BLS WORKING PAPERS

U.S. Department of Labor
U.S. Bureau of Labor Statistics
Office of Productivity and Technology

Job Tasks, Worker Skills, and Productivity

G. Jacob Blackwood,
Amherst College
Cindy Cunningham,
U.S. Bureau of Labor Statistics
Matthew Dey,
U.S. Bureau of Labor Statistics
Lucia Foster,
Center for Economic Studies, U.S. Census Bureau
Cheryl Grim,
Center for Economic Studies, U.S. Census Bureau
John Haltiwanger,
University of Maryland
Rachel Nesbit,
RAND Corporation
Sabrina Wulff Pabilonia,
U.S. Bureau of Labor Statistics
Jay Stewart,
Upjohn Institute for Employment Research
Cody Tuttle,
University of Texas at Austin
Zoltan Wolf,
New Light Technologies

Working Paper 585
September 2, 2025

Job Tasks, Worker Skills, and Productivity

G. Jacob Blackwood, Cindy Cunningham, Matthew Dey, Lucia Foster, Cheryl Grim, John Haltiwanger, Rachel Nesbit, Sabrina Wulff Pabilonia, Jay Stewart, Cody Tuttle, and Zoltan Wolf*

September 2, 2025

Abstract

We present new empirical evidence suggesting that we can better understand productivity dispersion across businesses by accounting for differences in how tasks, skills, and occupations are organized. This aligns with growing attention to the task content of production. We link establishment-level data from the Bureau of Labor Statistics Occupational Employment and Wage Statistics survey with productivity data from the Census Bureau's manufacturing surveys. Our analysis reveals strong relationships between establishment productivity and task, skill, and occupation inputs. These relationships are highly nonlinear and vary by industry. When we account for these patterns, we can explain a substantial share of productivity dispersion across establishments.

JEL codes: D24, J24

Keywords: productivity dispersion, tasks, skills, occupations

*Blackwood: Amherst College; Cunningham, Dey, and Pabilonia: Bureau of Labor Statistics; Foster and Grim: Center for Economic Studies, U.S. Census Bureau; Haltiwanger: University of Maryland; Nesbit: RAND Corporation; Stewart: Upjohn Institute for Employment Research; Tuttle: University of Texas at Austin; Wolf: New Light Technologies. Blackwood, Haltiwanger, and Tuttle were also part-time employees of the U.S. Census Bureau at the time of the writing of this paper. Any opinions and conclusions expressed herein are those of the authors and do not represent the views of the U.S. Census Bureau or the U.S. Bureau of Labor Statistics. The Census Bureau and the Bureau of Labor Statistics have ensured appropriate access and use of confidential data and have reviewed these results for disclosure avoidance protection (Project 7526914: CBDRB-FY24-CES020-006, CBDRB-FY25-CES007-10). We thank Enghin Atalay, Peter Meyer, Thomas Niebel, Pascual Restrepo, Jon Samuels, and Leo Sveikauskas for helpful comments and Lea Rendall for research assistance.
Corresponding author: JStewart@Upjohn.org.

I. Introduction

It is well known that measured productivity varies substantially across establishments, even within narrowly defined industries (Syverson 2004; Syverson 2011; Blackwood et al. 2021). Publicly available statistics from the Dispersion Statistics on Productivity (DiSP) [1] show that, on average, an establishment at the 90th percentile of the measured total factor productivity (TFP) distribution is nearly 2.9 times as productive as one at the 10th percentile within four-digit NAICS manufacturing industries (Cunningham et al. 2023).

Syverson (2011) reviews possible sources of productivity dispersion, including hard-to-measure factors such as managerial ability and input quality. Cunningham et al. (2023) show that firm characteristics commonly used to study firm dynamics—such as state, age class, and size class—explain little of the observed dispersion, suggesting a need to examine alternative explanations. We address this by linking establishment-level data on productivity and occupations to provide new insights into how establishments' organization of skills, tasks, and occupations contribute to productivity differences across establishments.

Standard productivity measures typically measure labor input as total hours worked, as in the DiSP data. However, accounting for variation in worker skills and the types of tasks performed may be essential for accurately measuring both productivity levels and dispersion.[2] Differences in measured productivity may partly reflect the

[1] DiSP was developed jointly by the Bureau of Labor Statistics (BLS) and the Census Bureau. See Cunningham et al. (2023) for a detailed description of the development of DiSP. DiSP is available at: https://www.bls.gov/productivity/articles-and-research/dispersion-statistics-on-productivity/ and https://www.census.gov/disp.

[2] A few empirical studies allow workers' skills to vary (Iranzo, Schivardi, and Tosetti 2008; Stoyanov and Zuanov 2022).

occupational mix and task content of the workforce (see, e.g., Acemoglu and Restrepo, 2019b).

This paper builds on a prior study (Blackwood et al. 2025) that was conducted before these establishment-level micro-datasets could be linked. That work examined, at the industry level, the relationship between dispersion in productivity and dispersion in measures of tasks, skills, and occupations within four-digit NAICS manufacturing industries. We use most of the task/skill/occupation measures developed in that study, which include two composite measures constructed using data from the OEWS survey and Occupational Information Network (O*NET),[3] five aggregate task measures derived from O*NET, and the percent of STEM workers. We also include occupation shares for two additional occupation groups—production workers and management.

The two composite measures are counterfactual mean wages that summarize establishment-level variation in the distribution of occupations, tasks, and skills. One measure captures variation in the occupational mix and is related to—but distinct from—the BLS skill-adjusted labor input used in official TFP estimates.[4] The other captures variation in the use of five aggregate tasks derived from O*NET's work activity and work-context-importance scales, as described in Acemoglu and Autor (2011): nonroutine analytical, nonroutine interpersonal, nonroutine manual physical, routine manual, and routine cognitive tasks.

[3] O*NET 12.0, 13.0, 19.0, and 22.1. U.S. Department of Labor, Employment & Training Administration (2007–2017).

[4] See https://www.bls.gov/productivity/technical-notes/labor-composition-for-total-factor-productivity-using-new-method-nov-2022.htm for a description of the official labor composition measure. For a more detailed discussion of the theory and measurement issues behind the labor composition index, see Zoghi (2007).

Blackwood et al. (2025) find that within-industry dispersion in tasks, skills and occupations are positively associated with dispersion in both labor productivity (LP) and TFP. These associations differ across types of manufacturing industries but are especially strong among high-tech industries. The remarkably high dispersion of both productivity and task/skill intensities in high-tech industries implies there is considerable heterogeneity in both productivity levels and the ways of organizing production, especially in the most innovative sectors of the economy.

This prior analysis was limited to industry-level variation and did not investigate establishment-level relationships. This paper overcomes that limitation by linking occupational data from the Occupational Employment and Wage Statistics (OEWS) survey to establishment productivity data from the Collaborative Micro-productivity Project (CMP), resulting in the CMP-OEWS dataset that we use in this study.[5]

This linked dataset allows for novel establishment-level analyses. We start by showing that adjusting total hours using a simple scalar measure of task/skill/occupation content has little effect on measured productivity dispersion. While that might seem surprising, the initial analysis assumes a (log) linear and uniform relationship across establishments within the same industry—a restrictive assumption. Relaxing this assumption, we find that these relationships are much more complex. The relationship between these task/skill/occupation measures and productivity turns out to be highly nonlinear, with the strongest relationships occurring at the extremes—among establishments in the top and bottom quintiles of the task/skill and occupation

[5] A restricted-access dataset, the Collaborative Micro-productivity Project (CMP) (Bureau of Labor Statistics and Census Bureau 1972–2020), is available for use by qualified researchers on approved projects in the Federal Statistical Research Data Centers (https://www.census.gov/fsrdc).

measures. These nonlinear patterns are even more pronounced among larger establishments.

Motivated by these findings, we re-examine within-industry variation in the OEWS task, skill, and occupation measures. Two key findings emerge from this analysis. First, most of the variation in broad occupation shares occurs within, not between, industries. Second, within-industry variation in occupation shares is especially large in the top and bottom quintiles of the task/skill distribution.

Finally, we examine the role of our task/skill/occupation measures in explaining the variation in TFP. While not causal, the analysis shows that a substantial portion of establishment-level TFP dispersion within industries can be accounted for by differences in occupational mix, tasks, and skills. These relationships are nonlinear, vary by industry, and are particularly strong among larger establishments and those in high-tech sectors.

The rest of the paper is organized as follows. Section II describes how we construct the occupation, task, and skill measures. Section III discusses our data sources and the matching procedure. Section IV presents our main results, which illustrate the complex relationship between TFP and the task, skill, and occupation measures. Section V concludes and outlines directions for future research.

II. Measuring occupations, tasks, and skills

We start by defining some basic concepts.[6] **Tasks** are activities that when combined with intermediate goods create a good or service. Tasks can be performed by

[6] These descriptions are based on the nomenclature from the Revised Handbook of Analyzing Jobs (U.S Department of Labor, Employment and Training Administration 1991) and Acemoglu and Autor (2011).

capital equipment, such as robots, or by workers. Our focus is on tasks done by workers. Because we do not observe time spent on different tasks, we use occupations as proxies. An **occupation** can be thought of as a bundle of tasks.

Skill refers to a worker's ability to perform various tasks and is commonly measured as a function of education and experience. However, due to data constraints, we use wages to proxy for occupational skill requirements.[7] Complex tasks generally require greater skills, although the relationship between skills and tasks can vary over time and across businesses. That presents a challenge for productivity measurement and highlights a need for more detailed information on tasks and skills.

We now turn to defining our two composite measures, five task measures, and three broad occupation groups (see also Blackwood et al. (2025)).

Bundled Task/Skill Index (TSB): Counterfactual Wages

Our first composite measure of task/skill intensity is a counterfactual wage equal to the average wage an establishment would pay if it paid each worker the average national wage for their occupation. This measure accounts for differences in the occupational mix across establishments by attaching a different wage to each occupation. By using the average national wage, the price of every occupation is the same in all establishments. We refer to this as a "bundled" task/skill index (TSB) because tasks are bundled into occupations.

[7] We believe that mean occupation wages are a reasonable single-dimensional proxy for skills. As a check, using the 2011–2019 Current Population Survey Outgoing Rotation Groups (Economic Policy Institute 2025), we estimate three separate regressions of log wage on 1) a quartic in age and 4 education groups, 2) detailed Census occupation groups, and 3) a quartic in age, 4 education groups, and detailed Census occupation groups. The R-squared values are 0.32, 0.38, and 0.46, respectively, suggesting that occupation on its own explains much of the variation in wages due to skill variation.

More formally, let $\overline{w}_{ej}$ and L_{ej} denote the mean log wage and the number of workers in occupation *j* at establishment *e*. Suppressing time subscripts for simplicity, the average national wage for occupation *j* is:

$$\overline{w}_{nj} = \frac{1}{\sum_{e \in E_n} L_{ej}} \sum_{e \in E_n} (\overline{w}_{ej} \times L_{ej}) \tag{2}$$

where E_n is the set of all establishments. The counterfactual mean log wage for establishment *e*, $\widetilde{w}_e$, is then:

$$\widetilde{w}_e = \frac{1}{L_e} \sum_{j \in J_e} (\overline{w}_{nj} \times L_{ej}) \tag{3}$$

where J_e is the set of all occupations in establishment *e* and L_e is total employment in establishment *e*.

TSB is a simple measure that summarizes the types of tasks employed in an establishment using wages as a proxy for skills to price those tasks. Because TSB is based on occupation-specific average national wages, cross-establishment differences in this measure reflect variation in the occupation mix. However, this measure does not distinguish between different occupations (with different task sets) paying the same wage. Two establishments might have the same TSB but quite different mixes of occupations.[8]

Unbundled Task/Skill Index (TSU): Task-Adjusted Counterfactual Wages

Our second composite measure focuses on tasks and follows Acemoglu and Autor (2011), who use O*NET data to operationalize the Autor, Levy, and Murnane

[8] Blackwood et al. (2025) illustrate this point by plotting TSB against a dissimilarity index that quantifies how the occupational mix of the establishment differs from the average occupation mix of its four-digit industry. The dissimilarity index measures the absolute value of the sum over all occupations (two-digit Standard Occupational Classification (SOC)) of the distances between the establishment's payroll share for that occupation and the industry-wide payroll share for that occupation. It takes on values between zero and one, with larger values indicating that the occupational distribution in an establishment differs substantially from the typical establishment in the industry.

(2003) taxonomy of tasks. Autor, Levy, and Murnane (2003) develop a two-dimensional categorization of tasks based on whether they are (1) routine or non-routine and (2) cognitive or manual. They subdivide non-routine cognitive tasks into analytical and interpersonal tasks. This leads to five task categories: non-routine cognitive (analytical), non-routine (interpersonal), non-routine manual physical, routine cognitive, and routine manual.[9]

We use O*NET data to create these five task indexes for each occupation for the four years for which the index variables are available for most occupations (2007, 2008, 2014, and 2017).[10] We then merge these five task indexes to the OEWS by occupation and estimate the following regression of the average national wage on the five task indexes for each occupation:

$$\overline{w}_{nj} = \alpha + \sum_{k=1}^{5} \beta_k \tau_{jk} + \varepsilon \tag{4}$$

where τ_{jk} is the O*NET measure of task *k* for occupation *j*, and $\overline{w}_{nj}$ is the average wage for each occupation as defined in equation (2).[11] The coefficients on the task indexes, β_k, are akin to prices in a hedonic regression. We calculate the counterfactual average establishment wage as:

$$\widehat{w}_e = \frac{1}{L_e} \sum_{k=1}^{5} \hat{\beta}_k \left[\sum_{j \in J_e} \left(L_{ej} \times \tau_{jk} \right) \right] \tag{5}$$

where the summation in square brackets is the total amount of task *k* employed by the establishment and $\hat{\beta}_k$ is the estimated "price" of task *k* estimated from equation (4). That

[9] Acemoglu and Autor (2011) include a sixth category, offshorability, which we do not include here because it is not a task.

[10] We match two prior years of OEWS data to a given O*NET year to obtain the employment weights. When an occupation is covered in both OEWS years, we average the two years; otherwise, we take the value for the one OEWS year with coverage for that occupation. Thus, the 2007 O*NET is matched to 2005 and 2006 OEWS; 2008 O*NET to 2006 and 2007 OEWS; 2014 O*NET to 2012 and 2013 OES; and 2017 O*NET to 2015 and 2016 OEWS.

[11] We first aggregate occupations to a time-consistent SOC classification.

is, the TSU measure can be thought of as the average price of the tasks performed by employees in the establishment.

We refer to this second measure as an "unbundled" task/skill index (TSU) because tasks (weighted by prices) are aggregated without accounting for how they are bundled into occupations. As with the TSB index, many combinations of tasks can result in the same TSU value.

Both TSB and TSU reflect task/skill differences across establishments and account for the prices of those tasks in the labor market, where prices reflect the skills required to accomplish those tasks (among other things that determine wages). The major difference between these two measures is that TSB reflects how the tasks are organized into occupations, indirectly accounting for complementarities between tasks that make up an occupation and the benefit of having them performed by the same person, while TSU prices the tasks individually and ignores any complementarities between tasks within occupations.

Individual Average Task Indexes

In addition to TSB and TSU, we also construct five task measures based on the average values of the individual O*NET task indexes. For each of the five task indexes described above, we measure an employment-weighted establishment-level average for task index *k* as follows:

$$\bar{\tau}_{ek} = \frac{1}{L_e} \sum_{j \in J_e} \tau_{jk} \times L_{ej} \tag{6}$$

where *k* = 1, …, 5. Thus, $\bar{\tau}_{ek}$ is the average task *k* content of all jobs in establishment *e*. We construct these measures for each establishment for each year in our sample.

Occupation Groups

In addition to the measures described above, we also perform analysis with three major occupation groups: Production Workers, STEM Occupations, and Management, as defined by the Standard Occupational Classification (SOC) codes in the OEWS data.[12] Our definition of production workers focuses on workers who do actual production (including material moving) and is therefore narrower than the ASM definition, which includes occupations that are not directly involved with production.[13] The broader task-based definition of production work in the ASM likely means workers in STEM occupations sometimes fall under the production worker definition (e.g., under product development), while most workers in STEM occupations fall into the nonproduction category. Management occupations are included in ASM total employment, but not in the production worker count.

III. Data and matching

This section describes the two datasets we use, the CMP data and the OEWS occupation data, and how we match them.

[12] Production Workers include Production Occupations (51-0000) and Material Moving Workers (53-7000). STEM Occupations include Computer and Mathematical Occupations (15-0000), Architecture and Engineering Occupations (17-0000), and Life, Physical, and Social Science Occupations (43-5000). Management includes Management Occupations (11-0000) and Business and Financial Operations Occupations (13-0000).

[13] The ASM definition includes workers engaged in fabricating, processing, assembling, inspecting, receiving, packing, warehousing, shipping (but not delivering), maintenance, repair, janitorial, guard services, product development, auxiliary production for plant's own use (e.g., power plant), recordkeeping, and other closely associated services (including truck drivers delivering ready-mixed concrete). It also includes first-line supervisors.

CMP Data

As part of the CMP, BLS and the Census Bureau created an establishment-level productivity database for the manufacturing sector (Cunningham et al. 2023). Data on inputs and output are from the Annual Survey of Manufactures (ASM) and the Census of Manufactures (CM) and are longitudinally linked using information from the Longitudinal Business Database (LBD), which is based on the Census Bureau's Business Register (Chow et al. 2021). The LBD provides high-quality longitudinal links and information on the universe of active nonfarm private sector employer establishments. The ASM is collected annually and is a five-year panel of manufacturing establishments updated by births in each year.[14] The CM collects data from all manufacturing establishments, except those that are very small, every five years.[15] In this paper, our CMP dataset covers 2001–2020 and consists of approximately 999,000 establishment-year observations.

Before we could match OEWS data to the CMP data, it was necessary to address some disagreements between the CM/ASM data and the LBD. Because production functions are calculated for each industry, the most relevant issue is disagreements in industry codes, which can occur for several reasons. For example, ASM and CM industry codes are based on responses collected at the time the survey was conducted, whereas LBD codes are updated with a lag. Consequently, an establishment that changes its industry may appear in different industries in the LBD, the ASM, and the CM. Thus, there could be up to three different "raw" industry codes in the CMP data: the LBD industry code, the CM industry code, and the ASM industry

[14] ASM panels start in years ending in "4" and "9."
[15] The CM is collected in years ending in "2" and "7."

code. Changes to the NAICS industry classification system over time and the differential timing of those changes in the various datasets are another potential problem. For example, the transition from the 2007 to the 2012 NAICS codes reduced the number of manufacturing industries substantially, from 473 to 364 six-digit industries. The LBD provides an additional longitudinally consistent industry code—the vintage consistent (VC) industry code (Chow et al. 2021), which is a single vintage of NAICS (in this case, the 2017 vintage) and is extended backward so that industry codes are consistent with the 2017 classification system throughout the sample period. The VC process sometimes entails consolidation of codes or imputation based on establishment characteristics.

Because an establishment's industry code is an integral part of our matching process, it is important to have the best possible chance of matching an establishment's NAICS code to the code that the BLS assigns in the OEWS. Accordingly, we use four different NAICS codes in our CMP dataset: the LBD code, the ASM code (which is available for manufacturing observations only), the VC code, and a "combined" code created by combining information from the CM, ASM, and LBD. Appendix A provides details of this procedure.

Occupational Employment and Wage Statistics (OEWS) Data

Our occupation data come from the OEWS survey, a semi-annual survey of approximately 200,000 establishments conducted each May and November.[16] This

[16] From 1999 to 2001, the program surveyed approximately 400,000 establishments in November of each year. Starting in November 2002, the program switched to semi-annual sampling with 200,000 establishments sampled each May and November. To keep sample sizes roughly consistent across the various years, we combine November and May panels to create a pseudo-annual sample, with November establishments being assigned the year in which the May establishments were surveyed. For this reason, we do use data for 2002.

survey covers wage and salary workers in private, nonfarm industries. Employer Identification Numbers (EINs) and NAICS codes come from the BLS Quarterly Census of Employment and Wages (QCEW), which is the sample frame for the OEWS survey.

The survey asks establishments to provide a complete payroll record for the pay period that includes the 12th of the sample month. For each occupation, respondents report the number of full- and part-time employees in each of 12 wage intervals.[17] The OEWS survey uses the Office of Management and Budget's Standard Occupational Classification (SOC) to categorize workers into over 800 detailed occupations.

The sample contains both certainty and non-certainty units. The former are generally sampled every three years, while the latter are selected randomly and tend to be smaller establishments. Given this sample design, six consecutive panels can be used to create a representative sample that covers every three-year period.[18] We use this aspect of the sampling scheme in our matching process as described below.

We make the same time-consistent adjustment to the OEWS industry codes as we make to the LBD and the ASM industry codes. This results in two versions of the OEWS survey NAICS codes: one is the original version, and the other is a time-

---

[17] Wages in the OEWS data represent straight-time, gross pay, exclusive of premium pay. Base rate, cost-of-living allowances, guaranteed pay, hazardous-duty pay, incentive pay including commissions and production bonuses, tips, and on-call pay are included, while back pay, jury duty pay, overtime pay, severance pay, shift differentials, non-production bonuses, employer cost for supplementary benefits, and tuition reimbursements are excluded from the reported wage. For a description of the wage intervals, see https://www.bls.gov/oes/mb3-methods.pdf.

[18] Official estimates are typically published in May. These estimates are based on data from the May panel and the previous five panels. Note that although official estimates are published, they are not a true time series. In year-to-year comparisons of consecutive years, data from approximately 2/3 of units appear in both years. For these units, the wages are updated using the Employment Cost Index, but employment counts are not adjusted.

consistent version where some six-digit industry codes are aggregated into quasi-five-digit codes.

Linking OEWS and CMP Data

Linking OEWS data to the CMP is not straightforward because the establishment identifiers are not the same in the two datasets. Both datasets include the taxpayer ID (EIN) and the industry code for each establishment. However, the EIN does not necessarily correspond to the Census definition of an enterprise, which depends on operational control. A Census enterprise may comprise multiple EINs. For each establishment in our augmented CMP sample, our goal is to identify the best candidate in the OEWS survey, where the best candidate is defined based on the EIN, NAICS code, geography (state FIPS code), and size (as measured by employment). Loosely speaking, a match occurs if the values of these variables are the same for any two records in the two datasets.

EIN-based matches are, in principle, exact for single-unit firms. However, even among single-unit EINs, matches may not be exact for several reasons. First, the two business registers have slightly different criteria for classifying establishments according to single- or multi-unit status.[19] This implies that a single-unit CMP establishment may have multiple candidates in the OEWS survey that share the same EIN. Second, NAICS codes may differ because of differences between the BLS and Census Bureau business registers. This can occur because the two agencies use slightly different criteria for classifying establishments into industries. Third, there can be temporal mismatches in

[19] This discrepancy occurs in part because the timing of single-unit growth into multi-unit, or of multi-unit contraction into single-unit, can be difficult to infer (Chow et al. 2021).

the data collected for the establishment, because the two surveys may be conducted at different times and for different reference periods.[20] The OEWS survey sampling frame produces a representative sample with every three years of OEWS surveys. We use establishments from three years of OEWS data as possible matches for establishments in each single CMP year, where the three years of OEWS data are centered on the CMP year. For example, establishments in the 2014–2016 OEWS provide possible matches for the 2015 CMP.

We require OEWS establishments to match on EIN and be similar in size in every step of our matching procedure. We measure size using employment and calculate the "employment difference" as $|L_{CMP} - L_{OEWS}|/((L_{CMP} + L_{OEWS})/2)$ where L_{CMP} is employment in the CMP and L_{OEWS} is employment in the OEWS. Our matching procedure is hierarchical in that we prioritize potential donors that match on the most detailed information on industry and geography. We begin with the most stringent criteria and successively relax them.

For establishments with 100 or more employees, the matching criteria are:

(1) EIN, six-digit industry, state, employment difference less than 0.5
(2) EIN, time-consistent six-digit industry, state, employment difference less than 0.5
(3) EIN, six-digit industry, employment difference less than 0.5
(4) EIN, time-consistent six-digit industry, employment difference less than 0.5
(5) EIN, four-digit industry, employment difference less than 0.5

For smaller establishments (those with fewer than 100 employees), we modify the employment-threshold criterion so that it depends on the length of time between

[20] The reference periods for the ASM and OEWS survey can differ by up to 18 months.

observations in the OEWS and ASM.[21] To calculate these modified thresholds, we start with single units whose EINs match exactly. Using these establishments, we calculate the 90th percentile of the absolute difference in employment, $|L_{CMP} - L_{OEWS}|$, allowing it to vary by the time difference between when establishments are sampled in the OEWS and CMP. CMP contains employment values for the week of March 12, while the OEWS is conducted in May and November. Because we allow OEWS establishments in adjacent years to be donors, the length of time between samples can be 2, 4, 8, 10, 14, or 16 months. Using single-unit establishment matches, we calculate the 90th percentile of the absolute difference in employment for each time length across establishment size categories: 0–4, 5–9, 10–19, 20–49, and 50–99. Rules for matching then follow the previous hierarchy with slight modifications:

(1) EIN, six-digit industry, state, employment difference less than 0.5 OR absolute employment difference <= p90 (time-varying absolute employment difference)
(2) EIN, time-consistent six-digit industry, state, employment difference less than 0.5 OR absolute employment difference <= p90 (time-varying absolute employment difference)
(3) EIN, six-digit industry, employment difference less than 0.5 OR absolute employment difference <= p90 (time-varying absolute employment difference)
(4) EIN, time-consistent six-digit industry, employment difference less than 0.5 OR absolute employment difference <= p90 (time-varying absolute employment difference)
(5) EIN, four-digit industry, employment difference less than 0.5 OR absolute employment difference <= p90 (time-varying absolute employment difference)

[21] We introduce this criterion because small employment changes in such establishments can cause large percentage changes.

In both cases (1–99 and 100+ employees), criterion (1) begins with the original six-digit NAICS codes, whereas (2) relies on the time-consistent codes described in Section III.A, which are slightly less detailed than the original six-digit codes in some instances because of aggregation when the NAICS vintages differ. In (3), we return to our original six-digit codes but relax the geographic requirement, and (4) repeats (3) but instead uses the time-consistent codes. Finally, (5) allows for matches with four-digit industry codes (as well as EIN and size, as in all cases). As mentioned in Section III.A, multiple industry codes may be used in the matching procedure. Therefore, in each step, we iterate over three industry codes: beginning with the combined NAICS code, then using the ASM NAICS code if we find no potential donors with the combined code, then finally the LBD NAICS code.

In many cases, multiple potential donors in the OEWS satisfy the same criteria for a match. When this occurs, we break ties by choosing the donor closest in size to the CMP establishment. When multiple donors are the same size, the second tiebreaker is to choose the donor closest in time to the CMP establishment.[22] Finally, in cases where both the employment and time difference are both the same, we randomly choose a donor from among those that meet all the criteria. Appendix A provides an example that illustrates these steps.

This matching process ensures that each donor from the OEWS has at least the same EIN, four-digit industry, and size as its CMP recipient. The resulting subsample of the CMP contains matched OEWS information on the occupation distribution. We believe our current approach balances match quality with sample size requirements.

[22] For each year of CMP establishments, we consider potential matches from three years of OEWS survey establishments.

Final Analysis Sample

Our final matched dataset consists of approximately 333,000 establishment-year observations covering 2001 through 2020 that contain information on both the occupation distribution and productivity. This is about 33% of the original CMP observations. Our analysis sample incorporates the following modifications. First, we use the 2017 vintage-consistent NAICS code from the LBD instead of the codes used in matching, because we need consistent codes over time to allow us to remove industry-year effects by demeaning.[23] Second, we create inverse propensity score weights (PW), because the CMP-OEWS data are not a representative sample of all manufacturing establishments. We estimate a logistic regression, where the covariates are industry, establishment size, and payroll, to generate the predicted probability of each LBD establishment being included in the linked dataset. The inverses of the predicted values from this regression are the PW.[24]

Table 1 presents descriptive statistics for the OEWS, the CMP, and the matched sample (CMP-OEWS). Differences in the mean and standard deviation of employment help us illustrate the different sample characteristics across the three datasets: employment moments are largest in the unweighted linked data, much smaller in the PW-weighted CMP data, and still smaller in the sample-weighted OEWS survey data. But employment moments for the linked data are broadly similar to those of unmatched data when we apply the PW weights.

[23] The vintage-consistent industry codes used are the same industry codes used to create the publicly available DiSP.
[24] Details of the PW construction can be found in Cunningham et al. (2023).

Table 1. Descriptive statistics of employment in manufacturing

Summary statistic	OEWS (weighted)	CMP (PW)	CMP-OEWS (unweighted)	CMP-OEWS (PW)
Mean	32.5	53.1	221.5	61.9
Standard deviation	167.5	204.8	543.1	253.6

Notes: OEWS survey weights account for the probability of selection, the fact that six panels of data are combined to form the full sample, and differences in employment totals between the sample and the QCEW frame. PW refers to inverse propensity weights. N (in 1000s): 593 (OEWS), 999 (CMP), and 333 (CMP-OEWS). CMP refers to the combined ASM, CM, and LBD data. CMP-OEWS refers to the matched dataset. Source: Authors' calculations based on OEWS and CMP.

To provide a better sense of the distributional differences between the datasets, Table 2 shows standard deviations of the demeaned variables used in the analyses. Employment dispersion after demeaning exhibits the same patterns as those in Table 1. Task/skill/occupation variation in the linked sample is smaller relative to the OEWS survey sample size without PW but becomes closer with PW. Dispersion in measures of productivity, earnings-per-worker, and capital intensity are similar in the CMP and matched data.

Table 2. Standard deviations of key variables

Key variables	OEWS (weighted)	CMP (PW)	CMP-OEWS (unweighted)	CMP-OEWS (PW)
Employment	164.2	198.6	507.6	239.8
Analytical	0.435	-	0.3118	0.3320
Interpersonal	0.501	-	0.3190	0.3866
Physical	0.510	-	0.3802	0.4331
Routine cognitive	0.501	-	0.3189	0.3885
Routine manual	0.718	-	0.4800	0.5486
TSU	0.153	-	0.1078	0.1147
TSB	0.189	-	0.1392	0.1486
Log(TFP)	-	0.4808	0.4952	0.5073
Log(LP)	-	0.7472	0.7283	0.7461
Production worker share	-	-	0.2022	0.2407
STEM share	-	-	0.0777	0.0734
Management share	-	-	0.0684	0.0885

Notes: OEWS survey weights account for the probability of selection, the fact that six panels of data are combined to form the full sample, and differences in employment totals between the sample and the QCEW frame. PW refers to inverse propensity weights. Industry-year effects are removed. N (in 1000s): 593 (OEWS), 999 (CMP), and 333 (CMP-OEWS). CMP refers to the combined ASM, CM, and LBD data. Source: Authors' calculations based on OEWS and CMP.

Tables 1 and 2 show that the matched dataset retains the basic properties of the OEWS and CMP data, which gives us confidence in the results of the analysis that follows.

IV. Relationship between productivity and occupations, tasks, and skills

We start our analysis by making a simple adjustment to labor hours. We then allow for more complex relationships. Through a series of empirical exercises, we highlight the importance of allowing for more complex relationships when trying to understand the connection between productivity and occupations, tasks, and skills.

A Simple Dispersion Exercise

We start by making a simple multiplicative adjustment to the labor input used in the DiSP data (total hours) as shown in equation (7):

$$\log TFP_{et} = logQ_{et} - \alpha_K logK_{et} - \alpha_L \log(Z_{et}L_{et}) - \alpha_M logM_{et} \quad (7)$$

where TFP is total factor productivity, *Q* is real output measured as deflated revenues, *K* is real productive capital stock, *M* is the deflated value of expenditures on intermediate inputs (materials, resales, contract work, electricity, and fuels), *Z* is a normalized version of TSB (or TSU), *L* is total hours, and *e* and *t* index establishments and time. The parameters $\alpha_K, \alpha_L,$ and α_M are factor elasticities measured by the share of expenditures of each input in total cost in each six-digit NAICS industry. For more details on the construction of these variables, see Cunningham et al. (2023).

In constructing *Z*, we normalize TSB (TSU) so it has a mean of one in each industry-year cell. We calculate mean TSB (TSU) for each four-digit industry × year cell, then divide each establishment's TSB (TSU) by the industry-year mean. To adjust total hours, we multiply total hours by this normalized measure of TSB (TSU), yielding a labor input measure that incorporates task-skill intensity in a simple manner.

Table 3. Accounting for skill/task intensity in log TFP dispersion—Average IQR (CMP-OEWS sample)

Labor input	Propensity weighted	Activity weighted
Total Hours	0.456	0.496
Total Hours × TSB	0.456	0.494
Total Hours × TSU	0.456	0.494

Source: Authors' calculations based on OEWS and CMP.

Table 3 shows the average interquartile range (IQR) over our sample period for the distribution of log TFP using unadjusted hours data and hours data to efficiency units using TSB and TSU. Column 1 shows results using only PW weights. In column 2, we report results using activity weights (AW), where AW = PW × employment. Regardless of adjustment or weighting, this simple approach to convert labor input into efficiency units does not reduce measured dispersion.

This simple one-dimensional adjustment to account for skill and tasks does not take into account the potentially rich interaction of tasks and other inputs in the type of task content of production approach to productivity. Establishments that are organized differently likely have different production technologies that are not well captured by this simple Cobb-Douglas function with only a one-dimensional multiplicative adjustment of the labor input.[25] We are not prepared to implement an approach along the lines of Acemoglu and Restrepo (2019b) in this measurement-oriented paper. Instead, we explore the relationship between productivity and the occupational and task mix in a multi-dimensional, nonlinear manner within industries.

Importance of Nonlinearities

We now examine the potential for nonlinearities in the relationship between TFP and measures of skills, tasks and occupations. After sweeping out industry-year effects, we compute percentiles of each task/skill and occupation measure. In turn, we compute the average (log) TFP for each percentile. For disclosure avoidance reasons, we report the results as a quartic relationship relating average TFP to the percentile ranking of the task/skill/occupation measure. We conduct these exercises separately using PW and

[25] This simple exercise imposes the same factor elasticities across all establishments in the same industry.

AW. Comparing the results using the two weighting methods gives us insight into the differences between large and small establishments.

Figure 1. Relationship between log TFP and TSB

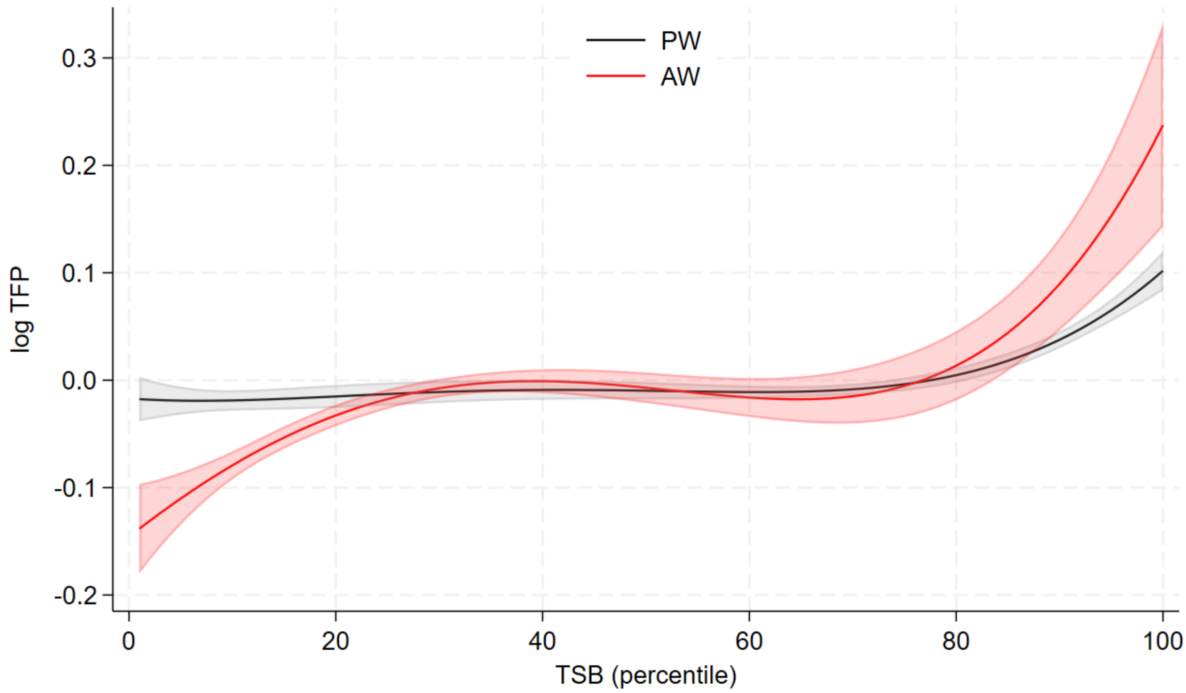

Notes: N = 333,000 establishment-years. The label “PW” refers to inverse propensity score weights and label “AW” refers to activity weights. Source: Authors’ calculations based on the OEWS survey, O*NET, and CMP.

Figure 1 shows the relationship between TFP and TSB (percentiles). Notably, there is a highly nonlinear relationship, especially when we weight by employment (AW). In the top TSB quintile, TFP rises rapidly with TSB using either weighting method, but the increase is much faster using AW. In the AW quartic, TFP rises rapidly over the lowest 20 percentiles as well. Together, these imply that much of the nonlinearity is due to the stronger relationship between TFP and TSB among large establishments.

Figure 2 shows analogous relationships between TFP and the O*NET analytical task measure (left panel) and the routine manual task measure (right panel). The results for analytical tasks in Figure 2 mimic those for TSB. The patterns for routine manual tasks are the mirror image, with TFP declining with routine manual over the bottom and top quintiles, but with a relatively flat relationship in the middle of the distribution. Again, these figures imply that large establishments are driving the nonlinearities in the tails.

Figure 2. Relationships between log TFP and ONET analytical and routine manual shares

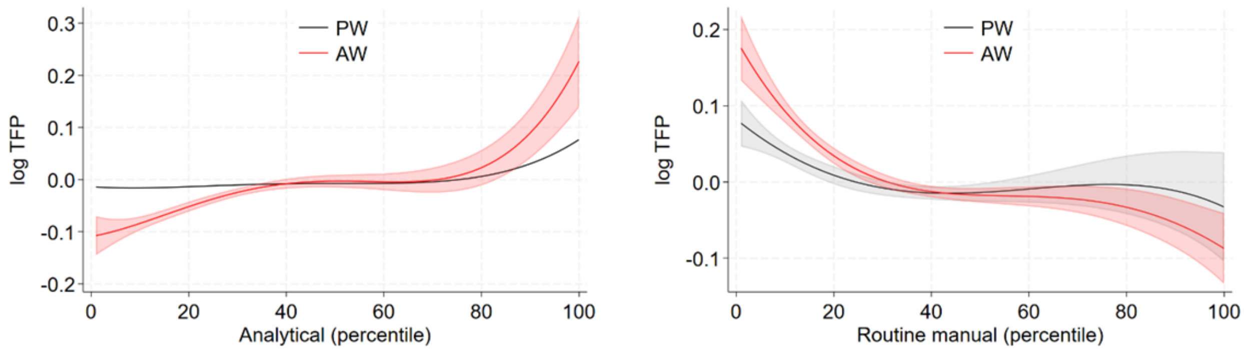

Notes: N = 333,000 establishment-years. The label “PW” refers to inverse propensity score weights and the label “AW” refers to activity weights. Standard error bands for the PW quartic in left panel are not available for disclosure reasons. Source: Authors’ calculations based on the OEWS survey, O*NET, and CMP.

Turning to occupation groups in Figure 3, we see similar nonlinear relationships between TFP and occupational shares. The STEM occupational share is positively related to TFP, especially in the tails of the distribution—with evidence for a relationship in the lower tail for larger establishments. The mirror image holds for the production worker share at establishments.

The nonlinear patterns in Figures 1–3 suggest complex relationships between measured TFP and the task/skill and occupational mix across establishments. We explore these relationships further below. But we first take a closer look at the relationship between TSB and occupations both between and within industries.

Figure 3. Relationships between log TFP and STEM and production worker shares

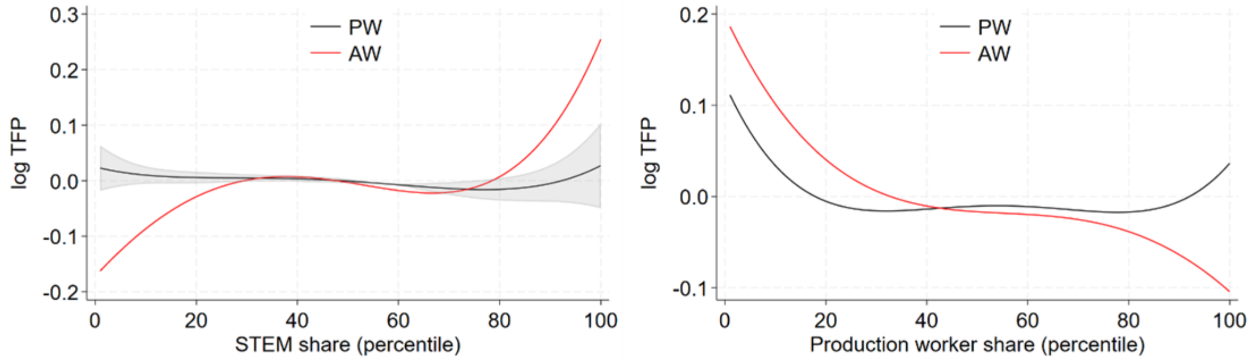

Notes: N = 333,000 establishment-years. The label “PW” refers to inverse propensity score weights and the label “AW” refers to activity weights. Standard error bands for the AW quartic in left panel and for both quartics in the right panel are not available for disclosure reasons. Source: Author’s calculations based on the OEWS survey, O*NET, and CMP.

A Closer Look at Occupations

Given that the positive relationship between TFP and TSB shows up mainly in the top and bottom TSB quintiles, we start by looking at occupation shares by TSB quintile. We then examine the sources of variation in employment shares by quintile.

For these exercises, we employ the OEWS research dataset used in Blackwood et al. (2025). This dataset includes all establishments in the QCEW. Occupation data from the OEWS are used for OEWS respondents and are imputed for establishments not in the OEWS.[26] The advantage of this approach over using only OEWS sample data is that each establishment has a weight of one.[27]

[26] See Appendix A for a brief description of the dataset construction. See Blackwood et al. (2025) and Dey, Piccone and Miller (2019) for detailed descriptions.
[27] The sample design of the OEWS survey is geared toward measuring employment, which complicates the construction of establishment weights.

Figure 4. Mean occupation share of employment by TSB quintile

A. Production worker

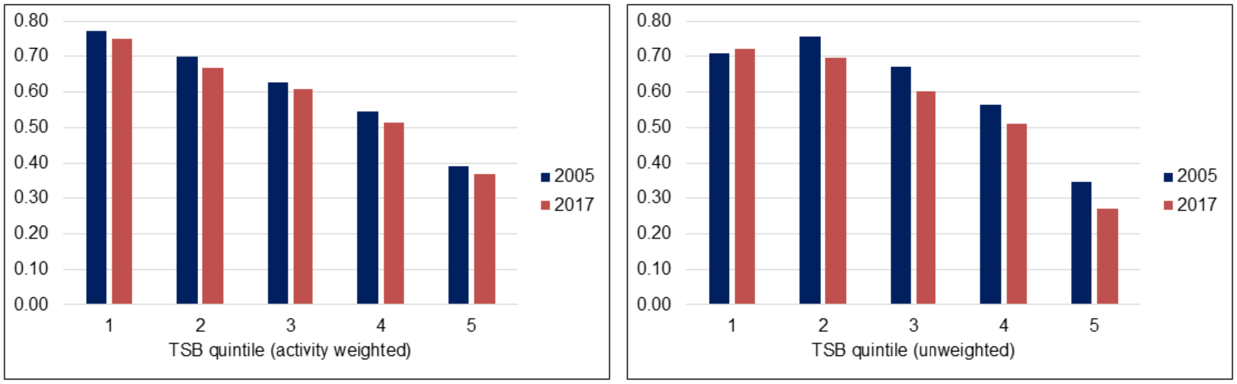

B. STEM worker

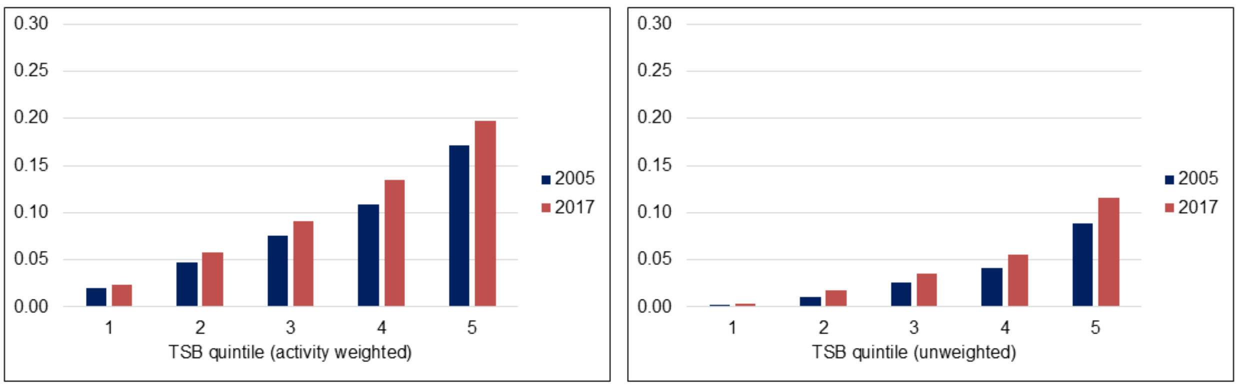

C. Management

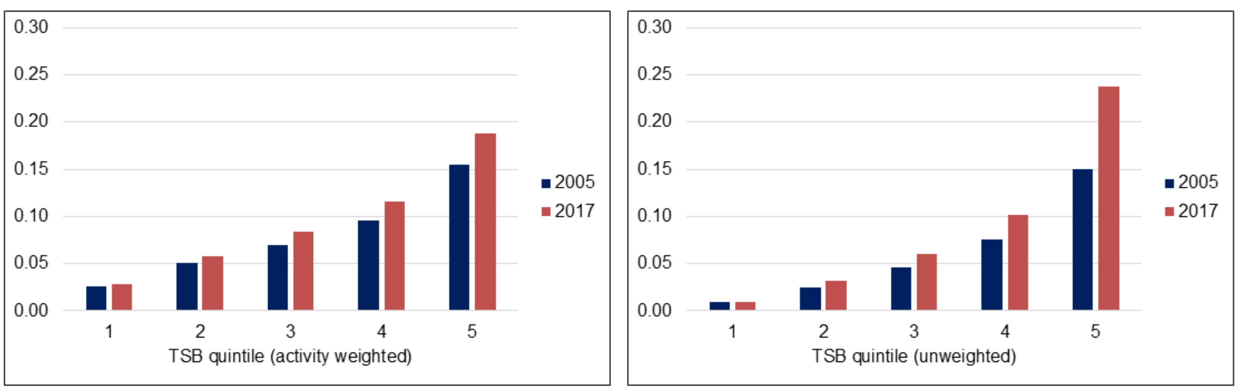

Source: OEWS-QCEW research dataset, authors' calculations.

Figure 5. Within-industry variation in occupation share of employment by TSB quintile

A. Production worker

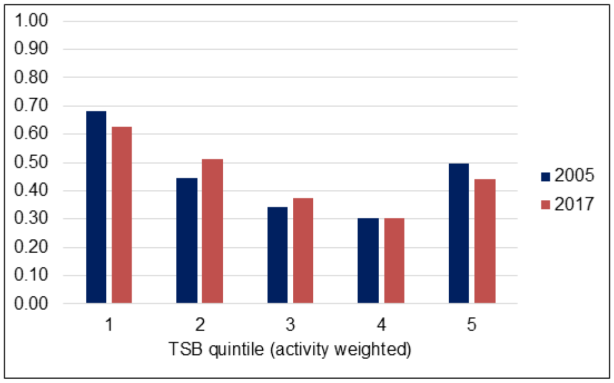

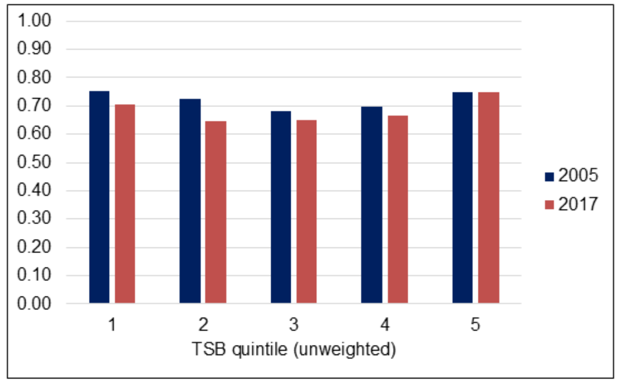

B. STEM worker

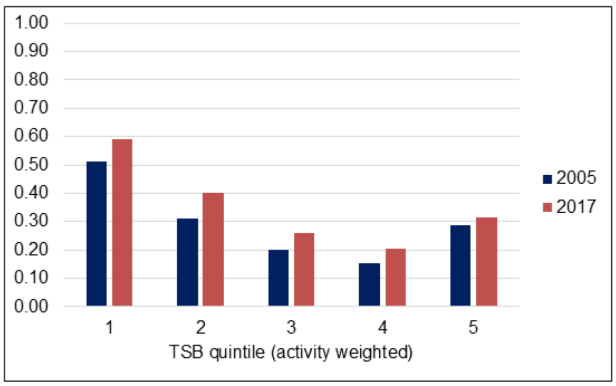

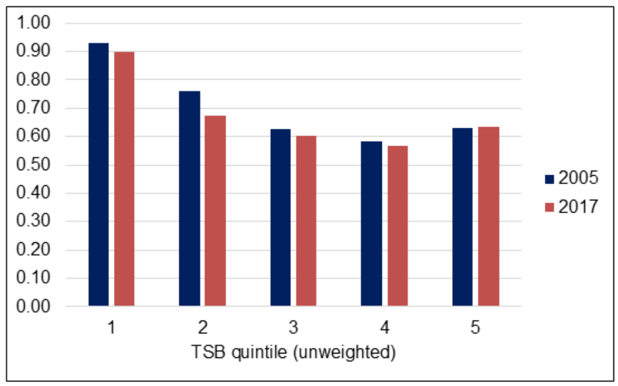

C. Management

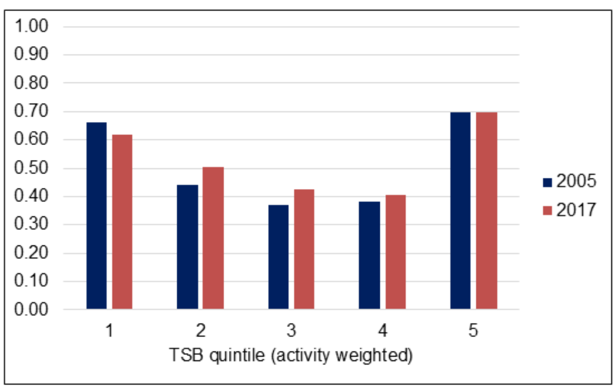

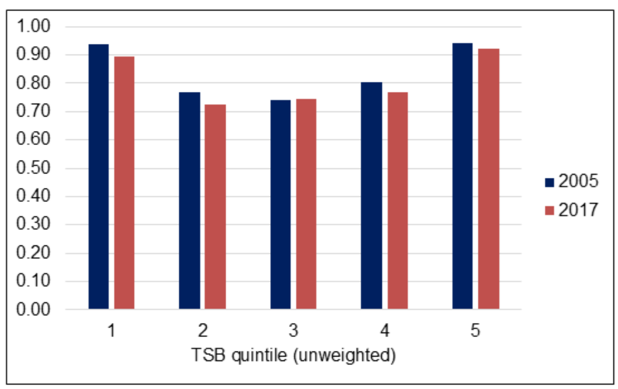

Source: OEWS-QCEW research dataset, authors' calculations.

The six panels in Figure 4 show the mean occupation shares of production workers, STEM workers, and management by quintile on an activity-weighted (i.e., employment) and unweighted basis. Looking at the activity-weighted graphs on the left, we see that the relationship between the occupational shares and TSB quintiles are as expected, with higher TSB establishments employing relatively fewer (lower-paid) production workers and relatively more (higher-paid) STEM and management employees (note the difference in scales for the STEM and management shares). Comparing these graphs to the unweighted graphs on the right provides insight into the difference between large and small establishments. The higher employment shares of STEM and management employees in the weighted graphs implies that within TSB quintiles, larger establishments employ relatively more of these workers.

In addition to the between-quintile variation in employment shares, we also expect to see between-industry variation. Blackwood et al. (2025) note that many occupational distributions are consistent with a given level of TSB. Thus, it is possible that there is considerable within-industry variation in occupation shares.

The graphs in Figure 5 show the within-industry share of total variation in employment shares for the three occupation groups.[28] Starting with the activity-weighted graphs, for all three occupation groups, the within-industry share of total variation is greatest in the first quintile—from around 50 percent for STEM occupations to nearly 70 percent for production workers. The graphs for all three occupation groups also exhibit a general U-shape, with troughs around the third and fourth quintiles. The within-industry

[28] Within each TSB quintile, we regress the occupation share on industry dummy variables. The figures show values of $(1 - R^2)$ from each regression, which are the within-industry shares of total variation in employment shares.

share of variation in employment shares is lowest for STEM occupations, which is not very surprising given that there is a lot of industry variation in technology intensity. For all three occupation groups, the within-industry share of total variation in employment shares is greater in the unweighted graphs, again reflecting the difference between large and small establishments. That the within-industry share is greater in the unweighted graphs suggests that there is more within-industry variation in employment shares among small establishments compared with large establishments.

Together, Figures 4 and 5 point to significant heterogeneity in how establishments organize production—even using these broad occupation groups. As expected, there is significant variation by TSB quintile and across industries, but there is also significant variation in employment shares within industries. Thus, establishments in the same industry do things very differently.

Figure 6. Percent of establishments with zero production workers by employment size class

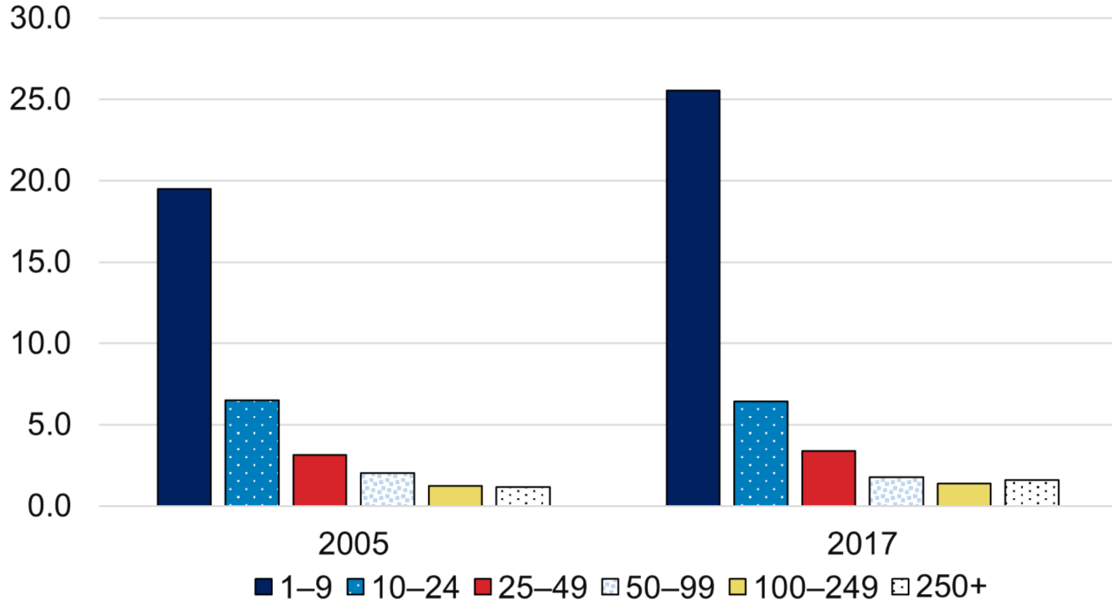

Notes: Estimates are unweighted. Source: OEWS-QCEW research dataset, authors' calculations.

Table 4. Percent of establishments with zero production workers for 13 four-digit industries with the highest percentage (SOC definitions)

Industry	25–99 workers (2005)	25–99 workers (2017)	100+ workers (2005)	100+ workers (2017)
Pharmaceuticals	14.1	6.5	1.1	2.2
Machinery	2.2	6.7	4.7	4.1
Computer	17.2	25.3	25.4	22.0
Telecommunications equipment	17.1	17.9	7.1	13.0
AV equipment	7.1	7.4	0.0	7.8
Semiconductor	8.7	4.5	1.3	6.1
Instruments	4.6	9.7	1.7	8.4
Magnetic & optical equipment	28.1	9.5	16.5	38.2
Aerospace	4.6	9.2	3.7	4.3
Dairy	5.3	2.5	0.0	0.0
Tobacco	11.4	12.1	0.0	0.0
Concrete/cement	11.9	13.4	4.5	3.4
Autos/trucks	12.5	0.0	3.9	0.0

Notes: Estimates are unweighted. Source: OEWS-QCEW research dataset, authors' calculations.

Given the apparent differences between large and small establishments with respect to the employment share of production workers, we investigate these shares along a different dimension—the percentage of establishments that have zero production workers. There are several reasons why an establishment might report employing zero production workers: they may have automated away their production workers, their production workers may do multiple tasks and be classified by their non-production tasks, or they may be design and marketing establishments that contract out the actual production of goods.

Figure 6 shows the percentage of establishments that have zero production workers by size. Not surprisingly, the largest percentage is among smaller establishments. The most likely explanation is that production workers may not be coded as such because they do other activities in these very small establishments. Still, even among larger establishments (25+ employees), there is a non-trivial percentage that does not employ any production workers.[29]

Finally, Table 4 shows the percentage of establishments with zero production workers for the 13 industries with the highest percentages. To avoid possible classification issues with small establishments, we focus on establishments with 25+ employees. Most of these industries are high-tech, although there are some non-tech industries as well (the bottom four rows). Among high-tech industries, some of the highest numbers show up in industries such as Computers, Telecommunications Equipment, and Magnetic & Optical Equipment. The numbers are quite large and, in some cases, changed substantially between 2005 and 2017.

[29] We plan to explore this avenue in the next iteration of this paper—for example, recreating Figures 4 and 5 excluding the very smallest establishments.

A More General Accounting for TFP Dispersion

We now return to our analysis of the relationship between occupations, tasks, skills, and productivity. The analysis above suggests that a simple (log) linear relationship between TFP and TSB within industries is inadequate. This section examines how much within-industry log TFP variation can be explained by these industry-specific nonlinear relationships. We do this using standard regression analysis, focusing on the adjusted R-squared as a metric. This analysis is descriptive and not causal, because TFP, skills/tasks, and occupations are all endogenous. Still, the analysis provides insights into how indicators of businesses "doing business differently" (through their skill/task and occupational share differences) account for dispersion in measured log TFP. In these regressions, we sweep out industry and year effects and present the results using both PW and AW weights.

Tables 5A and 5B present results for all industries. We first consider each skill, task, and occupation measure separately in a linear fashion. We then allow the coefficient on each measure to vary by industry, and then estimate specifications with groups of the skill, task and occupation measures (with and without industry interactions). Finally, we present results from a "full" specification where all measures are included, with both linear, quadratic, and cross terms, all of which are permitted to vary by industry.

Table 5A. **The relationship between the distribution of log TFP, occupations, tasks, and skills for all industries, propensity weighted**

Explanatory variable(s)	**Coefficient**	**Std. error**	**Adj. R^2**	**Adj. R^2 (indINT)**
TSU	0.133	0.022	0.0009	0.0116
TSB	0.155	0.018	0.0021	0.0110
Routine manual	-0.037	0.005	0.0016	0.0105
Routine cognitive	-0.037	0.006	0.0008	0.0075
Non-routine manual physical	-0.046	0.006	0.0015	0.0152
Interpersonal	0.035	0.007	0.0007	0.0072
Analytical	0.054	0.008	0.0012	0.0102
Production share	-0.058	0.011	0.0007	0.0183
STEM share	0.115	0.040	0.0003	0.0091
Management share	0.084	0.034	0.0002	0.0058
TSB + TSU + O*NET	-	-	0.0027	0.0512
TSB + TSU + O*NET + occ. shares	-	-	0.0032	0.0660
Polynomial(TSB + TSU + O*NET)	-	-	-	0.1109
Polynomial(TSB + TSU + O*NET + occ. shares)	-	-	-	0.1603

Notes: N=333,000 establishment-years. The dependent variable, log TFP, is demeaned by industry and year. All task/occupation measures also are demeaned by industry and year. Coefficients are not reported for multivariate regressions. The final column, titled "Adj. R^2 (indINT)," refers to regressions in which the explanatory variables are interacted with four-digit industry fixed effects. Source: Authors' calculations based on the OEWS survey, O*NET, and CMP.

Several interesting patterns emerge. First, without industry interactions, bivariate relationships are statistically significant with expected signs but virtually no explanatory power. For these specifications, results using AW yield more explanatory power. The bivariate linear results are consistent with our simple decomposition exercise above yielding relatively little explanatory power. Second, when we include industry-specific relationships, even while maintaining (log) linear relationships, explanatory power increases notably—especially in the AW weighted regressions. Third, the "full" specification (last row in the table) yields an adjusted R-squared of 0.16 for the PW regression and 0.22 in the AW case. We are also struck by how much the explanatory power increases with the addition of the broad occupational shares (production worker, STEM share, and management share) even after controlling for the TSB, TSU and O*NET measures (the second to last row). Although our final specification is far from parsimonious, the changes in the adjusted R-squared as we add more variables and interactions provide some guidance about what matters in explaining productivity dispersion.

We investigate these patterns further by re-estimating these regressions separately for establishments in high-tech and low-tech industries (Tables 6A, 6B, 7A and 7B).[30] There are several things to note in these tables. First, we can account for a larger fraction of TFP variation in high-tech industries compared with low-tech industries. Nearly every R-squared in the high-tech regressions (Tables 6A and 6B) is larger than the corresponding R-squared in the low-tech regressions (Tables 7A and 7B). Second, activity weighting matters more for high-tech industries than for low-tech industries. In low-tech industries, activity weighting matters only for the multivariate regressions (the last four rows), whereas it matters for all of the high-tech regressions. Third and most notably, our fully interacted regressions (the last rows) explain more than one-fifth of TFP variation in low-tech industries and more than one-third of TFP variation in high-tech industries.

[30] The high-tech group contains the following four-digit NAICS codes: 3241, 3251, 3252, 3254, 3332, 3333, 3336, 3339, 3341, 3342, 3343, 3344, 3345, 3346, 3353, and 3364. All other industries are classified as low-tech.

Table 5B. **The relationship between the distribution of log TFP, occupations, tasks, and skills for all industries, activity weighted**

Explanatory variable(s)	Coefficient	Std. error	Adj. R^2	Adj. R^2 (indINT)
TSU	0.575	0.023	0.0155	0.0975
TSB	0.446	0.019	0.0161	0.0938
Routine manual	-0.111	0.005	0.0102	0.0593
Routine cognitive	-0.069	0.006	0.0015	0.0118
Non-routine manual physical	-0.125	0.006	0.0074	0.0663
Interpersonal	0.089	0.006	0.0025	0.0484
Analytical	0.197	0.008	0.0153	0.0950
Production share	-0.306	0.012	0.0124	0.0819
STEM share	0.782	0.036	0.0215	0.0855
Management share	0.513	0.034	0.0041	0.0184
TSB + TSU + O*NET	-	-	0.0199	0.1391
TSB + TSU + O*NET + occ. shares	-	-	0.0247	0.1590
Polynomial(TSB + TSU + O*NET)	-	-	-	0.1797
Polynomial(TSB + TSU + O*NET + occ. shares)	-	-	-	0.2181

Notes: N=333,000 establishment-years. The dependent variable, log TFP, is demeaned by industry and year. All task/occupation measures also are demeaned by industry and year. Coefficients are not reported for multivariate regressions. The final column, titled “Adj. R^2 (indINT),” refers to regressions in which the explanatory variables are interacted with four-digit industry fixed effects. Source: Authors’ calculations based on the OEWS survey, O*NET, and CMP.

Table 6A. **The relationship between the distribution of log TFP, occupations, tasks, and skills in high-tech industries, propensity weighted**

Explanatory variable(s)	Coefficient	Std. error	Adj. R^2	Adj. R^2 (indINT)
TSU	0.330	0.063	0.0049	0.0235
TSB	0.303	0.043	0.0086	0.0236
Routine manual	-0.078	0.013	0.0060	0.0167
Routine cognitive	-0.087	0.021	0.0033	0.0074
Non-routine manual physical	-0.061	0.015	0.0026	0.0159
Interpersonal	0.046	0.022	0.0010	0.0074
Analytical	0.110	0.021	0.0047	0.0202
Production share	-0.162	0.028	0.0054	0.0273
STEM share	0.158	0.064	0.0012	0.0216
Management share	0.038	0.082	0.0000	0.0117
TSB + TSU + O*NET	-	-	0.0106	0.0596
TSB + TSU + O*NET + occ. shares	-	-	0.0160	0.0864
Polynomial(TSB + TSU + O*NET)	-	-	-	0.1347
Polynomial(TSB + TSU + O*NET + occ. shares)	-	-	-	0.2073

Notes: N=49,500 establishment-years. The dependent variable, log TFP, is demeaned by industry and year. All task/occupation measures also are demeaned by industry and year. Coefficients are not reported for multivariate regressions. The final column, titled "Adj. R^2 (indINT)," refers to regressions in which the explanatory variables are interacted with four-digit industry fixed effects. The high-tech industries include the following four-digit NAICS codes: 3241, 3251, 3252, 3254, 3332, 3333, 3336, 3339, 3341, 3342, 3343, 3344, 3345, 3346, 3353, and 3364. All other industries are classified as low-tech.
Source: Authors' calculations based on the OEWS survey, O*NET, and CMP.

Table 6B. **The relationship between the distribution of log TFP, occupations, tasks, and skills in high-tech industries, activity weighted**

Explanatory variable(s)	**Coefficient**	**Std. error**	**Adj. R^2**	**Adj. R^2 (indINT)**
TSU	1.088	0.076	0.0472	0.2082
TSB	0.715	0.038	0.0473	0.1995
Routine manual	-0.203	0.012	0.0276	0.1111
Routine cognitive	-0.216	0.024	0.0066	0.0137
Non-routine manual physical	-0.229	0.016	0.0220	0.1244
Interpersonal	0.301	0.022	0.0134	0.1070
Analytical	0.355	0.019	0.0450	0.2074
Production share	-0.663	0.032	0.0456	0.1640
STEM share	0.910	0.049	0.0502	0.1814
Management share	0.807	0.084	0.0093	0.0343
TSB + TSU + O*NET	-	-	0.0570	0.2540
TSB + TSU + O*NET + occ. shares	-	-	0.0683	0.2798
Polynomial(TSB + TSU + O*NET)	-	-	-	0.3046
Polynomial(TSB + TSU + O*NET + occ. shares)	-	-	-	0.3531

Notes: N=49,500 establishment-years. The dependent variable, log TFP, is demeaned by industry and year. All task/occupation measures also are demeaned by industry and year. Coefficients are not reported for multivariate regressions. The final column, titled "Adj. R^2 (indINT)," refers to regressions in which the explanatory variables are interacted with four-digit industry fixed effects. The high-tech industries include the following four-digit NAICS codes: 3241, 3251, 3252, 3254, 3332, 3333, 3336, 3339, 3341, 3342, 3343, 3344, 3345, 3346, 3353, and 3364. All other industries are classified as low-tech.
Source: Authors' calculations based on the OEWS survey, O*NET, and CMP.

Table 7A. **The relationship between the distribution of log TFP, occupations, tasks, and skills in low-tech industries, propensity weighted**

Explanatory variable(s)	Coefficient	Std. error	Adj. R^2	Adj. R^2 (indINT)
TSU	0.080	0.021	0.0003	0.0078
TSB	0.103	0.019	0.0009	0.0071
Routine manual	-0.026	0.005	0.0008	0.0085
Routine cognitive	-0.026	0.006	0.0004	0.0075
Non-routine manual physical	-0.042	0.007	0.0013	0.0149
Interpersonal	0.032	0.007	0.0006	0.0072
Analytical	0.038	0.008	0.0006	0.0070
Production share	-0.028	0.011	0.0002	0.0155
STEM share	0.061	0.042	0.0000	0.0052
Management share	0.097	0.037	0.0003	0.0039
TSB + TSU + O*NET	-	-	0.0016	0.0485
TSB + TSU + O*NET + occ. shares	-	-	0.0019	0.0595
Polynomial(TSB + TSU + O*NET)	-	-	-	0.1032
Polynomial(TSB + TSU + O*NET + occ. shares)	-	-	-	0.1450

Notes: N=283,500 establishment-years. The dependent variable, log TFP, is demeaned by industry and year. All task/occupation measures also are demeaned by industry and year. Coefficients are not reported for multivariate regressions. The final column, titled “Adj. R^2 (indINT),” refers to regressions in which the explanatory variables are interacted with four-digit industry fixed effects. The high-tech industries include the following four-digit NAICS codes: 3241, 3251, 3252, 3254, 3332, 3333, 3336, 3339, 3341, 3342, 3343, 3344, 3345, 3346, 3353, and 3364. All other industries are classified as low-tech.
Source: Authors’ calculations based on the OEWS survey, O*NET, and CMP.

Table 7B. **The relationship between the distribution of log TFP, occupations, tasks, and skills in low-tech industries, activity weighted**

Explanatory variable(s)	Coefficient	Std. error	Adj. R^2	Adj. R^2 (indINT)
TSU	0.233	0.019	0.0013	0.0173
TSB	0.617	0.016	0.0135	0.0390
Routine manual	0.056	0.005	0.0014	0.0288
Routine cognitive	-0.049	0.006	0.0005	0.0365
Non-routine manual physical	0.110	0.007	0.0029	0.0315
Interpersonal	0.121	0.006	0.0031	0.0197
Analytical	0.101	0.007	0.0021	0.0163
Production share	-0.072	0.011	0.0004	0.0323
STEM share	1.220	0.034	0.0132	0.0399
Management share	0.482	0.033	0.0018	0.0141
TSB + TSU + O*NET	-	-	0.0431	0.1345
TSB + TSU + O*NET + occ. shares	-	-	0.0548	0.1599
Polynomial(TSB + TSU + O*NET)	-	-	-	0.1780
Polynomial(TSB + TSU + O*NET + occ. shares)	-	-	-	0.2283

Notes: N=283,500 establishment-years. The dependent variable, log TFP, is demeaned by industry and year. All task/occupation measures also are demeaned by industry and year. Coefficients are not reported for multivariate regressions. The final column, titled "Adj. R^2 (indINT)," refers to regressions in which the explanatory variables are interacted with four-digit industry fixed effects. The high-tech industries include the following four-digit NAICS codes: 3241, 3251, 3252, 3254, 3332, 3333, 3336, 3339, 3341, 3342, 3343, 3344, 3345, 3346, 3353, and 3364. All other industries are classified as low-tech.
Source: Authors' calculations based on the OEWS survey, O*NET, and CMP.

Putting all the pieces together, we conclude that there is enormous between-establishment, within-industry variation in measures of skill, task and occupations that are correlated with the between-establishment, within-industry variation in TFP. The fraction of the variance accounted for by skill, task, and occupation variation is higher for larger establishments (based on the activity-weighted results). This may reflect genuine differences in the role of these differences for large establishments, but it may also reflect that large establishments have more defined differences in these dimensions. An important feature of the findings is that the relationship between these skill, task, and occupation measures and TFP varies substantially by industry and is nonlinear. Moreover, these relationships are more pronounced in high-tech industries.

V. Concluding remarks

Measured productivity differences among establishments are large even within narrowly defined industries. Apart from true differences in efficiency, dispersion in measured productivity can be due to unobserved differences in organizational characteristics, input responsiveness, markups, measurement error, and production function specification. In addition, unobserved differences in inputs—for example, capital and labor characteristics and/or composition—are also subsumed in the productivity residual.

In this paper, we examine the relationship between establishment-level TFP within industries and measures of the skill, task, and occupational differences across these same establishments. We construct establishment-level measures of occupation composition and task/skill intensity using data from the OEWS survey and the O*NET and then link these to the CMP establishment-level productivity data. We match the OEWS survey data and CMP data using a hierarchical algorithm that prioritizes information on EINs, narrowly-defined industry, and state—a significant challenge given the lack of a common establishment identifier in the two data sources.

Our empirical results indicate that there are highly nonlinear, industry-specific relationships between measured TFP and the skill, task, and occupational measures within industries. It is evident that businesses are organized differently via the information on the skill, task, and occupational differences. These organizational differences are related in a complex manner with measured TFP variation.

The standard approach to measure and study within-industry variation in productivity between establishments is to assume that all establishments use the same production technology. The benchmark is often a Cobb-Douglas specification with industry-level factor elasticities. Our results indicate that this approach is inadequate. Understanding measured productivity variation across establishments within industries requires understanding how businesses organize themselves differently in terms of their mix of skills, tasks, and occupations. There are many challenges going forward.

First, it will be important to understand the driving forces underlying these differences. A potentially fruitful next step will be to integrate measures of technology adoption from the Annual Business Survey (e.g., automation) into the linked CMP-OEWS data. This integration will allow us to explore how these observable differences are related to the skill/task/occupation differences we identify. This would permit us to begin exploring the hypotheses from the work of Acemoglu and Restrepo (2018a, 2018b, 2019a, 2019b, 2020) that emphasize the complex relationship between automation, tasks, and factors of production such as capital and labor.

Second, there is a large literature studying the creative destruction process using establishment-level differences in measured productivity. Consistent with canonical models of firm dynamics, measured productivity differences across establishments are closely connected to growth and survival dynamics. The latter patterns indicate that measured productivity differences have a systematic relationship with key outcomes. However, the differences in organizational structure in terms of skills, tasks, and occupations raise questions about how to think about these well-documented connections between productivity and reallocation dynamics. Third, and relatedly, our results suggest progress needs to be made to specify parsimonious but much richer production specifications for studying establishment-level variation in measured productivity.

This is an ambitious to-do list, and this paper has only taken initial steps. The insights from these first steps suggest that the integrated CMP-OEWS data infrastructure has great potential for making progress on these issues. The research and development activities that led to the integrated CMP-OEWS were possible through the continued collaboration between the two statistical agencies. We hope this paper also serves to highlight the benefits of this collaboration.

References

Acemoglu, Daron, and David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." In *Handbook of Labor Economics*, vol. 4, Part B, edited by Orley Ashenfelter and David Card, 1043–171. Elsevier.

Acemoglu, Daron, and Pascual Restrepo. 2018a. "Modeling Automation." *AEA Papers and Proceedings* 108: 48–53. https://doi.org/10.1257/pandp.20181020.

Acemoglu, Daron, and Pascual Restrepo. 2018b. "The Race Between Machine and Man: Implications of Technology for Growth, Factor Shares, and Employment." *American Economic Review* 108 (6): 1488–542. https://doi.org/10.1257/aer.20160696.

Acemoglu, Daron, and Pascual Restrepo. 2019a. "Artificial Intelligence, Automation and Work." In *The Economics of Artificial Intelligence: An Agenda*, edited by Ajay Agrawal, Joshua Gans, and Avi Goldfarb, 197–236.

Acemoglu, Daron, and Pascual Restrepo. 2019b. "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives* 33 (2): 3–30. https://doi.org/10.1257/jep.33.2.3.

Acemoglu, Daron, and Pascual Restrepo. 2020. "Robots and Jobs: Evidence from US Labor Markets." *Journal of Political Economy* 128 (6): 2188–244. https://doi.org/10.1086/705716.

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *The Quarterly Journal of Economics* 118 (4): 1279–333. https://doi.org/10.1162/003355303322552801.

Bernard, Andrew B., and Teresa C. Fort. 2013. "Factoryless Goods Producing Firm." *American Economic Review: Papers & Proceedings* 105(5): 518–523. https://doi.org/10.1257/aer.p20151044.

Blackwood, G. Jacob, Lucia S. Foster, Cheryl A. Grim, John Haltiwanger, and Zoltan Wolf. 2021. "Macro and Micro Dynamics of Productivity: From Devilish Details to Insights." *American Economic Journal: Macroeconomics* 13 (3): 142–72. https://doi.org/10.1257/mac.20170282.

Blackwood, G. Jacob, Cindy Cunningham, Matthew Dey, Lucia Foster, Cheryl Grim, John Haltiwanger, Rachel Nesbit, Sabrina Wulff Pabilonia, Jay Stewart, Cody Tuttle, and Zoltan Wolf. 2025. "Opening the Black Box: Task and Skill Mix and Productivity Dispersion." In *Technology, Productivity, and Economic Growth*, edited by Susanto Basu, Lucy Eldridge, John Haltiwanger, and Erich Strassner, 323–360. University of Chicago Press. https://www.nber.org/books-and-chapters/technology-productivity-and-economic-growth.

Chow, Melissa, Teresa C. Fort, Christopher Goetz, Nathan Goldschlag, James Lawrence, Elisabeth Ruth Perlman, Martha Stinson, and T. Kirk White. 2021. "Redesigning the Longitudinal Business Database." Center for Economic Studies

Discussion Paper No. 21-08. https://www.census.gov/library/working-papers/2021/adrm/CES-WP-21-08.html.

Bureau of Labor Statistics. 2001–2020. Occupational Employment and Wage Statistics. [Restricted-access dataset.] https://www.bls.gov/oes/.

Bureau of Labor Statistics and Census Bureau. 1972–2020. Collaborative Micro-productivity Project. [Restricted-access dataset version 7.] https://www.census.gov/fsrdc.

Cunningham, Cindy, Lucia Foster, Cheryl Grim, John Haltiwanger, Sabrina Wulff Pabilonia, Jay Stewart, and Zoltan Wolf. 2023. "Dispersion in Dispersion: Measuring Establishment Level Differences in Productivity." *Review of Income and Wealth* 69 (4): 999–1032. https://doi.org/10.1111/roiw.12616.

Dey, Matthew, David S. Piccone, Jr., and Stephen M. Miller. 2019. "Model-Based Estimates for the Occupational Employment Statistics Program." *Monthly Labor Review* August. https://www.bls.gov/opub/mlr/2019/article/model-based-estimates-for-the-occupational-employment-statistics-program.htm.

Economic Policy Institute. 2025. Current Population Survey Extracts, Version 2025.7.10. https://microdata.epi.org.

Iranzo, Susana, Fabiano Schivardi, and Elisa Tosetti. 2008. "Skill Dispersion and Firm Productivity: An Analysis with Employer-Employee Matched Data." *Journal of Labor Economics* 26 (2): 247–85. https://doi.org/10.1086/587091.

Stoyanov, Andrey, and Zubanov, Nick. 2022. "Skill Complementarity in Production Technology: New Empirical Evidence and Implications." *German Economic Review* 23 (2): 233–74. https://doi.org/10.1515/ger-2020-0102.

Syverson, Chad. 2004. "Market Structure and Productivity: A Concrete Example." *Journal of Political Economy*, 112(6): 1181–1222. https://doi.org/10.1086/424743.

Syverson, Chad. 2011. "What Determines Productivity?" *Journal of Economic Literature* 49 (2): 326–65. https://www.doi.org/10.1257/jel.49.2.326.

Zoghi, Cindy. 2007. "Measuring Labor Composition: A Comparison of Alternate Methodologies." In *Labor in the New Economy*, edited by Katherine G. Abraham, James R. Spletzer, and Michael Harper, University of Chicago Press. https://www.doi.org/10.7208/chicago/9780226001463.001.0001.

United States Department of Labor, Employment and Training Administration. 1991. *The Revised Handbook of Analyzing Jobs*.

United States Department of Labor, Employment & Training Administration. 2007–2017. O*NET 12.0, 13.0, 19.0, and 22.1. [Dataset.] https://www.onetcenter.org/db_releases.html.

Data Appendix

Details on CM-ASM Industry Codes

Our goal is to find, for each establishment-year observation in the CMP, the best match in the OEWS. An establishment's industry code in the CMP is an integral part of the matching procedure, and it is therefore important to have the best possible chance of matching an establishment's NAICS code to the industry code in the OEWS. We use four different NAICS codes for our CMP dataset: the LBD code, the ASM code (only available for manufacturing observations), the VC code, and a "combined" code created by combining information from the CM, ASM, and LBD. We create our combined code as follows: (1) for establishments that are surveyed by the CM, we use the industry code from the CM year that is closest to the reference year; (2) if no CM code is available, we use the ASM code; and (3) if no ASM code is available, we use the LBD code. We think this combined code most closely aligns with the timing of industry code updates in the OEWS survey.[31] Finally, we make a "time-consistent code" correction to the LBD, ASM, and combined codes. The correction aggregates six-digit codes to five-digit codes in cases where there is consolidation or other changes in the classification between different NAICS vintages. Note that this correction differs from the VC code approach taken by Chow et al. (2021). Our time-consistent codes do not aim to put everything in terms of the 2017 classification vintage, but instead to simply aggregate any codes that disappear or are broken up between vintages so that we can abstract from vintage differences. We describe below how we use these seven versions (ASM,

[31] The OEWS program occasionally updates industry codes based on the information collected from each establishment's answers to the survey. When this occurs, the OEWS industry code will differ from that in the Quarterly Census of Employment and Wages (QCEW), which is the BLS business register.

time-consistent ASM, LBD, time-consistent LBD, combined, time-consistent combined, and VC) of NAICS codes in our matching procedure.

Matching Procedure Example

Consider a CMP establishment within a given EIN and suppose that there are 12 potential donors in the OEWS that have the same EIN. In step 1, we start by using the combined NAICS code. If none of the 12 candidates is in the same six-digit industry code as defined by our combined NAICS measure, we then check for agreement using the six-digit industry ASM NAICS code. If there is still no match, we use the LBD NAICS code. If there is still no agreement, we move to step 2, allowing for possible mismatches in industry vintage changes by instead using the time-consistent combined, ASM, and LBD NAICS in a similar iterative manner. The algorithm continues through step 5 or until a match is found. If there is no match in step 5, the observation is not used.

In this example, suppose we identify three candidate OEWS donors that match in Step 1. We need to narrow them down to one final donor. Among these three candidates, we first look for the one most similar in size to the CMP establishment. Suppose that eliminates one potential donor, but the two other candidates have the same employment. We next compare the years in which those donors were surveyed. If the CMP year was 2013, we would be evaluating potential matches from the 2012–2014 OEWS surveys. Suppose both donors were surveyed in 2013, so that this tiebreaker does not help us narrow down our candidates. The final step would be to randomly choose one of the two remaining candidates to be our preferred donor. By following our step-by-step matching and tiebreaking processes, we have identified one match out of the original 12 candidates. In the end, 33% of CMP observations have been augmented with occupation data from the OEWS.

*The O*NET Data*

The O*NET data are collected from workers in targeted occupations at establishments and contain over 275 variables that describe each occupation. The O*NET database is sponsored by the Employment and Training Administration of the Department of Labor and is collected through the National Center for O*NET Development and the Research Triangle Institute. O*NET first began surveying job holders in 2001. Prior to that, past Dictionary of Occupational Titles data, collected sometimes decades earlier by job analysts visiting workplaces, were recoded into O*NET variables. Because new surveying was rolled in gradually, the first O*NET completely based on surveys was released in 2008. O*NET resurveys occupations on a rolling basis over a five-year period. The number of respondents per occupation varies, and respondents are randomly selected to answer a subset of the questionnaire. The value of a particular O*NET variable is the average response over the jobholders who answered that question, so within-occupation variation cannot be observed. See Handel (2016) for more about the history, strengths, and weaknesses of O*NET.

Acemoglu and Autor (2011) use 16 O*NET variables corresponding to the five task categories described in the text: non-routine cognitive (analytical), non-routine (interpersonal), non-routine manual physical, routine cognitive, and routine manual. Non-routine cognitive (analytical) includes analyzing data/information, thinking creatively, and interpreting information for others. Non-routine cognitive (interpersonal) includes establishing and maintaining personal relationships; guiding, directing, and motivating subordinates; and coaching/developing others. Non-routine manual physical includes operating vehicles, mechanized devices, or equipment; tasks where workers use their hands to handle, control, or feel objects, tools, or controls; manual dexterity; and spatial orientation. Routine cognitive includes importance of repeating the same tasks, importance of being exact or accurate, and structured vs. unstructured work (reverse). Routine manual includes tasks where the pace of work is determined by speed of equipment, controlling machines and processes, and tasks requiring repetitive motions (Acemoglu and Autor 2011, p. 1163). The O*NET-SOC occupational categories are aggregated to SOC categories, and each variable is scaled and then standardized to mean zero and standard deviation one using employment weights from the OEWS survey. The five indexes are created by summing the standardized variables for each task category, which are then once again normalized.

The OEWS Research Dataset

Because the sample design of the OEWS is geared toward measuring employment, creating establishment weights is complicated. As an alternative to reweighting, we use a research dataset that is a modified version of the dataset developed by Dey, Piccone, and Miller (2019). This research dataset supplements the OEWS data by imputing occupation data for the entire Quarterly Census of Employment and Wages (QCEW), which is the BLS business register and is the sample frame for BLS establishment surveys. The main advantage of this approach is that all establishments are represented and have a weight of one.

The imputation process involves two stages, a matching stage where potential donors are identified and a selection stage where the best donor is selected. The process is hierarchical, where the conditions for finding acceptable matches are sequentially relaxed. At the most detailed level of the hierarchy, a donor and frame unit will match on industry (six-digit NAICS), ownership (private or type of government), state, and county and will have similar employment levels. As the process continues through the hierarchy, geography is relaxed first and then ownership. It is not until late in the process, after most of the frame units have already found an acceptable donor, that industry and employment proximity are relaxed. The matching stage often results in multiple potential donors. To preserve variance, the selection of a particular donor from the set of acceptable matches is random. Wages are adjusted to account for differences by MSA and industry. In contrast to the published statistics, the research dataset centers the sample on the reference year instead of using data from the five panels prior to May of the reference year. For example, under this approach, the sample for May 2017 is constructed using data from the following panels: May 2018, November 2017, May 2017, November 2016, May 2016, and November 2015. This results in a nationally representative sample centered on May 2017. To avoid overlap, these “year samples” are constructed at three-year intervals. This effectively assumes the occupational mix within an establishment is fixed over the three-year interval.

Data Appendix References

Acemoglu, Daron and David Autor. 2011. “Skills, Tasks and Technologies: Implications for Employment and Earnings.” In *Handbook of Labor Economics*, vol. 4, Part B, edited by Orley Ashenfelter and David Card, 1043–171. Elsevier. https://doi.org/10.1016/S0169-7218(11)02410-5.

Chow, Melissa, Teresa C. Fort, Christopher Goetz, Nathan Goldschlag, James Lawrence, Elisabeth Ruth Perlman, Martha Stinson, and T. Kirk White. 2021. “Redesigning the Longitudinal Business Database.” Center for Economic Studies Discussion Paper No. 21-08. https://www.census.gov/library/working-papers/2021/adrm/CES-WP-21-08.html.

Dey, Matthew, David S. Piccone, Jr., and Stephen M. Miller. 2019. “Model-Based Estimates for the Occupational Employment Statistics Program.” *Monthly Labor Review* August. https://www.bls.gov/opub/mlr/2019/article/model-based-estimates-for-the-occupational-employment-statistics-program.htm.

Handel, Michael. 2016. “The O*NET Content Model: Strengths and Limitations.” *Journal of Labor Market Research* 49:157–76. https://doi.org/10.1007/s12651-016-0199-8.