professional, scientific, and technical services	0.21	0.09	2.41
Management of companies and enterprises	0.32	0.09	3.09
Administrative and support services	0.16	0.08	3.14
Waste management and remediation services	0.19	0.03	0.87
Educational services	0.21	0.20	1.36
Ambulatory health care services	0.26	0.20	2.10
Hospitals and nursing and residential care facilities	0.13	0.16	2.14
Social assistance	0.16	0.36	2.19
Performing arts, spectator sports, museums, and related activities	0.13	0.09	-0.93
Amusements, gambling, and recreation industries	0.03	0.00	-0.76
Accommodation	0.05	0.01	-1.96
Food services and drinking places	0.06	0.02	0.16
Other services, except government	0.15	0.09	-0.38
Federal	0.11	0.07	0.93
State and local	0.10	0.09	1.06

Table 1 presents aggregates of the underlying detail which subsumes education and class of worker. These tabulations show that older workers contributed significantly across many sectors but the industries in which these contributions were focused differed by sex. For example, the industry where older male workers had the largest contribution was the legal services sector, but older female workers had the largest contribution in the social services sector. Older male workers also contributed significantly to output growth in management of companies and enterprises, farms, and ambulatory health care services. Older female workers contributed significantly in the educational services, ambulatory health care services, and the hospitals and nursing and residential care industries. Obviously, this single tabulation does not tell a definitive story about how male and female workers are used across industries. For example, it could be the case that younger female workers are growing in the management of companies sector relative to male workers even though table 1 shows that older male workers had a larger contribution than female workers. The cross-classified data allows for examining this type of question and the basic idea is that making this data available to the public would permit such analysis of this and a large set of similar questions.

Table 2 shows that industries' use of workers varies by level of educational attainment. The fastest growing industry over this time period was data processing, internet publishing, and other information services; growth of workers with at least a bachelor's degree was particularly important for this sector. The second fastest growing industry over this period was warehousing and storage; this is the sector where workers with some college but no bachelor's made the largest contribution to output growth. This reinforces the usefulness of integrating labor data with industry growth accounting data. Workers with some college but no bachelor's were also important for growth in the waste management and remediation services sector.

It is also useful to distinguish between workers with a bachelor's degree and those with a higher level of educational attainment. Both workers with a bachelor's degree and at least a master's degree made large contributions in the management of companies and enterprises sector and computer systems design and related services, but workers with at least a master's degree made larger contributions in the management of companies and enterprises, educational services, computer and electronic products, and the government sectors.

Table 2: Contributions of Labor Input to Industry Output Growth 2005-2020

	Workers with Some College but no Degree	Workers with BA	Workers with MA+	Industry Output Growth
Farms	0.06	0.26	0.05	1.49
Forestry, fishing, and related activities	0.08	0.21	0.18	-1.07
Oil and gas extraction	0.01	0.16	0.12	5.15
Mining, except oil and gas	-0.09	0.02	0.05	-2.20
Support activities for mining	0.08	0.22	0.09	-2.04
Utilities	-0.05	0.16	0.10	0.55
Construction	0.08	0.32	0.11	-0.67
Wood products	-0.17	-0.07	0.01	-1.44
Nonmetallic mineral products	-0.18	0.12	0.10	-1.45
Primary metals	-0.13	0.07	0.01	-0.64
Fabricated metal products	0.00	0.14	0.07	-0.65
Machinery	-0.18	0.16	0.14	-0.43
Computer and electronic products	-0.34	-0.27	0.04	2.12
Electrical equipment, appliances, and components	-0.19	0.30	0.17	-0.99
Motor vehicles, bodies and trailers, and parts	-0.13	0.00	0.02	0.35
Other transportation equipment	-0.09	0.18	0.17	-0.49
Furniture and related products	-0.38	-0.19	0.07	-3.28
Miscellaneous manufacturing	-0.16	0.15	0.23	-0.04
Food and beverage and tobacco products	0.10	0.21	0.11	0.27
Textile mills and textile product mills	-0.32	-0.11	0.07	-4.89
Apparel and leather and allied products	-0.81	-0.77	0.00	-6.77
Paper products	-0.13	0.00	0.02	-0.99
Printing and related support activities	-0.70	-0.33	-0.04	-2.59
Petroleum and coal products	-0.02	-0.01	0.00	-0.23
Chemical products	-0.05	0.05	0.14	-0.67
Plastics and rubber products	-0.08	0.16	0.04	-0.80
Wholesale trade	-0.08	0.22	0.13	1.99
Retail trade	-0.10	0.16	0.18	1.58
Air transportation	-0.45	-0.40	0.01	-1.93
Rail transportation	-0.51	-0.06	0.08	-0.65
Water transportation	-0.06	0.06	0.03	-0.30
Truck transportation	0.13	0.16	0.08	0.69
Transit and ground passenger transportation	0.31	0.28	0.17	1.53
Pipeline transportation	0.08	0.04	0.08	2.22
Other transportation and support activities	0.92	0.89	0.33	0.97
Warehousing and storage	1.44	0.72	0.18	6.35
Publishing industries, except internet (includes software)	-0.37	0.40	0.58	2.49
Motion picture and sound recording industries	-0.33	-0.09	0.02	0.22
Broadcasting and telecommunications	-0.28	-0.19	-0.03	3.05

Table continues on next page

	Workers with Some College but no Degree	Workers with BA	Workers with MA+	Industry Output Growth
Data processing, internet publishing, and other information services	0.06	1.18	1.01	11.30
Federal Reserve banks, credit intermediation, and related activities	-0.22	0.33	0.37	0.10
Securities, commodity contracts, and investments	-0.11	0.51	0.53	1.03
Insurance carriers and related activities	0.02	0.62	0.43	4.05
Funds, trusts, and other financial vehicles	0.02	0.12	0.09	1.16
Real estate	0.00	0.08	0.04	1.25
Rental and leasing services and lessors of intangible assets	-0.06	-0.03	0.02	1.06
Legal services	-0.18	0.23	-0.26	-1.06
Computer systems design and related services	0.39	2.65	1.66	5.83
Miscellaneous professional, scientific, and technical services	-0.02	0.93	0.84	2.41
Management of companies and enterprises	0.26	1.13	1.78	3.09
Administrative and support services	0.19	0.69	0.29	3.14
Waste management and remediation services	0.48	0.44	0.24	0.87
Educational services	-0.08	0.65	1.43	1.36
Ambulatory health care services	0.42	0.90	1.29	2.10
Hospitals and nursing and residential care facilities	0.08	0.89	0.99	2.14
Social assistance	0.44	1.48	1.18	2.19
Performing arts, spectator sports, museums, and related activities	-0.54	-0.33	-0.18	-0.93
Amusements, gambling, and recreation industries	-0.47	0.12	0.10	-0.76
Accommodation	-0.31	-0.09	-0.01	-1.96
Food services and drinking places	0.21	0.22	0.08	0.16
Other services, except government	-0.14	0.32	0.23	-0.38
Federal	-0.18	0.29	0.62	0.93
State and local	-0.16	0.20	0.51	1.06

3.2. Part 2: Industry growth and Demographic Groups

The first part of this section provided two examples of the usefulness of decomposing labor input by demographic group as a separately measured contribution to industry economic growth. This section explores some of the relationships between industry growth and the labor market. An example of this type of question is: are male and female hours converging in industries that are growing relatively more rapidly. Is this true for all workers or only for relatively young workers? Another question is: what is the relationship between TFP growth and the use of labor by demographic group? Is TFP growth associated with hiring more young workers, less educated workers? What is the relationship between the use of information technology (IT) and demographic groups? Is IT associated with hiring more male workers with a college degree and less older workers, for example? This section does not claim to answer these questions but provides examples of how the data can be used to start to explore questions like these.

For example, figure 2 shows that industries with relatively higher growth rates have faster hours growth for female older (65+) workers relative to hours worked by male workers. There is evidence of conditional convergence as well; that is regressing hours growth of female 65+ workers relative to men on the initial ratio of hours worked and industry output growth is consistent with convergence in the hours worked of older female workers relative to men.[11]

Another question is how the ratio of female to male hours is impacted by an industry's use of IT. Figure 3 shows that the ratio of female to male hours for 35–54-year-old workers grew in most IT-intensive industries over the 2005–2019 period.[12] A large literature examines the role of information technology on the workforce. Data that cross-classifies the workforce by demographic group has the potential to advance this analysis. The plot in figure 3 suggests, for example, that the growth of IT is not biased toward hiring of men, for example, because in many industries where the use of IT is relatively important, the hours of women workers are also growing relative to men. Importantly, both the IT capital data and the labor data are constructed to be consistent with the production function approach to growth accounting. Constructing data with consistent output and inputs often needed to be constructed by individual researchers (for example, Krusell, Ohanian, Rios-rull, and Violante 2000) even though pieces of this are available, so having a merged dataset could enable much more analysis of how changes in production impact certain groups of workers.

---

11. This shows the regression results of growth rate of relative female 65+ hours worked on its initial level and output growth, weighted by industry size.

	(1)
VARIABLES	femhoursrat_grwth
femhoursrat2005	-0.109***
	(0.0184)
Qgrwth	3.475***
	(1.256)
Constant	0.0733**
	(0.0336)
Observations	63
R-squared	0.386

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

12. IT-intensive industries are defined as those within the upper quartile of industries in measured contribution of IT equipment and software to economic growth based on the contribution tables of the ILPA.

Figure 4 indicates that having a master's degree or higher is important for IT-intensive industries. In all of the IT-intensive industries, except legal services, hours of master's degree workers grew faster than workers with a bachelor's degree. Interestingly, the legal services sector was an outlier, the growth of hours worked of bachelor's workers outpaced those with a higher degree over the period. This could reflect an elevated level of hours worked by master's+ workers at the beginning of the sample, or it could reflect a noteworthy difference in the production process of legal services.

Figure 5 displays the trend in compensation per hour for female 35–54-year-old workers relative to male workers of the same age for the same IT-intensive industries. This tabulation shows a broad distribution of changes in relative compensation rates between female and male workers for these industries. For example, female workers in the legal services, management of companies, and insurance carries exhibited wage gains relative to men while relative wages fell in the pipeline transportation, publishing, and computer systems design sectors.

Figure 6 provides suggestive evidence that TFP growth did not have a large differential impact on female versus male workers aged 65 and older across industries. If industries with high TFP growth hired mostly male older workers instead of female workers, for example, then this would likely raise concerns about the impact of TFP on workers. The result in figure 6 shows that industries with higher TFP growth actually had higher growth of female hours relative to male workers. Of course this could simply be a catch up in hours worked of female labor hours, but the evidence that TFP does not appear to be replacing older female workers is indicative of the type of analysis that can be conducted with integrated data. Similarly, figure 7 shows that TFP growth did not appear to have a differential impact on workers with a bachelor's degree versus workers with some college and no degree. Questions and data like this are relevant because there is an influential literature that analyzes the impact of TFP growth on workers with a college degree. Having labor data cross-classified by different workers allows for more nuanced assessment of the labor market impacts of TFP growth.

Finally, figure 8 shows that the relationship between TFP growth and relative compensation per hour for female workers relative to male in the 35–54 age group is heterogeneous across industries. This figure plots the change in relative compensation per hour for this group of workers for industries in the top quartile ranked by TFP growth over the period. For example, in the pipeline and publishing sectors, labor compensation per hour fell for female workers relative to males over this period, but in the water transportation and insurance carrier industry relative labor compensation rose over this period.

4. Conclusions and Next Steps

The purpose of this paper is to present a tractable approach to integrating labor data by demographic into the U.S. national accounts. An advantage of this is approach is that it is based on established methods with clean links to measuring the sources of economic growth. Another advantage is that historical datasets exist that could extend the work back until 1947. Linking the two would permit important longer-term analysis of how labor market changes related to industry-level changes in the economy.

This paper also serves as a proof of concept that labor market data by demographic group has the potential to yield useful results on the interactions between the industry production process and labor market outcomes by demographic group, including the impact of capital investment and changes in technology and input structure. There are many potential extensions to integrating labor market data by demographic group into the national accounts. The most obvious is to provide additional demographic groups. There are measurement challenges as well that have not been the focus of this paper, including measuring fringe benefits by demographic group, wages of self-employed workers, and incorporating stock options into the compensation measures.

Figure 2: Relative Hours Growth of 65+ Female Workers, 2005–2019. Each point is an industry.

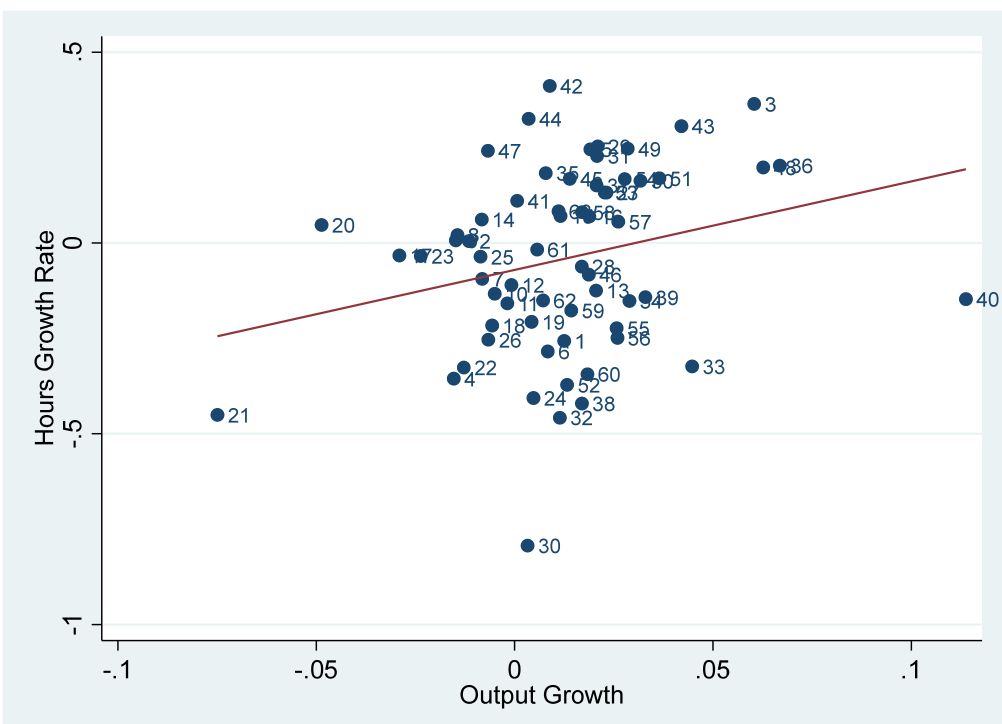

Figure 3: Growth Rate of Female Hours Relative to Male for 35–54-Year-Old Workers for IT-Intensive Industries

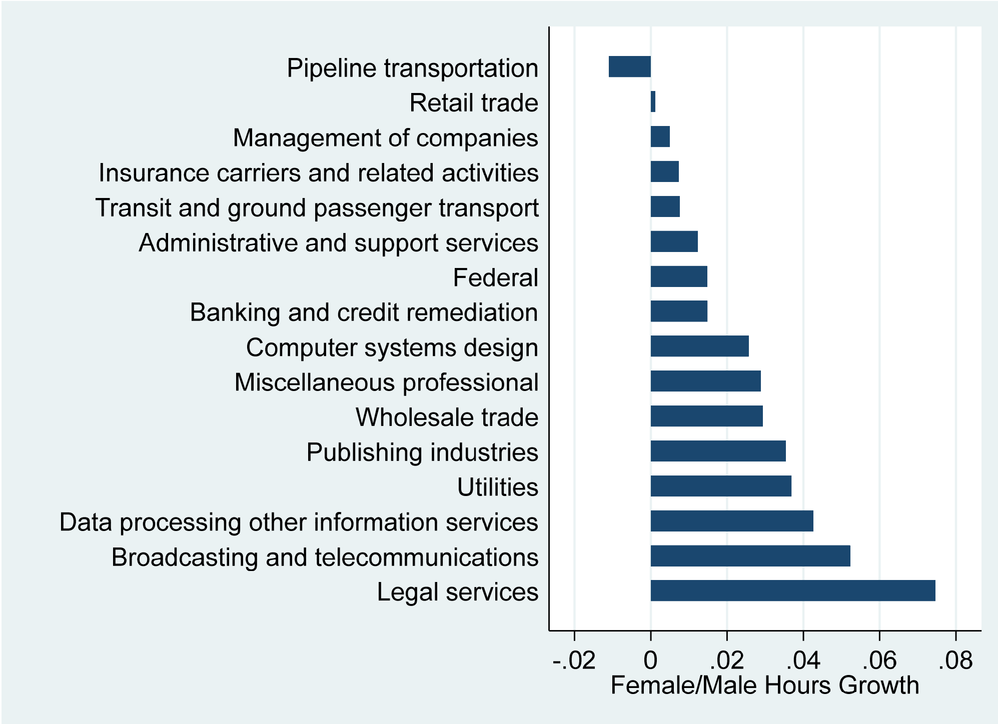

Figure 4: Hours Growth of Workers with at least a Master's Degree Relative to Bachelor's Workers (less the economy-wide average) for IT-Intensive Industries

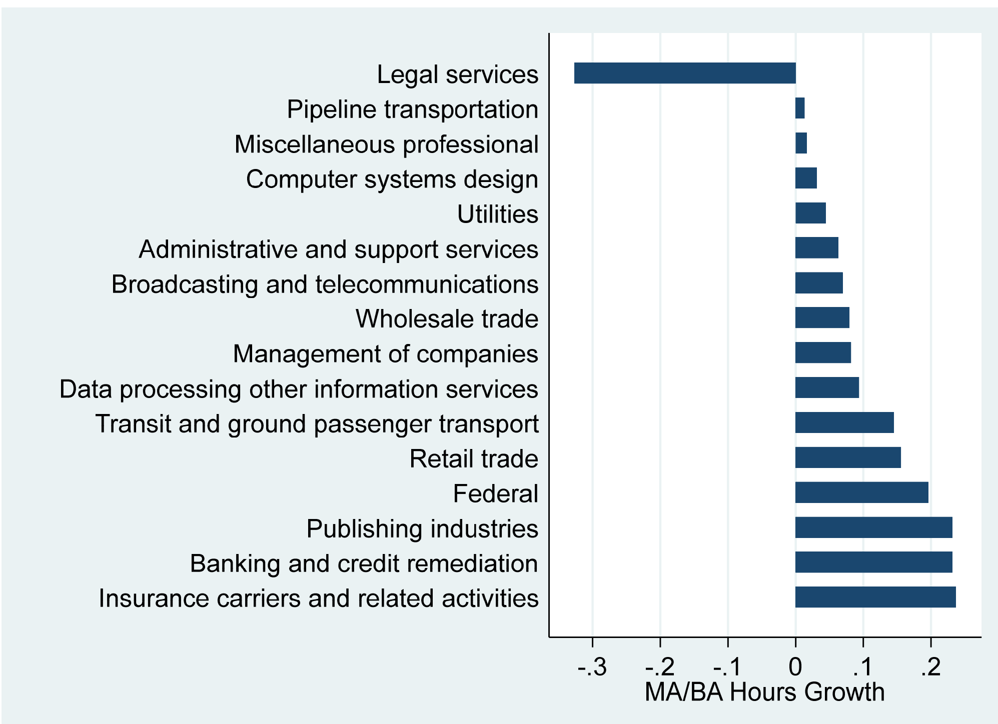

Figure 5: Growth in Compensation per Hour for Female workers Relative to Male, 2005–2019 for IT-Intensive Industries

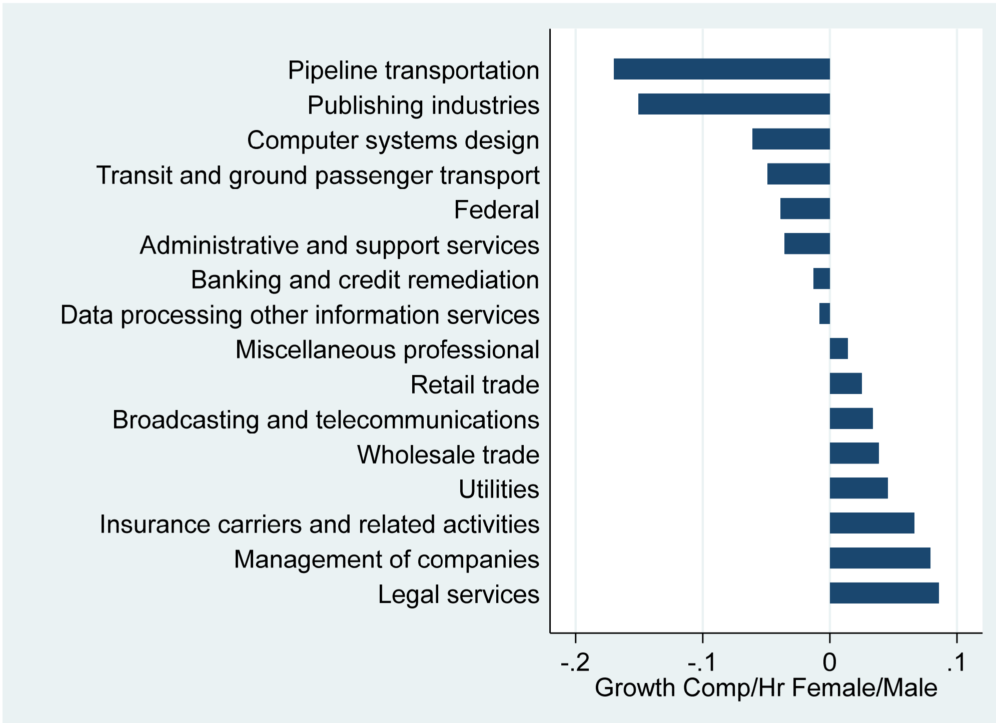

Figure 6: Relative Hours Growth of 65+ Female Workers, 2005–2019. Each point is an industry.

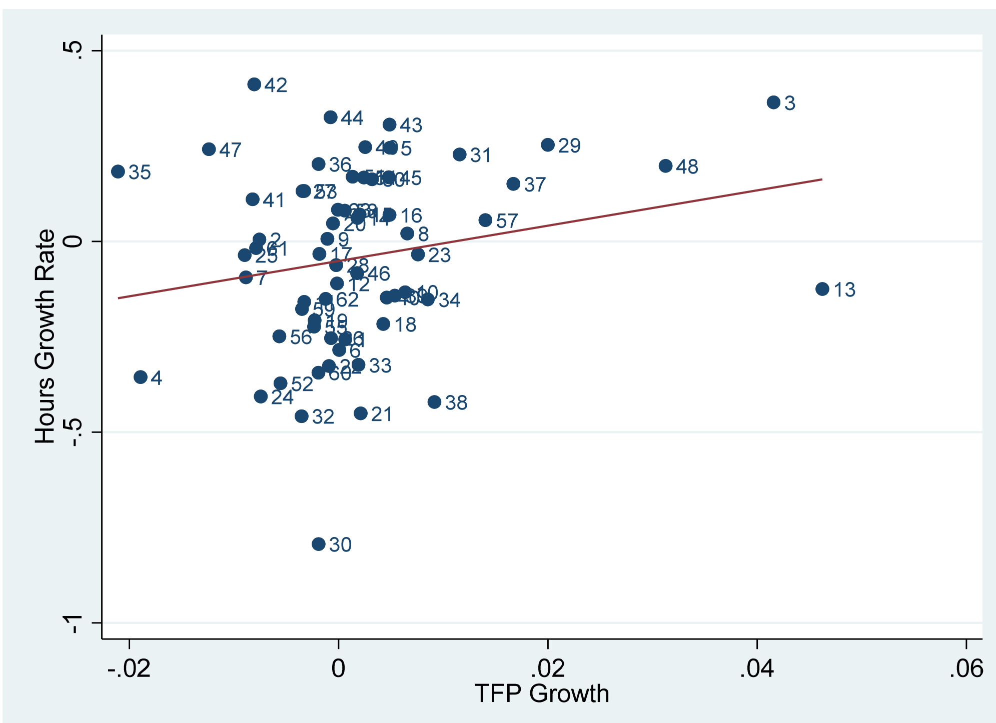

Figure 7: Relative Hours Growth of Workers with a Bachelor's Degree Relative to Workers with Some College but No Degree, 2005–2019. Each point is an industry.

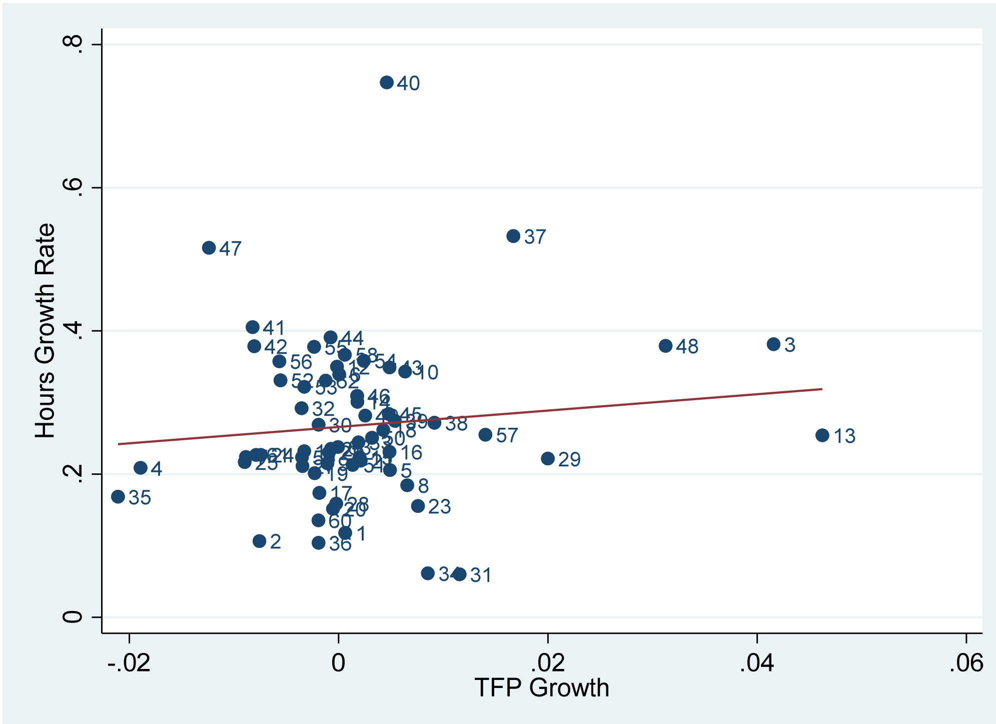

Figure 8: Growth in Compensation per Hour of Female Workers relative to Male workers 35–54 for relatively high TFP growth industries

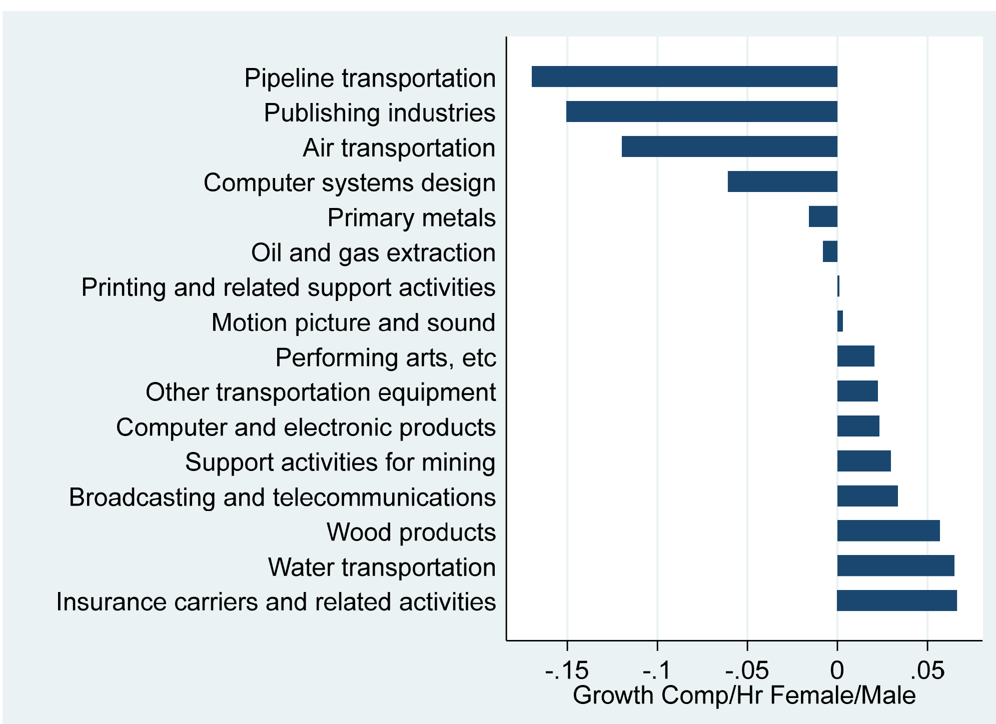

References

Acemoglu, D., Mühlbach, N. S., and Scott, A. J. 2022. The Rise of Age-Friendly Jobs. *NBER Working Paper*.

Bean, S. C. 2016. Independent Review of UK Economic Statistics.

Dingle, J., & Neiman, B. 2020. How Many Jobs can be done at Home. *NBER Working Paper*.

Fairlie, R., Couch, K., and Xu, H. 2021. The Evolving Impacts of the COVID-19 Pandemic on Gender Inequality in the U.S. Labor Market: The COVID Motherhood Penalty. *NBER Working Paper*.

Garner, C., Harper, J., Russell, M., and Samuels, J. 2021, April. New Statistics for 2019 and Updated Statistics for 1987–2018, Including Extended Capital Detail. *Survey of Current Business*.

Goldin, C. 2014. A Grand Gender Convergence: Its Last Chapter. *American Economic Review*, 1091–1119.

Goldin, C. 2022. Understanding the Economic Impact of COVID–19 on Women. *NBER Working Paper*.

Isenberg, E., Landivar, L. C., and Mezey, E. 2013. A Comparison Of Person-Reported Industry To Employer-Reported Industry In Survey And Administrative Data. *CES Working Paper Series*.

Jorgenson, D. W., Kuroda, M., and Nishimizu, M. 1987. Japan-U.S. Industry-Level Productivity Comparison, 1960–1979. *Journal of the Japanese and International Economies*.

Jorgenson, D., Ho, M., and Stiroh, K. 2005. *Information Technology and the American Growth Resurgence* . MIT Press.

Jorgenson, D., Nomura, K., and Samuels, J. 2016. A Half Century of Trans-Pacific Competition:Price level indices and productivity gaps for Japanese and U.S. industries. In D. Jorgenson, K. Fukao, & M. Timmer, *The World Economy, Growth or Stagnation?* Cambridge University Press.

Krusell, P., Ohanian, L., Rios-rull, J.-v., and Violante, G. 2000. Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis. *Econometrica*, 1029–1053.