Opening the Black Box: Task and Skill Mix and Productivity Dispersion

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 15, 2022

An important gap in most empirical studies of establishment-level productivity is the limited information about workers' characteristics and their tasks. Skill-adjusted labor input measures have been shown to be important for aggregate productivity measurement. Moreover, the theoretical literature on differences in production technologies across businesses increasingly emphasizes the task content of production. Our ultimate objective is to open this black box of tasks and skills at the establishment-level by combining establishment-level data on occupations from the Bureau of Labor Statistics (BLS) with a restricted-access establishment-level productivity dataset created by the BLS-Census Bureau Collaborative Micro-productivity Project. We take a first step toward this objective by exploring the conceptual, specification, and measurement issues to be confronted. We provide suggestive empirical analysis of the relationship between within-industry dispersion in productivity and tasks and skills. We find that within-industry productivity dispersion is strongly positively related to within-industry task/skill dispersion.

*Blackwood: Amherst College; Cunningham, Dey, Pabilonia, and Stewart: U.S. Bureau of Labor Statistics; Foster and Grim: Center for Economic Studies, U.S. Census Bureau; Haltiwanger: University of Maryland; Nesbit: University of Maryland and Center for Economic Studies, U.S. Census Bureau; Tuttle: University of Texas at Austin; Wolf: New Light Technologies. Blackwood, Haltiwanger, and Tuttle were also Schedule A 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. This paper uses public domain DiSP data (the public domain data have been cleared by the Census Bureau Disclosure Review Board, Clearance Number: CBDRB-FY21-305). The authors would like to thank Emin Dinlersoz, Lucy Eldridge, John Eltinge, Susan Houseman, Erich Strassner, and participants at the NBER/CRIW Conference on Technology, Productivity, and Economic Growth for comments. Corresponding author: Stewart.Jay@bls.gov.

1. Introduction

It is well known that productivity varies across establishments, even within detailed industries. For example, Cunningham et al. (2022) found that on average an establishment at the 90th percentile of the total factor productivity (TFP) distribution is about 2.9 times as productive as an establishment at the 10th percentile within four-digit NAICS manufacturing industries. In a survey article, Syverson (2011) discusses several possible sources of productivity dispersion, including difficult-to-measure factors such as differences in managerial talent and differences in the quality of labor and other inputs. In this paper, we explore the role of differences in the characteristics of workers and the tasks they perform across establishments within the same industry.

Our paper builds on a joint Bureau of Labor Statistics (BLS) and Census Bureau project that developed publicly available productivity dispersion statistics (Dispersion Statistics on Productivity, DiSP) posted on the BLS and Census Bureau websites and a restricted-access dataset available to researchers.[1] Consistent with the prior literature, the DiSP data show that the degree of within-industry dispersion varies considerably across industries and over time. A limitation of the DiSP data (as well as much of the related literature) is that the labor input is measured as total hours worked by all workers. Ideally, measurement of establishment-level productivity would also account for the types of workers the establishment employs and the tasks they perform. Our ultimate goal is to address this limitation by integrating establishment-level

[1] See Cunningham et al. (2022) for a detailed description of the development of the datasets. DiSP is available at: https://www.bls.gov/productivity/articles-and-research/dispersion-statistics-on-productivity/ and https://www.census.gov/disp. The restricted-access dataset is available for use by qualified researchers on approved projects in the Federal Statistical Research Data Centers (https://www.census.gov/fsrdc).

data from the BLS Occupation Establishment and Wage Survey (OEWS) with Census Bureau establishment-level business data.

In this paper, we explore the conceptual, measurement, and specification issues to be addressed for this integration to be successful. In addition, we include an empirical analysis relating within-industry dispersion of productivity measures to within-industry dispersion of task and skill measures for four-digit NAICS manufacturing industries over the 2000-to-2017 period.

Our productivity dispersion statistics come from the DiSP data (described in detail in Cunningham, et al. (2022)). We create four types of measures of tasks and/or skills using OEWS and Occupational Information Network (O*NET) data. Our first task/skill measure is a composite index accounting for the pricing of occupations in the labor market. It is related to, but distinct from, the skill-adjusted labor input measure BLS publishes as part of its official TFP measures.[2] Conceptually, it is a counterfactual average establishment wage: the average wage of an establishment if the establishment paid the national mean occupational wage for each occupation it employs. Importantly, this measure reflects the share of employment in each occupation at the establishment level each year. We refer to this as a bundled task/skill intensity index (TSB), because the pricing of tasks is bundled through the occupations.

Our second task/skill intensity index is similarly a counterfactual wage but based on predicted wages from a linear regression of OEWS national occupational wages on five task indexes constructed from work activities and work-context-importance scales in the O*NET (as described in Acemoglu and Autor (2011), p. 1163). The five task indexes are: non-routine cognitive (analytical), non-routine cognitive (interpersonal), routine cognitive, routine manual,

[2] See https://www.bls.gov/productivity/technical-notes/changes-in-composition-of-labor-total-factor-productivity-2014.pdf for a description of the official measure. For a more detailed discussion of the theory and measurement issues behind the labor composition index, see Zoghi (2007).

and non-routine manual physical. This counterfactual wage is the average wage of an establishment if the establishment paid the predicted wage for each task it employs (again using the employment shares of each task at the establishment level each year). We refer to this index as an unbundled task/skill index (TSU) because it prices the tasks directly regardless of which occupations accomplish these tasks.

The TSB and TSU reflect task differences across establishments as well as the prices of those tasks in the labor market, where prices reflect the skills required to accomplish those tasks (among other things that may determine wages). The major difference between these two measures is the bundled measure (TSB) implicitly accounts for how the tasks are organized into occupations, while the unbundled measure (TSU) does not.

We also create measures that do not use wage information but rather use direct information on the tasks being performed at individual establishments based on the occupational mix. For these measures, we use the establishment values of the five O*NET task indexes described above individually. Finally, we create a measure based on the percentage of STEM workers in each establishment (%STEM), based on the occupational mix.

We compare the within-industry labor productivity (LP) and TFP dispersion measures from DiSP to the within-industry dispersion in these task/skill measures. To preview our results, we find the TSB and TSU are highly correlated with each other at the establishment level within industries. The TSB is also positively correlated with the analytical task index, the interpersonal task index, and the %STEM, but negatively correlated with the non-routine manual physical, routine manual, and routine cognitive tasks indexes. These establishment-level correlations conform with our intuition about the skills required to perform these composite tasks.

Turning to the relationship between within-industry dispersion in task/skill indexes and dispersion in productivity, we find higher within-industry productivity dispersion is associated with higher within-industry dispersion of TSB, TSU, the analytical task index, and %STEM. These patterns differ quantitatively across different groupings of manufacturing industries but they are especially strong in the high-tech industries. For example, the elasticity of within-industry TFP dispersion with respect to TSB dispersion is about three times larger in the high-tech manufacturing industries than in the non-high-tech ("non-tech" hereafter) manufacturing industries.

The remarkably high within-industry dispersion of both productivity and task/skill intensities across establishments in high-tech industries implies there is considerable heterogeneity in both the outcomes and the ways of doing business, especially among the most innovative sectors of the economy. It is well known high-tech industries exhibit high productivity growth and are more STEM occupation intensive (see, e.g., Decker et al. (2020)). Our findings suggest the strong relationship in first moments carries over to the corresponding second moments.

While our results are suggestive, they are promising for the longer-run objective of integrating the OEWS data with the Census Bureau business microdata at the establishment-level. The establishment-level analysis will enable us to explore related issues such as the relationship between technology adoption and the task/skill mix of businesses. We will examine this in the next phase of our research, but we provide an overview of the potential for such analysis with establishment-level data integration.

While the integration of the OEWS and CMP microdata awaits future research, in an initial step to this integration we find that more than 80 percent of between-establishment variation in our skill/task indexes among single- and multi-unit EINS is accounted for by between-taxpayer ID (EIN) variation. In other words, establishments within the same firm and industry exhibit considerable similarity in their skill and task mix indexes. Relatedly, the productivity dispersion literature has found that high-productivity establishments are part of high-productivity firms, for example, see Baily, Hulten, and Campbell (1992). These findings are interesting in their own right, but they also facilitate the planned integration of the OEWS and ASM data at the micro level, because a common identifier is the EIN.

The paper proceeds as follows. Section 2 presents a conceptual framework largely through a review of the literature relating productivity to the skills of workers and tasks they perform. Section 3 describes the productivity data, the OEWS data, and our task/skill intensity measures. Section 4 presents the results relating within-industry productivity and task/skill dispersion. Section 5 provides an overview of next steps. Our concluding remarks are in Section 6.

2. Background and Conceptual Framework

In a standard production function, $Q_{et}=A_{et}\cdot F(L_{et},K_{et})$—where Q_{et} is output, L_{et} is labor input, K_{et} is capital input, A_{et} is a Hicks-neutral productivity term (often interpreted as technical change), and *e* and *t* index establishments and time respectively—the different ways businesses use labor may show up as differences in L, $\frac{\partial F}{\partial L}$, or both. Our objective is to take a first step toward exploring heterogeneity in labor input and its potential effect on productivity dispersion. An establishment may have higher measured productivity than its competitors because it uses a given set of labor inputs more efficiently, or because its production process consists of more advanced tasks (generally) accompanied by more skilled labor.

Within the simple production function specification above, differences in skills and tasks across establishments can be captured by introducing a multiplier, Z_{et}, that can be interpreted as an adjustment that converts labor hours into efficiency units based on skills and tasks. The first argument of $F(.)$ then becomes $Z_{et} \cdot L_{et}$, which accounts for differences across establishments due to differences in the efficiency with which L_{et} is used in production.[3] Both A_{et} and Z_{et} are efficiency parameters. The difference is that while A_{et} increases the productivity of both factors of production, Z_{et} affects only the productivity of labor.

Another approach is to explicitly model labor types in order to analyze the returns to different skills. For ease of exposition, we drop *e* and *t* subscripts for the remainder of this section. The canonical model for understanding skill premia assumes two types of workers, low-skilled and high-skilled (see equation (2) in Acemoglu and Autor (2011)):

$$F = [(A_l L)^\rho + (A_h H)^\rho]^{1/\rho} \tag{1}$$

where L and H denote low-skilled and high-skilled workers and ρ is a constant that characterizes the substitutability between L and H. A_l and A_h are factor-specific augmenting technology terms. The elasticity of substitution, $\sigma = 1/(1-\rho)$, is defined as the percentage change in relative demand for low-skilled workers for a percentage change in the relative price of high-skilled workers. If the two types are substitutes, an increase in the relative price of one leads to an increase in demand for the other skill. In the limiting case of perfect substitution, or $\sigma \to \infty$ ($\rho \to 1$), relative wages are constant. In the other extreme, $\sigma \to 0$ ($\rho \to -\infty$), the two labor types are perfect complements, i.e., they can be used only in fixed proportions. The third special case, $\sigma \to 1$ ($\rho \to 0$), yields the Cobb-Douglas production function where fixed shares are paid to each

[3] Gollop et al. (1987) first demonstrates the potential importance of using efficiency units of labor. The methods developed from this early work have been widely adopted by statistical agencies around the world (see Schreyer 2001). BLS uses a related approach in their total factor productivity measures (see https://www.bls.gov/opub/hom/msp/home.htm).

factor. In the constant elasticity of substitution (CES) framework above, σ is a crucial parameter because it determines not only how changes in the supply of labor types affect wages and demand, but also how changes in technology affect L and H. In particular, factor-augmenting technology is encapsulated in A_l and A_h in the sense that technical change may affect the productivity of labor types independently of each other. But new technologies do not replace either type of labor.

A more general version of the model above allows for the possibility of technological advances complementing or fully replacing either labor type (Acemoglu (1998)):

$$F = [\,(1-\alpha)(A_l L + B_l)^{\,\rho} + \alpha(A_h H + B_h)^{\,\rho}]^{\,1/\rho} \tag{2}$$

where B_l and B_h denote skill-replacing technologies that are perfect substitutes for the two labor types,[4] and α reflects the distribution of tasks between different types of labor (Card and DiNardo (2002)). An increase in α increases the productivity of H and at the same time lowers the productivity of L, while A_l and A_h only affect the productivity of the corresponding factor.[5]

Although the models described above have been empirically successful in studying the supply and demand for skills, they have been less successful in explaining phenomena such as the different wage dynamics of high- and low-skill workers, job polarization, diffusion of low-skill-replacing technologies, and automation.[6] Acemoglu and Autor (2011) and others have examined alternative models that define the production process as a set of tasks done by workers with the appropriate skills (e.g., Autor, Levy, and Murnane 2003; Acemoglu and Restrepo 2018a, 2019b).

---

[4] In the case of skill-complementing technologies, the linear terms of the inner parentheses in equation (2) would have the same harmonic sum structure as in equation (1) but with exponents less than zero, in which case, $F = \left[(1-\alpha)\left([(A_l L)^{\rho_l} + {B_l}^{\rho_l}]^{1/\rho_l}\right)^{\rho} + \alpha\left([(A_h H)^{\rho_h} + {B_h}^{\rho_h}]^{1/\rho_h}\right)^{\rho}\right]^{1/\rho}$ with $\rho_l, \rho_h < 0$.

[5] The endogenous technology choice model in Dinlersoz and Wolf (2018) uses a similar but simpler structure.

[6] See Acemoglu and Autor (2011), and references therein, for a list of previous studies and overview of this class of models with a detailed description of their strengths and caveats.

A common property of these more recent approaches is to make a distinction between tasks and skills. Autor, Levy, and Murnane (2003) shows, for example, that information and communication technologies can substitute for workers who perform routine tasks and at the same time complement workers who perform more complex or non-routine tasks. They also show such shifts in tasks associated with computerization and their effects on education demand can explain 60 percent of the shift in demand for college-educated labor in 1970s and 1980s.

In the same spirit, Acemoglu and Autor (2011) conceptualizes the production process as a set of activities (tasks) that produce output and frame the firm's decision problem as one where workers' capabilities (skills) are allocated to tasks. They argue the conceptual distinction is justified whenever the mapping between skills and tasks is not one-to-one. This would be the case if a certain worker type can be allocated to different tasks, or if the firm can change the task content of production. The latter gives rise to the possibility of an endogenous response of technology to market conditions, automation, or endogenous technology choice (Acemoglu (1998), Acemoglu and Restrepo (2018a, 2018b, 2019a, 2019b, 2020), and Dinlersoz and Wolf (2018)).

It is instructive to compare the task-content approach of Acemoglu and Restrepo (2019b) in terms of the implied production function with those discussed above. They write:

$$Q = \Pi\,(I, N)\left[\Gamma(I, N)^{\frac{1}{\sigma}}\,(A_L L)^{\frac{\sigma-1}{\sigma}} + \big(1 - \Gamma(I, N)\big)^{\frac{1}{\sigma}}(A_K K)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}} \quad (3)$$

where I is innovation, N is the number of tasks, $\Gamma(I, N)$ is the labor task content of production, and $(1-\Gamma(I, N))$ is the capital task content of production. In their model, N tasks are ordered by automatability. All tasks can be performed by either labor or capital. New tasks are represented by an increase in N. Innovation, I, is the cutoff—tasks below I are performed by capital and tasks above I are performed by workers. Thus, innovation (an increase in I) represents an increase in the number and types of tasks that can be performed by capital. The Hicks-neutral productivity term, $\Pi(I, N)$, is assumed to be a function of innovation and tasks as well. This approach highlights that proxies for the number and nature of tasks and the relationship between innovation and tasks are important for understanding the production process. This specification also helps clarify that the firm's choice of a production function is defined by the set of tasks, the level of innovation, and the types of factor inputs.

We draw on these papers to help map out a potential link between differences in worker types and productivity. Such a link is relevant because existing measures of productivity dispersion are calculated without explicitly controlling for possible heterogeneity in labor, even though it would be ideal to control for these differences when calculating establishment-level differences in productivity. We take the first steps in this direction by comparing information on occupations and wages from OEWS and O*NET and exploring the relationship between within-industry dispersion in task/skill intensity and within-industry dispersion in productivity.

3. Data and Measurement

3.1 The Establishment-level Productivity Database

BLS and the Census Bureau have collaborated to create an establishment-level productivity database (called CMP hereafter) for the manufacturing sector from 1972 to 2018.[7] These data are based on the Annual Survey of Manufactures (ASM), the Census of Manufactures (CM), and the Longitudinal Business Database (LBD). The ASM collects data annually and is a five-year panel of manufacturing establishments updated by births in each year.[8] The CM collects data from all manufacturing establishments, except those that are very small, every five

[7] CMP stands for Collaborative Micro-productivity Project.
[8] ASM panels start in years ending in "4" and "9."

years, in years ending in “2” and “7.” The LBD is a longitudinally-linked version of the Census Bureau’s Business Register (see Chow et al. (2021)). It provides high-quality longitudinal links and information on the universe of manufacturing establishments. The CMP includes measures of inputs, output, and productivity (see Cunningham et al. (2022)). The ASM establishments in the productivity database form the basis for creating the DiSP, which we describe below.

3.2 Dispersion Statistics on Productivity (DiSP)

The productivity dispersion measures we use here come from the DiSP, a joint BLS-Census *public-use* experimental data product, which currently includes annual measures for all 86 four-digit NAICS manufacturing industries from 1987 to 2018. These measures tell us how much more productive one establishment is from another between different points in the productivity distribution within four-digit NAICS manufacturing industries. We use the activity-weighted interquartile range (IQR) and 90–10 dispersion measures for both LP and TFP in our analyses. LP is calculated as the log of real revenue per hour, while TFP is calculated as the log of real revenue per combined unit of all factor input costs (capital, labor hours, energy, and materials).[9]

[9] Revenue is based on the value of shipments with adjustments for resales and changes in inventories from the ASM and deflated using the industry implicit price deflator from BLS. We use deflators at the highest level of industry detail available in BLS, BEA, and NBER-CES data. See Cunningham et al. (2022) for more details on deflators and the construction of the DiSP. An important issue is how much observed revenue productivity dispersion reflects output and input price differences across establishments in the same narrow industry. Foster, Haltiwanger and Syverson (2008) finds considerable price dispersion across establishments for the limited number of detailed industries with comparable physical product data in the Census manufacturing data. They find that dispersion in TFPQ (physical productivity adjusting for plant-level differences in prices) is actually larger than dispersion in the type of revenue productivity measures (they refer to these as TFPR) that we use given the negative covariance between prices and TFPQ (i.e., establishments that are more productive have lower marginal costs yielding lower prices). They also find a high positive correlation between TFPQ and TFPR, which is reassuring for the analysis in the current paper. There may be variable markups across establishments in the same industry that influence revenue productivity dispersion measures. Foster, Haltiwanger and Tuttle (2022) shows that identifying such markup variation is complex in the presence of differences in production technologies across establishments. This insight leads us back to the current paper and project seeking to understand such differences by integrating the OEWS microdata with the ASM.

3.3 Task and Skill Concepts

Before discussing our data on tasks and skills, we summarize the basic concepts outlined in section 2, relying on the nomenclature from the Revised Handbook of Analyzing Jobs (Employment and Training Administration (1991)) and Acemoglu and Autor (2011). *Tasks* are activities that when combined with intermediate goods create a good or service and are the true factors of production we would like to measure. However, because we do not observe time spent in different tasks, we use occupations as proxies. An *occupation* is a job in which "a common set of tasks are performed or are related in terms of similar objectives, methodologies, materials, products, worker actions, or workers characteristics" (Employment and Training Administration (1991), p. 9). Thus, an occupation can be thought of as a bundle of tasks. In contrast, a *skill* "is a worker's endowment of capabilities for performing various tasks" (Acemoglu and Autor (2011), p. 1045). Skill is commonly conceptualized in the economics literature as a function of education (and sometimes also experience). Operationally, it is often proxied by some measure of wages projected on observable indicators such as education and experience or, alternatively, wages are projected on occupations as in Acemoglu and Autor (2011) (see Figure 10 therein). While there is not a one-to-one correspondence between skills and tasks in this framework, complex tasks generally require greater skills. The relationship between skills and tasks can vary over time and across businesses, presenting a challenge for productivity measurement and highlighting a need for detailed data on tasks and skills.

3.4 Occupational Employment and Wage Statistics (OEWS) Data

We use occupation data from the BLS Occupational Employment and Wage Statistics (OEWS) survey, which is a semi-annual mail survey sampling approximately 200,000 establishments in May and November of each year.[10] This survey covers all workers, both full time and part time, in private non-agricultural industries.

The survey instrument asks establishments to provide what amounts to a complete payroll record for the pay period that includes the 12th of the sample month. Respondents report occupational wage information for each occupation by recording the number of employees in each of 12 wage intervals.[11] The OEWS survey uses the Office of Management and Budget's occupational classification system, the Standard Occupational Classification (SOC), to categorize workers into around 800 detailed occupations.[12] The SOC system provides much more occupational detail than do household surveys, which are the main sources of information about occupations.

The sample is composed of certainty units, which are generally sampled every three years, and non-certainty units. Official estimates are based on data from the current panel and the previous five panels (because the OEWS is typically released in May, the five panels extend back to the November panel three years earlier). Thus, although estimates are published every year, the OEWS data are not a true time series, because about two-thirds of the sample are the same in any two adjacent years.

---

[10] 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 and assign it the May year value. For this reason, we do not have data for 2002.

[11] Wages for the OEWS survey 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.

[12] From 1999 to 2013, the SOC structure has expanded from 770 occupations to its current 821 occupations.

The OEWS sampling and weighting methods guarantee total weighted employment equals the BLS frame—the Quarterly Census of Employment and Wages (QCEW)—employment, but there is nothing in the methods to guarantee the implied number of establishments equals the number of establishments on the frame. This makes it difficult to develop dispersion statistics at the establishment level. Therefore, any analysis attempting to measure establishment-specific effects will have to address this feature of the OEWS weighting scheme. As an alternative to reweighting the data, we use a research dataset created using a modified version of the imputation approach developed by Dey, Piccone, and Miller (2019).

This research dataset imputes data for the entire QCEW. For each reference year, they use the same dating convention as is used for the official OEWS release (that is, May of the reference year combined with the five previous panels). For each observation in the QCEW not in the OEWS, they identify 5–10 donor observations based on the characteristics of the establishments. The characteristics include employment, industry (six-digit NAICS), ownership, metropolitan statistical area (MSA), and the amount of time between reference periods of the observations. Donor establishments are evaluated on each attribute and weights are assigned based on closeness to the recipient on that attribute. In the experimental data series, the weights of the donor establishments are rescaled so they sum to one. The recipient's employment in each occupation is a weighted average of the donor establishments. Wages are determined similarly but are also adjusted for differences in wages by area and wage growth by area and industry.[13]

The Dey, Piccone, and Miller (2019) approach is designed as a potential replacement for the current method for generating official estimates. The main advantage of this approach is every establishment in the QCEW is represented and has an establishment weight of one. The

[13] These adjustments are not controls for industry and location. Rather, they are designed to convert the wages of the donor observations so they more-closely approximate the recipient establishment's actual wages.

disadvantage is that the staffing pattern for an establishment is an *average* of similar establishments. This makes sense for constructing aggregate estimates, but not for analyzing distributions. The research dataset incorporates two key modifications to this method for our analysis, which focuses on distributions.

The primary modification to this method is that occupation employment and wage data at the establishment-level are imputed from a single donor. 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 very 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 dispersion, the selection of a particular donor from the set of acceptable matches is random. As above, wages are adjusted to account for differences by MSA and industry.

The second modification that we make is to center the sample on the reference year instead of using data from the five panels prior to May of the reference year (as in the published statistics), so we can integrate the OEWS with other data. 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, we construct these "year samples" at three-year intervals. We are effectively assuming the occupational mix within an establishment is fixed over the three-year interval.

The result is a full dataset that we can use to estimate dispersion statistics on an establishment-weighted basis (weight = 1) or an employment-weighted basis (weight = employment). Our sample covers the years 2000, 2005, 2008, 2011, 2014, and 2017.[14] Much of our analysis is restricted to establishments in manufacturing industries, so we can make comparisons with the DiSP.

3.5 Bundled Task/Skill Intensity Index (TSB): Counterfactual Wages

Our first index of task/skill intensity is a counterfactual wage equal to the average wage paid by the establishment if the establishment paid the national average occupational wage for all workers in each occupation for each year in the sample. Thus, it accounts for differences in the occupational mix across establishments by attaching a different price to each occupation. By using the national average wage for each occupation, the price of each occupation is the same across establishments. We denote this as a "bundled" task/skill intensity index (TSB) because tasks are bundled into occupations.

Let w_{ejb} and L_{ejb} denote the mean log wage and the number of workers in wage interval b of occupation j in establishment e, respectively. All workers in the same wage interval are assigned the same log wage. Suppressing the time subscript for simplicity, the national mean log wage for occupation j is given by:

$$\overline{w}_{nj} = \frac{1}{\sum_{e \in E_n} L_{ej}} \sum_{e \in E_n} \sum_{b \in B} (w_{ejb} \times L_{ejb}) \tag{4}$$

[14] The five-year gap between 2000 and 2005 is due to a change in sampling from annual to semi-annual, which made it impossible to construct estimates for May 2002.

where B is the set of 12 wage intervals, E_n is the set of all establishments (nationwide, manufacturing, and non-manufacturing), and L_{ej} is the number of employees in occupation j for establishment e. The actual mean log wage for an establishment is:

$$\bar{w}_e = \frac{1}{L_e}\sum_{j\in J_e}(\bar{w}_{ej} \times L_{ej}) \tag{5}$$

where J_e is the set of occupations employed by establishment e, L_e is total employment in establishment e, and $\bar{w}_{ej} = \frac{1}{L_{ej}}\sum_{b\in B}(w_{ejb} \times L_{ejb})$. Substituting the national mean log occupational wage, $\bar{w}_{nj}$, for the establishment mean log wage for each occupation, the counterfactual mean log wage for establishment e, $\widetilde{w}_e$, is equal to:

$$\widetilde{w}_e = \frac{1}{L_e}\sum_{j\in J_e}(\bar{w}_{nj} \times L_{ej}) \tag{6}$$

The TSB measure is a simple measure that provides an index of the tasks employed by the establishment using wages, which proxy for skills, to price those tasks. Although it is a useful measure, it has the property that it does not distinguish between different occupations (with different task sets) paying the same wage. Thus, two establishments might have the same task/skill intensity but very different mixes of occupations. To illustrate this, Figure 1 plots the TSB measure (on the vertical axis) against an index of the dissimilarity of the occupational mix of the establishment relative to its industry (on the horizontal axis) for establishments in four industries: basic chemicals manufacturing, computers and peripherals manufacturing, semiconductor manufacturing, and big box retailers. The dissimilarity index we use is the absolute value of the sum over all occupations (2-digit SOC) of the distances between the establishment's payroll share for that occupation and the industry-wide payroll share for that occupation.[15] It takes on values between zero and one, with higher values indicating an establishment has a much different occupational distribution than the typical establishment in the industry.

Perhaps the most notable feature of the graphs is the fanning out of the TSB as the dissimilarity index increases. Thus, establishments with occupation mixes differing significantly from the industry average mix also have a wide range of task/skill intensities. But more important for our purposes, for a given level of task/skill intensity, there is considerable variation in the occupational mix. For example, in the basic chemicals manufacturing industry, at the mean skill intensity of around 3.2, the dissimilarity index varies from about 0.1 to just over 0.75. There is less variation in the TSB measure and in the dissimilarity index in the big box retail industry than in the other three industries. But there is still considerable variation in the dissimilarity index for a given level of TSB. Thus, although the TSB tells us a lot about differences across establishments in the types of occupations that they employ, it does not account for all the variation in how establishments organize production. There is still much to be learned from looking at differences in the distribution of occupations across establishments.

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

Our second task/skill intensity index builds on Acemoglu and Autor (2011), who uses O*NET data to operationalize the Autor, Levy, and Murnane (2003) taxonomy of tasks

---

[15] The dissimilarity index we use is:

$$D_{ie} = \frac{1}{2}\sum_{j \in J_i}\left|\frac{\bar{w}_{ej}L_{ej}}{\sum_j \bar{w}_{ej}L_{ej}} - \frac{\bar{w}_{ij}L_{ij}}{\sum_j \bar{w}_{ij}L_{ij}}\right|$$

where the variables are defined as before and the *i* subscript indicates industry. The index is scaled to represent the fraction of payments to the different occupations that would have to be reallocated to match the industry distribution of payments across occupations. It is worth noting this index is sensitive to the level of occupational detail. The index will be larger the greater the level of occupational detail. We used two-digit occupation codes, which are fairly aggregated, to calculate these indexes. In addition, we restricted the sample to establishments with 20+ workers.

(developed with the O*NET predecessor, the Dictionary of Occupational Titles (DOT)). Autor, Levy, and Murnane developed a two-dimensional categorization of tasks based on whether they are 1) routine or non-routine and 2) cognitive or manual. Routine tasks are those that can be described using a set of rules or specifications; non-routine tasks are those that cannot be described in this manner. They further break down non-routine cognitive tasks into analytic and interpersonal. This yields five categories of tasks: non-routine cognitive (analytical), non-routine (interpersonal), routine cognitive, routine manual, and non-routine manual physical.[16]

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 Research Triangle Institute. The O*NET data are collected from workers in targeted occupations at establishments and contain over 275 variables that describe each occupation.[17] Acemoglu and Autor (2011) uses 16 of these variables corresponding to the five task categorizations described above.[18] 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. The five indexes are created by

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

[17] O*NET first began surveying job holders in 2001. Prior to that, past DOT 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 re-surveys 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 job holders who answered that question, so within-occupation variation cannot be observed. See Handel (2016) for more about the history of O*NET and its strengths and weaknesses.

[18] 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. 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. 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. (See page 1163 of Acemoglu and Autor (2011).)

summing the standardized variables for each task category, which are then once again normalized.

We use this methodology to create the same five task indexes for each of the O*NET years where the index variables are available for most occupations (2007, 2008, 2014, and 2017).[19] We merge these five task indexes to OEWS wage data by occupation and estimate a regression of the national occupational mean log wage for each year on these five task indexes as follows:

$$\overline{w}_{nj} = \alpha + \sum_{k=1}^{5} \beta_k x_{jk} + \varepsilon \tag{7}$$

where τ_{jk} is the O*NET measure of task k for occupation j, and $\overline{w}_{nj}$ is defined in equation (4).[20] The coefficients on the task indexes, β_k, are akin to prices in a hedonic regression. We then calculate the counterfactual average establishment wage as:

$$\widehat{w}_e = \frac{1}{L_e} \sum_{k=1}^{5} \hat{\beta}_k \left[\sum_{j \in J_e} \left(L_{ej} \times \tau_{jk} \right) \right] \tag{8}$$

where the summation in square brackets is the total amount of task k employed by the establishment and $\hat{\beta}_k$ is the "price" of task k estimated from the regression in equation (7). That is, the TSU measure can be thought of as the average price of tasks performed by employees in the establishment.

We refer to this second measure as an "unbundled" task/skill intensity index (TSU) because tasks (weighted by prices) are aggregated without accounting for how the tasks are bundled into occupations. In contrast, TSB uses the occupational mix of an establishment (and the prices of such occupations), so it implicitly takes into account that individual occupations

[19] 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.

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

reflect a bundle of tasks. Like the TSB index, there are many combinations of tasks that can result in the same value of the index. We discuss these differences and similarities further below.

3.7 Individual Average Task Indexes

In addition to the two task/skill intensity measures based on counterfactual wages, $\widetilde{w}_e$ and $\widehat{w}_e$, we also develop a set of five task measures based on the average value of the individual O*NET task indexes. Recall, we are using five categories of O*NET tasks: non-routine cognitive (analytical), non-routine (interpersonal), routine cognitive, routine manual, and non-routine manual physical. For each of the five task indexes, we measure an employment-weighted establishment-level average for task index *k* as follows:

$$\bar{\tau}_{ek} = \frac{1}{L_e} \sum_{j \in J_e} \tau_{jk} \times L_{ej} \tag{9}$$

where k = 1, …, 5. Thus, $\bar{\tau}_{ek}$ is the average task *k* content of all jobs in establishment *e*. Again, time subscripts are suppressed for expositional convenience. These measures are constructed for each establishment for each year in our sample.

3.8 STEM Intensity Index (%STEM)

In our final measure, we rely on the definition of occupation as a bundle of tasks and create another alternative task index based on occupation data in the OEWS. We calculate the percentage of STEM workers in each establishment as follows:

$$\%STEM_e = \frac{1}{L_e} \sum_{j \in J_s} L_{ej} \tag{10}$$

where J_s is the set of STEM occupations, with STEM occupations being defined according to the recommendations of the SOC Policy Committee (2010). The %STEM equals the percentage of workers in an establishment who are working in the following sub-domain occupations—life and physical science, engineering, mathematics, and information technology occupations, social science occupations, architecture occupations, health occupations—within the following larger occupation groups: research, development, design, or practitioner occupations, technologist and technician occupations, postsecondary teaching occupations, managerial occupations, and sales occupations. The STEM measure is constructed for each establishment for each year in our sample.

3.9 How Do These Task/Skill Measures Differ?

As discussed above, the TSB measure prices the tasks of each occupation as a bundle and therefore indirectly accounts for the fact that the sets of tasks that make up an occupation are complementary and there is a benefit to having them performed by the same person. In contrast, the TSU measure prices the tasks individually and ignores any complementarities between tasks within occupations. We would expect the two measures to be different but highly correlated.

The first row of Table 1 shows average Pearson correlations between the establishment-level TSB and TSU task/skill indexes for different major sectors. The table entries are the employment-weighted averages of the within-industry correlations pooled over time. As expected, the correlations are high, although there is some variation across sectors. For example, the correlation is higher for manufacturing than for non-manufacturing. There is also a sizeable difference between high-tech and non-tech manufacturing industries, with the correlations being considerably higher for high-tech industries ($\rho = 0.911$).[21]

The next five rows of Table 1 show the correlation of TSB with each of the five task groups. These correlations are insightful because they reveal which type of tasks are more

[21] Following Wolf and Terrell (2016), we define the high-tech industries as those industries whose share of STEM workers exceeds 2.5 times the national average. This group includes the following 16 four-digit NAICS manufacturing industries: petroleum and coal products; basic chemical; resin, synthetic rubber, and artificial and synthetic fibers and filaments; pharmaceutical and medicine; industrial machinery; commercial and service industry machinery; engine, turbine, and power transmission equipment; other general purpose machinery; computer and peripheral equipment; communications equipment; audio and video equipment; semiconductor and other electronic components; navigational, measuring, electromedical, and control instruments; manufacturing and reproducing magnetic and optical media; electrical equipment manufacturing; aerospace products and parts.

strongly related to the composite index task/skill intensity measures. These individual task indicators also shed light on the high correlation between the TSB and TSU measures. Looking at the second row, we see the correlation between TSB and the O*NET "analytical tasks" measure is nearly the same as the correlation between the TSB and TSU measures. The correlation between the TSB measure and the O*NET "interpersonal tasks" measure is lower, but still high. The other three individual task indicators are negatively correlated with TSB. That is, establishments that have a high composite task/skill intensity generally employ fewer occupations heavy in these tasks.

In the last row of Table 1, we show the correlation between %STEM and TSB. Again, the correlations are higher for manufacturing than for non-manufacturing and higher for high-tech than non-tech manufacturing industries, with the correlations being considerably higher for high-tech industries.

3.10 Dispersion in Tasks/Skills

We calculate two measures of dispersion in establishment-level tasks/skills—the interquartile range (IQR) and interdecile (90–10) range—for each four-digit NAICS industry in each sample year. To account for industry differences in average skills/tasks so we can compare within-industry dispersion across industries and time, we calculate establishment-level tasks/skills as the deviation from the average tasks/skills in that establishment's four-digit industry. These measures are weighted using establishment employment. We then calculate our dispersion measures, which tell us the degree of within-industry dispersion. For this analysis of summary statistics, we also include the activity-weighted within-industry dispersion in productivity measures from DiSP (see Cunningham et al., 2022, for details). Given the DiSP measures are only available for manufacturing industries, the summary statistics in Table 2 are only for manufacturing.

Table 2 shows summary statistics for our within-industry dispersion measures (productivity and task/skill) for four-digit manufacturing industries. As before, these measures are industry averages pooled over all years. In Figures 2A and 2B, we show the means of the IQR dispersion measures for productivity and select task/skill intensity indexes over time for high-tech and non-tech industries. For all our measures, we find dispersion is much greater among the high-tech manufacturing industries than the non-tech manufacturing industries. Among the high-tech manufacturing industries, we also observe a gradual increase in each dispersion measure over time, widening the cross-industry dispersion over time. Rising within-industry dispersion in task/skill intensity measures is also present in the non-tech industries, with the notable exception of the %STEM dispersion. Perhaps not surprisingly, there is very little dispersion in %STEM dispersion among the non-tech manufacturing industries.

In Figures 3A and 3B, we show the standard deviations in these dispersion measures, i.e., the dispersion in dispersion. We see higher dispersion in dispersion in the high-tech manufacturing industries than the non-tech manufacturing industries for five of the six measures shown. Figures 2 and 3 highlight there is dispersion in dispersion both in the cross-section and over time.

In Table 3, we show the Pearson correlations between the industry-level TSB dispersion measures and the other task/skill dispersion measures, pooled across years. Panel A shows the correlations between the IQR dispersion measures, while Panel B shows the correlations between the 90–10 dispersion measures. In both panels, we see that among all industries, the highest correlations are between dispersion in the TSB and dispersion in the TSU, the analytical tasks measure, and the interpersonal tasks measure. An industry with above-average TSB dispersion is likely to exhibit above-average dispersion in each of these task/skill measures. Among non-manufacturing industries (last column), dispersion measures for these task/skill measures are similarly highly correlated. Among manufacturing industries (second column), the correlations between the TSB and the TSU, analytical tasks, routine manual tasks, non-routine manual physical tasks, and %STEM all exceed 0.5. The interpersonal and routine cognitive tasks dispersions are less likely to be important for TSB dispersion in manufacturing relative to non-manufacturing industries. The main difference between the high-tech and non-tech manufacturing industries is that dispersion in TSU, analytical tasks, non-routine manual physical tasks, and %STEM varies more closely with dispersion in TSB in the high-tech group, while dispersion in routine manual tasks and interpersonal tasks varies more with TSB dispersion in the non-tech group. Those relationships are much stronger for the interdecile range than for the IQR.

With these summary statistics as background, we now turn to the analysis of primary interest—the relationship between within-industry dispersion of productivity and task/skill intensity measures.

4 The Relationship between Productivity, Skills, and Tasks

To analyze the link between productivity, skills, and tasks, we consider the relationship between within-industry dispersion measures of productivity, skills, and tasks. Our analysis is descriptive and provides no causal interpretation. As discussed in section 2, the relationship between productivity, skills, and tasks likely reflects endogenous relationships between the choice of tasks, factor mixes, and innovation, where innovation can be either product-quality-enhancing or process-enhancing.

We first calculate Pearson correlations between our dispersion measures for manufacturing industries (all, high-tech, and non-tech). In Table 4, we see the TSB dispersion measure has the highest correlations with our productivity dispersion measures for the entire manufacturing sector, but the %STEM, analytical task, and TSU dispersion measures also have strong positive relationships. Looking at high-tech and non-tech manufacturing industry groups separately, we see the TSB, TSU, and %STEM dispersion measures have higher correlations with the productivity dispersion measures for the high-tech industries than for the non-tech industries. These patterns hold for both the IQR and 90–10 dispersion measures.

It is notable that the correlations for all manufacturing industries are higher than for either high-tech or non-tech industries. This is because the correlations for all manufacturing industries are not an average of the high- and non-tech industry correlations. Rather, they also include differences in dispersion between high-tech and non-tech industries.

Tables 5 and 6 provide an alternative view of the relationships in Table 4. These tables present the coefficients, standard errors, and R-squared values from regressions of industry-specific productivity dispersion measures on task/skill-dispersion and task-dispersion measures. All regressions are estimated using ordinary least squares (OLS) by pooling data across years while controlling for year effects. The statistics from these regressions are useful because the coefficients are closely related to the Pearson correlations in Table 4 and inform us about the explanatory power of the variation in skills and task-content for productivity dispersion. Again, we see positive and statistically significant associations between LP dispersion and the TSB, TSU, and %STEM dispersion measures, with stronger relationships among the high-tech industries relative to non-tech industries.

These patterns are similar for TFP dispersion, but the coefficients are smaller in magnitude. A possible explanation for this lies in the fact that LP does not account for the mix of factors used—e.g., capital, which can be substituted for labor. Because factor inputs are embodied in output, all else equal, we would expect LP to be higher in establishments that are more intensive in other inputs. This implies that LP dispersion partly reflects dispersion in, for example, capital intensity. If workers at capital-intensive establishments are more skilled and engage in more complex tasks, then the establishments would have higher TSB. In terms of the second moment, higher dispersion in capital intensity can yield higher dispersion in LP. This discussion highlights reasons why the relationship between TSB dispersion may be weaker with TFP dispersion compared to LP dispersion. We plan to investigate such hypotheses with integration of the OEWS and CMP microdata.

The magnitude of the estimated coefficients in Table 5 varies considerably across different groupings of industries as well as the alternative task/skill intensity. For example, the elasticity of within-industry TFP dispersion with respect to the TSB dispersion measure (evaluated at means) is more than three times larger in the high-tech grouping of industries than in non-tech industries.[22] This is also true for the TSU and the analytical tasks dispersion measure. In contrast, non-routine manual physical task dispersion is substantially more related to TFP dispersion in the non-tech industries than the high-tech industries.

The regression results for the 90–10 dispersion measures in Table 6 yield even larger quantitative effects. For example, the elasticities of within-industry TFP dispersion with respect

[22] The elasticities can be computed combining information from Tables 2 and 5. For example, the estimated coefficients relating the IQR of TFP to the IQR of the TSB measure are 2.30 and 0.77 for high-tech and non -tech, respectively. Using the means from Table 2, the implied elasticities are 1.20 and 0.32, respectively.

to both the TSB and TSU dispersion measures (evaluated at means) are more than six times as large in the high-tech industries as in the non-tech industries.

Our results can be thought of as complementary to those of Cunningham et al. (2022) who finds that establishment characteristics from the firm dynamics literature (state, age class, and size class) have limited explanatory power for productivity dispersion. In particular, the results in this paper suggest that we will have a better chance of understanding productivity differences if we go beyond the standard establishment characteristics and look at the basic characteristics of workers and tasks.[23]

The explanatory power of the regressions estimated over all manufacturing industries indicates that variation in skills and task content is relevant for productivity dispersion. For example, the first row R-squared in Table 5 shows that the TSB dispersion accounts for about a quarter of the variation in LP dispersion and about one-fifth of the variation in TFP dispersion across manufacturing industries. The second R-squared in column 1 suggests that TSU dispersion accounts for almost one-fifth of the total variation in LP dispersion and about one-tenth of the variation in TFP dispersion. Out of the five task measures, dispersion in analytical task-content is the most relevant for productivity dispersion, indicated by more positive estimated coefficients and mostly higher R-squared values.[24]

[23] Appropriate caution is needed in making these comparisons, because Cunningham et al. (2022) investigate the relationship between establishment-level productivity and characteristics, while here we are relating within-industry dispersion in productivity and task/skill characteristics.

[24] Our results may be affected by the level of detail in the industry classification, because the regression coefficients in Tables 5 and 6 depend on productivity dispersion and the dispersion measures depend on the level of detail at which they are calculated. DiSP is based on four-digit NAICS, at which level the average within-industry standard deviation of TFP over the 1997–2015 period is 0.46 (see Table 3 in Cunningham et al. (2022). Using six-digit NAICS, Decker et al. (2020) finds this statistic for the relevant TFP measure is 0.35–0.37 (see Panel A of their Figure 3). Blackwood et al. (2021) finds similar values using four-digit SIC and six-digit NAICS but for a different time-period. Their standard deviation measure is 0.31 (see their Table 1). In other words, more detailed industry codes imply lower dispersion, which is intuitive. However, these dispersion measures are similar in magnitude, which suggests using detailed industry codes is likely to have a limited effect on our findings.

The R-squared values for the high-tech and non-tech industries are lower those for manufacturing overall. As we saw in Table 4, this implies that it is differences in dispersion between high-tech and non-tech industries driving much of this association rather than within sector differences in dispersion.

Figure 4 illustrates the relationships between TFP IQR dispersion and each task/skill IQR dispersion measure by industry (pooled over the years in our sample) for selected measures and also includes the slope of the relationships in 2000 and 2017. Focusing first on the TSB task/skill intensity dispersion measure, we find a positive relationship in the pooled data, but we find the relationship changes over time depending on whether we look at the high-tech or non-tech industries. For high-tech industries, the slope was slightly negative in 2000 but strongly positive in 2017. On the other hand, for the non-tech industries, the relationship was positive in both years but weakened between 2000 and 2017. We find similar patterns using the analytical tasks dispersion. Also, looking at high-tech industries, we find the correlation between TFP dispersion and %STEM dispersion becomes strongly positive in 2017; but for the non-tech industries, there is little dispersion in %STEM and no change over time in the slope.

Patterns are very different using the non-routine manual physical task dispersion measures. Between 2000 and 2017, the negative slope between the TFP dispersion measure and the physical task dispersion measure becomes weaker among the high-tech industries. Among non-tech industries, the relationship is strongly positive, although it is weaker in 2017 than in 2000. When pooling across all industries, we find a strong positive relationship that is similar in both 2000 and 2017. Pooling across industries yields a consistently positive relationship, which highlights again that there are some interesting differences between high-tech and non-tech industry effects at work.

The especially high degree of within-industry dispersion in productivity and task/skill intensity measures and their strong positive relationship in the high-tech industries is striking. It is already well known that high-tech industries have higher than average within-industry productivity growth and higher than average skill/task intensity (as measured for example by the STEM intensity of workers in the industry—see, e.g., Decker et al. (2020)). Novel to our analysis is that these first moment relationships have related analogues in within-industry second moments. These innovation-intensive industries exhibit high dispersion in productivity outcomes across businesses accompanied by indicators that these businesses are organized quite differently in terms of their mix of workers. While our results are only suggestive, they offer prima facie evidence that there is likely a high payoff to integrating productivity and occupation data at the establishment level. We turn to the prospects of that integration in the next section.

5 Integration of Establishment-Level Productivity and Occupation Data

5.1. Challenges in Integrating the OEWS and CMP Microdata

The next step in the larger project is to integrate the establishment-level data in the CMP and the occupation data in the OEWS. In this section, we discuss the challenges that will need to be overcome in this data integration project, along with plans for the types of analyses that can be conducted with this data.

One challenge is the OEWS is drawn from the BLS business register, while the underlying source CMP data are drawn from the Census Bureau Business Register. The common identifier on both files is the Employer Identification Number (EIN). Studies have shown that there is a high match rate of EINs across the registers (Fairman et al. 2008; Haltiwanger et al.

2014). Name and address matching can facilitate the next step of establishment-level matching within EINs, but we anticipate that instructive analysis can be conducted using the EIN matches.

When matching OEWS data to the CMP data, we will need to account for the fact that both are samples. We will take two approaches to the analysis. Initially, we will restrict the OEWS sample to non-imputed observations. Because this will result in a large number of non-matches, we will combine the matched CMP-OEWS data with the Census Bureau's Business Register to estimate propensity score weights along the lines of Cunningham et al. (2022). Unlike most reweighting situations, the assumption that data are randomly missing holds—missing data are due to randomness in the OEWS and ASM sampling procedures rather than to the behavior of establishments. It is worth noting the CMP data are essentially a universe for productivity in Economic Census years. Thus, in Census years, we are less constrained by the difference in the sampling of establishments in the OEWS and the ASM.

The second approach is to use the full QCEW-OEWS dataset that we use here (that is, including imputed observations). This will result in a match rate of nearly 100 percent and a substantially larger sample. We recognize the imputed QCEW observations may give a distorted picture of the relationship between productivity and the occupation mix. One way to examine this is to see how much the occupational mix varies across establishments within the same EIN in the two OEWS samples. Put differently, it is important to examine whether using imputed observations affects the ratio of within-EIN versus between-EIN variation. Our preliminary analysis discussed in the next subsection reveals some differences but suggests both approaches can shed light on the question.

5.2. Within- vs Between-EIN Variation in Skill and Task Indexes Across Establishments

We compute variance decompositions of the between-establishment distributions of the different task and skill measures for each industry in terms of the fraction of the variance accounted for by between-EIN variation. Figure 5 shows box-and-whisker graphs of the between-EIN share of total variance for the seven measures in the full (including imputations) QCEW-OEWS dataset using all manufacturing industries. Because the shares are very similar across measures, the rest of the analysis will focus on the TSB measure. Figure 6 presents four box-and-whisker graphs showing the distribution of the share of between-EIN variation in the four pooled samples defined by (1) whether imputations are included and (2) whether single-unit EINs are excluded. The left two graphs show the between-EIN shares are very high when we include all EINs (single- and multi-unit EINs) and the between-EIN share of total variation is smaller when we include imputed observations. The right two graphs show the between-EIN shares for multi-unit EINs. As expected, the between-EIN shares are lower than when we include single-unit EINs (within-EIN variation is zero for these firms). As with the full sample, the between-EIN share of total variation is lower when we include imputed observations. Put another way, the imputations increase the within-EIN share of total variation in the TSB measure.

Figures 7 and 8 expand on the distributional statistics shown in Figure 6 by plotting the between-EIN share of TSB variation when using all imputed data versus the between-EIN share when using only direct or EIN matches. Each dot represents a four-digit industry. Figure 7 shows the relationship for all EINs while Figure 8 shows the relationship for multi-unit EINs only. For all EINs, whether we include imputed data or not has very little impact on the between-EIN share of variation. There are more industries where the between-EIN share is lower when including the imputed data, but the observations generally lie fairly close to the 45-degree line.

On the other hand, when examining only multi-unit EINs in Figure 8, including imputed data seems just as likely to increase or decrease the between-EIN share of variation, and the data are not particularly close to the 45-degree line.

These findings are interesting in their own right because they highlight that there is an important common component of task and skill mix across establishments within the same firm. The findings also will greatly facilitate the planned integration of the OEWS and CMP data, because the EIN is a common identifier on the files.

5.3. Planned Future Analysis

We will start our planned future analysis of the integrated OEWS and CMP data by examining the joint distribution of establishment-level productivity and our task/skill measures. This exploratory data analysis will examine the joint distribution of productivity and occupational mix through the type of correlation analysis above as well as cluster analysis looking to see whether we can identify distinct patterns relating the productivity and occupational distributions. Exploring the joint distribution will be useful for understanding productivity dispersion given DiSP does not account for differences in the occupational mix. For example, if the occupation distribution in one group of establishments within an industry is skewed towards high-skilled labor while it is skewed towards low-skilled labor in another industry, then it is reasonable to calculate dispersion accounting for such differences either by adjusting the labor input measure or allowing elasticities to vary across clusters, or both. How such clusters should be identified is an open question, but information on occupations from OEWS could provide insights about whether businesses use labor differently within an industry.

For our more structured analysis, we will explore a range of alternative production technologies as considered in section 2. A first step is to use the existing DiSP approach but replace total hours with task/skill-adjusted hours. To fix ideas, consider the specification of production in section 2: $Q=A \cdot F(Z \cdot L,K)$ where time and establishment subscripts are suppressed for simplicity. One approach would be to use the TSB and TSU measures as proxies for Z. As discussed above, TSB and TSU are task/skill-based indexes that take into account the pricing of those tasks. TSB uses the bundling of tasks through occupations, while TSU uses the pricing of identified tasks associated with the occupation mix. Such an approach is related to but distinct from the discussion in section 2 on creating an efficiency units measure of the labor input. An interesting question is: does measured productivity dispersion decline when we account for differences in the task-content of jobs across establishments? Just as TFP dispersion is less than LP dispersion, we would expect that accounting for task/skill differences would further reduce TFP dispersion. In addition, differences between the TSB and TSU measures can tell us something about the importance of task bundling.

Building on this approach, we plan to estimate a range of functional forms with production technologies specified with tasks. One challenge is how to specify a production technology with multi-dimensional task inputs along with departures from Cobb-Douglas functional forms. A key issue in this regard is that under Cobb-Douglas, the elasticity of substitution between factors is by definition one; deviating from this assumption allows the elasticity to be estimated. Existing studies suggest such generalizations are likely important in this context (see, e.g., Dinlersoz and Wolf (2018), Raval (2019), Oberfield and Raval (2021)). The CES models in section 2 are a good starting place given the recent insights of Acemoglu and Restrepo (2019b) (e.g., see equation (3), which is from their paper). We anticipate considering even more flexible functional forms will be of interest. Foster, Haltiwanger, and Tuttle(2022) found much of the rising mean and variance of measured markups using the De Loecker, Eechkhout, and Unger (2020) production-function-based approach can be accounted for (within manufacturing) by permitting factor elasticities to vary across establishments and time in a flexible manner (e.g., a translog approach with time-varying coefficients). Their findings are relevant for our research because measured markup dispersion is closely related to measured revenue productivity dispersion observed in the DiSP.

A challenge for estimation of the production technology is, as emphasized in section 2, that all the factor inputs, tasks, and adoption of specific process- or product-enhancing innovation are endogenous. As such, estimating the production function using OLS is not informative. Much of the recent literature has used control function methods (e.g., De Loecker et al. (2016), De Loecker, Eechkhout, and Unger (2020), Blackwood et al. (2021), Eslava and Haltiwanger (2021)), which is where we plan to start. A complicating issue is that in the absence of establishment- or firm-level prices, we will need to estimate the revenue function and account for the endogeneity of prices including markups. While a challenge, this issue enhances the interest in helping us understand the causes and consequences of businesses doing business differently.

In related but distinct set of exercises, we will explore the relationship between observable indicators of technology and changes in the mix of skills and tasks. Berman, Bound, and Griliches (1994) and Dunne, Haltiwanger, and Troske (1997) pursued early versions of such approaches with much cruder data. Dinlersoz and Wolf (2018) is a more recent example where establishment-level indicators of automation are combined with CMP data from a set of manufacturing industries. With integrated CMP-OEWS data, there are rich possibilities. CMP data include time-series-consistent measures of capital, and the ASM contains periodic indicators of the use of advanced technologies (e.g., computer investment). In addition, in more recent years, the Annual Business Survey (ABS) has modules on the adoption of advanced technologies. Analysis of those modules by Zolas et al. (2020) highlights that the adoption of advanced technologies such as robotics, automation, cloud computing, and artificial intelligence is both relatively rare and highly heterogeneous across businesses within the same industries.[25] Using the CES specifications in section 2 combined with Shepherd's lemma will permit us to examine the relationship between skill and task mix and the adoption of technologies.

It is likely that factors other than the uneven patterns of adoption of advanced technologies underlie the different choices that businesses make when they organize production. Offshoring and outsourcing are likely at work as well. Of course, the uneven impact of globalization across establishments is likely related to the uneven impact of adoption of advanced technologies.

6 Conclusion

Our analysis has barely opened the black box of tasks and skills at the establishment level, but our results suggest that we can gain considerable insights about a potentially important source of differences in measured productivity across establishments within an industry and about the mix of tasks and skills being used in those establishments. Our evidence is only suggestive, as we find industries with high dispersion in within-industry productivity also exhibit high dispersion in a number of the task/skill intensity measures we construct. But we have good reason to believe that we will find a positive relationship between establishment-level productivity and our establishment-level task/skill measures.

[25]Although the data used by Dinlersoz and Wolf (2018) are from the 1990s (Survey of Manufacturing Technologies), they arrive at similar conclusions about robots, automation, and advanced technologies.

More progress on opening the black box awaits integration of the CMP and OEWS data. As we have described, such integration holds great promise for helping us understand how different businesses do business differently. Moreover, opening this black box will also permit a rich exploration of the connection between the adoption of advanced technologies and their impact on the workforce.

References

Acemoglu, Daron. 1998. “Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality.” *The Quarterly Journal of Economics* 113 (4): 1055–89.

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.

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.

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.

Acemoglu, Daron, and Pascual Restrepo. 2020. “Robots and Jobs: Evidence from US Labor Markets.” *Journal of Political Economy* 128 (6): 2188–244.

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.

Baily, Martin Neil, Charles Hulten, and David Campbell**.** 1992. “Productivity Dynamics in Manufacturing Plants.” *Brookings Papers on Economic Activity: Microeconomics*: 187–249.

Berman, Eli, John Bound, and Zvi Griliches. 1994. “Changes in the Demand for Skilled Labor within U. S. Manufacturing: Evidence from the Annual Survey of Manufactures.” *Quarterly Journal of Economics* 109 (2): 367–97.

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.

Card, David and John E. DiNardo. 2002. “Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles.” *Journal of Labor Economics* 20 (4): 733–83.

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.

Cunningham, Cindy, Lucia Foster, Cheryl Grim, John Haltiwanger, Sabrina Wulff Pabilonia, Jay Stewart, and Zoltan Wolf. 2021. “Chaos Before Order: Productivity Patterns in U.S. Manufacturing.” *International Productivity Monitor* 41: 138–52.

Cunningham, Cindy, Lucia Foster, Cheryl Grim, John Haltiwanger, Sabrina Wulff Pabilonia, Jay Stewart, and Zoltan Wolf. 2022. "Dispersion in Dispersion: Measuring Establishment Level Differences in Productivity." Forthcoming *Review of Income and Wealth*.

Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, and Javier Miranda. 2020. "Changing Business Dynamism and Productivity: Shocks versus Responsiveness." *American Economic Review* 110 (12): 3952–90.

Dey, Matthew, David S. Piccone, Jr., and Stephen M. Miller. 2019. "Model-Based Estimates for the Occupational Employment Statistics Program." *Monthly Labor Review* August.

De Loecker, Jan, Pinelopi K. Goldberg, Amit K. Khandelwal and Nina Pavcnik. 2016. "Prices, Markups, and Trade Reform." *Econometrica* 84 (2): 445–510.

De Loecker, Jan, Jan Eeckhout, and Gabriel Unger. 2020. "The Rise of Market Power and the Macroeconomic Implications." *The Quarterly Journal of Economics* 135 (2): 561–644.

Dinlersoz, Emin and Zoltan Wolf. 2018. "Automation, Labor Share, and Productivity: Plant-Level Evidence from U.S. Manufacturing." Center for Economic Studies Discussion Paper No. 18-39.

Dunne, Timothy, John Haltiwanger, and Kenneth R. Troske. 1997. "Technology and Jobs: Secular Changes and Cyclical Dynamics." Carnegie-Rochester Conference Series on Public Policy, 1997, 46 (1): 107–78.

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

Eslava, Marcela, and John Haltiwanger. 2021. "The Size and Life-cycle Growth of Plants: The Role of Productivity, Demand and Wedges." NBER Working Papers No. 27184.

Fairman, Kristin, Lucia Foster, C.J. Krizan, and Ian Rucker. 2008. "An Analysis of Key Differences in Micro Data: Results from the Business List Comparison Project." Center for Economic Studies Discussion Paper No. 08-28.

Foster, Lucia, John Haltiwanger, and Chad Syverson. 2008. "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" *American Economic Review*, 98(1): 394–425.

Foster, Lucia, John Haltiwanger, and Cody Tuttle. 2022. "Rising Markups or Changing Technology?" Unpublished.

Gollop, Frank M., Barbara M. Fraumeni, and Dale W. Jorgenson. 1987. *Productivity and U.S. Economic Growth*. Harvard University Press.

Haltiwanger, John C., Henry R. Hyatt, Erika McEntarfer, Liliana Sousa, and Stephen Tibbets. 2014. "Firm Age and Size in the Longitudinal Employer-Household Dynamics Data." Center for Economic Studies Discussion Paper No. 14-16.

Handel, Michael. 2016. "The O*NET Content Model: Strengths and Limitations." *Journal of Labor Market Research* 49:157–76.

Oberfield, Ezra and Devesh Raval. 2021. "Micro Data and Macro Technology." *Econometrica* 89 (2): 703–32.

Raval, Devesh. 2019. "The Micro Elasticity of Substitution and Non-Neutral Technology." *RAND Journal of Economics* 50(1): 147–67.

Schreyer, Paul. 2001. "OECD Productivity Manual: A Guide to the Measurement of Industry-Level and Aggregate Productivity Growth." Paris, Organisation for Economic Co-Operation and Development.

Standard Occupational Classification Policy Committee. 2010. "Attachment C: Detailed 2010 SOC Occupations Included in STEM." https://www.bls.gov/soc/2010/home.htm.

Syverson, Chad. 2011. "What Determines Productivity?" *Journal of Economic Literature* 49 (2): 326–65.

Wolf, Michael and Dalton Terrell. 2016. "The High-Tech Industry, What Is It and Why It Matters to Our Economic Future." *Beyond the Numbers: Employment and Unemployment* 5(8), (U.S. Bureau of Labor Statistics, May 2016).

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.

Zolas, Nikolas, Zachary Kroff, Erik Brynjolfsson, Kristina McElheran, David N. Beede, Cathy Buffington, Nathan Goldschlag, Lucia Foster, and Emin Dinlersoz. 2020. "Advanced Technologies Adoption and Use by U.S. Firms: Evidence from the Annual Business Survey." NBER Working Paper No. 28290.

Figure 1. TSB Task/Skill Intensity and Dissimilarity Indexes in Selected Industries (2017)

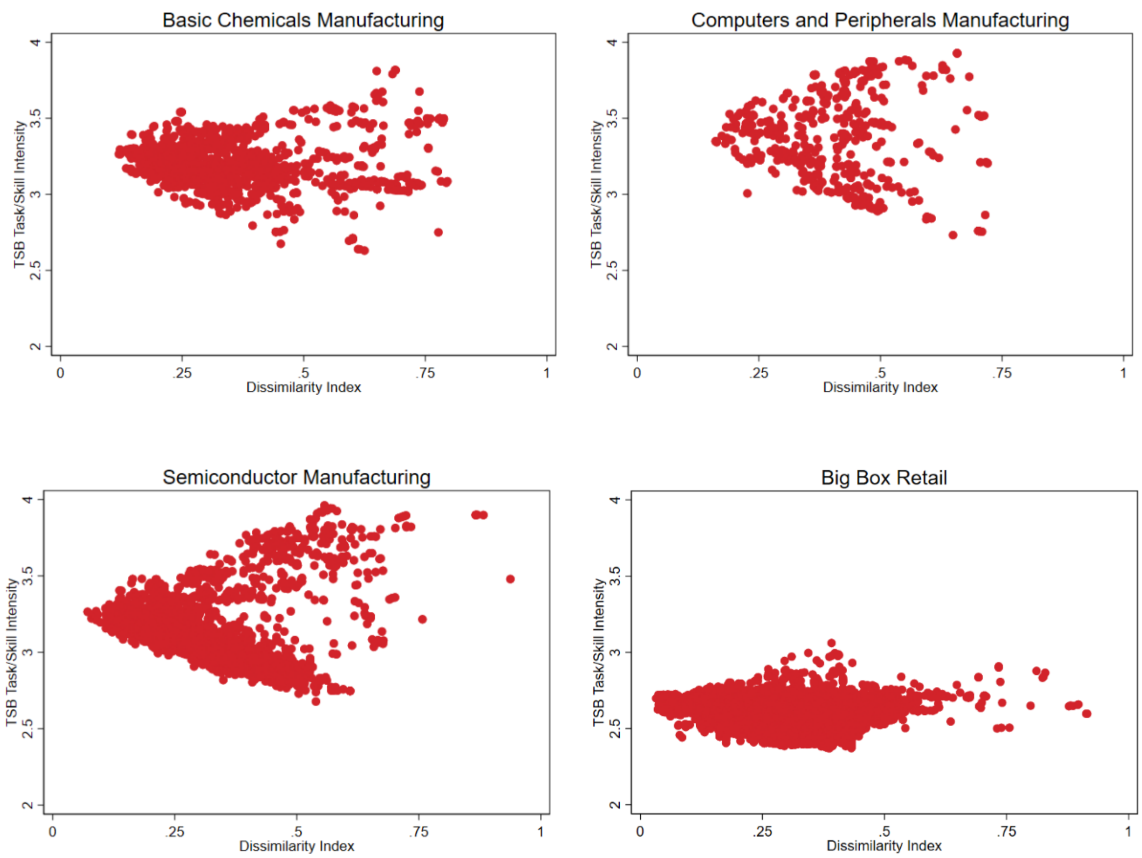

Note: Each dot represents an establishment's indexes in 2017. We exclude establishments with fewer than 20 employees.

Source: Authors' calculations based on the OEWS.

Table 1. Average Industry Correlations between Establishment-Level TSB Task/Skill Intensity and Other Task/Skill Measures

		Manufacturing Industries			
	All Industries	**All**	**High-tech**	**Non-tech**	**Non-manufacturing Industries**
TSU task/skill intensity	0.719	0.773	0.911	0.741	0.696
Non-routine cognitive: Analytical	0.716	0.752	0.904	0.717	0.700
Non-routine cognitive: Interpersonal	0.509	0.578	0.567	0.580	0.479
Routine cognitive	-0.236	-0.359	-0.571	-0.311	-0.182
Routine manual	-0.426	-0.647	-0.832	-0.605	-0.329
Non-routine manual physical	-0.372	-0.569	-0.828	-0.509	-0.286
%STEM workers	0.433	0.610	0.822	0.561	0.355

Note: Each cell represents the establishment-employment-weighted mean four-digit NAICS industry-wide Pearson correlations of establishment skill/task measures pooling the data across 2000, 2005, 2008, 2011, 2014, and 2017.

Source: Authors' calculations based on the OEWS and the O*NET.

Table 2. Summary Statistics of Industry Dispersion Measures in the Manufacturing Sector

	All		**High-tech**		**Non-tech**	
	Mean	**Std. Dev.**	**Mean**	**Std. Dev.**	**Mean**	**Std. Dev.**
Panel A. IQR Dispersion Measures						
TSB task/skill intensity	0.239	0.088	0.350	0.089	0.205	0.054
TSU task/skill intensity	0.151	0.052	0.194	0.051	0.138	0.045
Non-routine cognitive: Analytical	0.511	0.167	0.654	0.162	0.468	0.143
Non-routine cognitive: Interpersonal	0.551	0.120	0.524	0.082	0.560	0.128
Routine cognitive	0.560	0.165	0.564	0.111	0.558	0.178
Routine manual	0.909	0.220	1.063	0.139	0.862	0.218
Non-routine manual physical	0.630	0.166	0.770	0.105	0.587	0.157
%STEM workers	0.082	0.133	0.282	0.144	0.021	0.031
Labor productivity	0.838	0.328	0.967	0.435	0.798	0.277
Total factor productivity	0.530	0.221	0.672	0.301	0.487	0.169
Panel B. 90–10 Dispersion Measures						
TSB task/skill intensity	0.503	0.167	0.717	0.161	0.437	0.102
TSU task/skill intensity	0.302	0.089	0.390	0.086	0.276	0.071
Non-routine cognitive: Analytical	1.032	0.282	1.310	0.274	0.948	0.225
Non-routine cognitive: Interpersonal	1.115	0.224	1.095	0.166	1.121	0.239
Routine cognitive	1.175	0.289	1.276	0.219	1.145	0.300
Routine manual	1.870	0.383	2.092	0.171	1.803	0.404
Non-routine manual physical	1.313	0.304	1.591	0.188	1.228	0.282
%STEM workers	0.195	0.253	0.582	0.236	0.077	0.082
Labor productivity	1.689	0.504	1.954	0.683	1.608	0.403
Total factor productivity	1.119	0.449	1.414	0.648	1.030	0.319

Notes: Industry-employment-weighted means and standard deviations of industry-level dispersion measures for all 86 four-digit NAICS industries in the manufacturing sector pooling the data across 2000, 2005, 2008, 2011, 2014, and 2017.

Source: Authors' calculations based on the OEWS, the O*NET, and DiSP.

Figure 2. Cross-industry Mean IQRs by Time and Industry Type

Panel A. High-tech Manufacturing

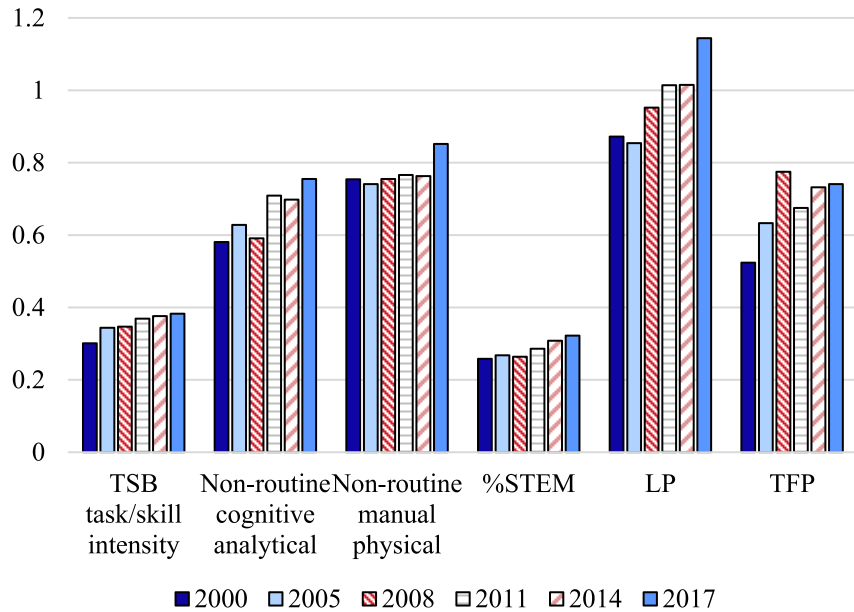

Panel B. Non-tech Manufacturing

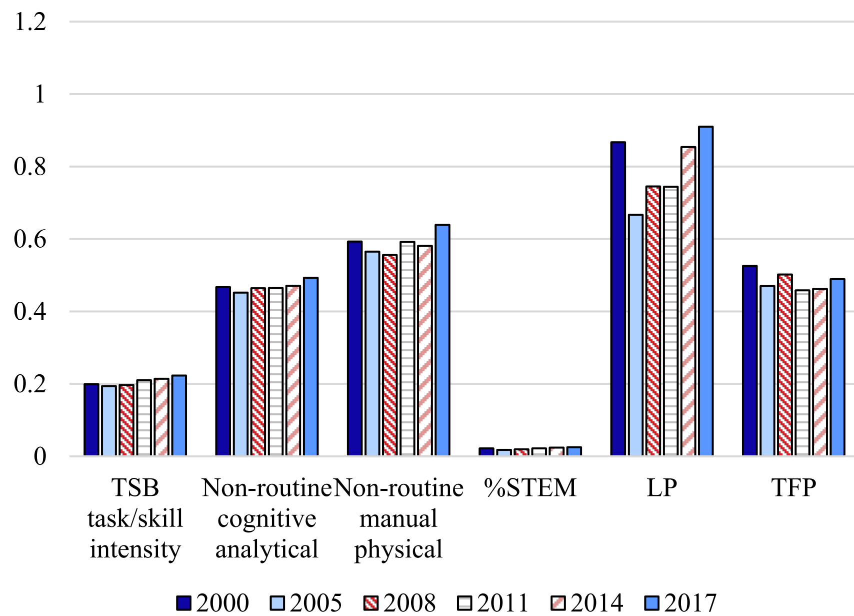

Notes: Industry-employment-weighted means of industry-level dispersion measures.

Source: Authors' calculations based on the OEWS, the O*NET, and DiSP.

Figure 3. Cross-industry Standard Deviation in IQRs by Time and Industry Type

Panel A. High-tech Manufacturing

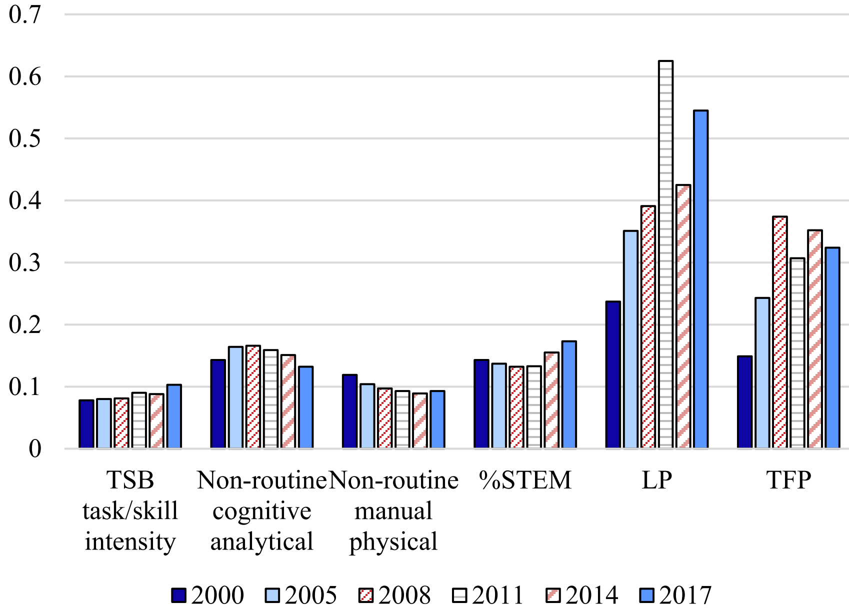

Panel B. Non-tech Manufacturing

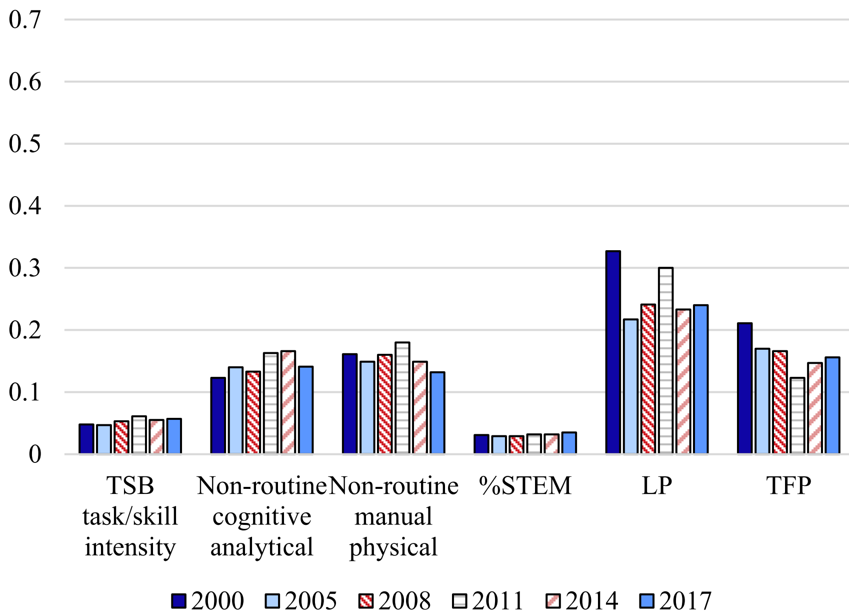

Notes: Standard deviations of IQRs of industry-level dispersion measures.

Source: Authors' calculations based on the OEWS, the O*NET, and DiSP.

Table 3. Correlations between Within-Industry Dispersion in TSB Task/Skill Intensity and Dispersion in Task Measures

		Manufacturing Industries			
	All Industries	**All**	**High-tech**	**Non-tech**	**Non-manufacturing Industries**
Panel A. IQR dispersion					
TSU task/skill intensity	0.642	0.748	0.855	0.587	0.636
Non-routine cognitive: Analytical	0.639	0.745	0.865	0.557	0.632
Non-routine cognitive: Interpersonal	0.596	0.057	0.208	0.226	0.608
Routine cognitive	0.294	0.071	-0.011	0.124	0.293
Routine manual	0.253	0.618	0.381	0.630	0.283
Non-routine manual physical	0.331	0.569	0.549	0.371	0.333
%STEM workers	0.280	0.815	0.841	0.245	0.255
Panel B. 90–10 dispersion					
TSU task/skill intensity	0.662	0.798	0.892	0.594	0.655
Non-routine cognitive: Analytical	0.664	0.799	0.874	0.606	0.657
Non-routine cognitive: Interpersonal	0.559	0.146	0.205	0.296	0.571
Routine cognitive	0.341	0.348	0.187	0.371	0.337
Routine manual	0.248	0.574	0.146	0.677	0.294
Non-routine manual physical	0.325	0.596	-0.0001	0.551	0.326
%STEM workers	0.536	0.832	0.856	0.361	0.521

Note: Each cell represents the industry-employment-weighted Pearson correlation between the four-digit NAICS industry-level TSB dispersion measure and the other skill/task dispersion measures correlations pooling the data across 2000, 2005, 2008, 2011, 2014, and 2017.

Source: Authors' calculations based on the OEWS and the O*NET.

Table 4. Correlations between Productivity Measures and Within-industry Dispersion Measures of Skill in the Manufacturing Sector

	Labor Productivity			**Total Factor Productivity**		
	All	**High-tech**	**Non-tech**	**All**	**High-tech**	**Non-tech**
Panel A. IQR dispersion						
TSB task/skill intensity	0.520	0.397	0.329	0.454	0.317	0.260
TSU task/skill intensity	0.415	0.314	0.306	0.310	0.229	0.144
Non-routine cognitive: Analytical	0.443	0.320	0.369	0.314	0.184	0.213
Non-routine cognitive: Interpersonal	-0.086	-0.082	-0.027	-0.077	-0.206	0.097
Routine cognitive	0.121	-0.101	0.307	0.091	-0.115	0.355
Routine manual	0.215	-0.026	0.143	0.231	0.249	0.041
Non-routine manual physical	0.275	0.122	0.133	0.243	0.036	0.169
%STEM workers	0.440	0.241	0.012	0.453	0.290	-0.049
Panel B. 90–10 dispersion						
TSB task/skill intensity	0.551	0.448	0.366	0.564	0.518	0.209
TSU task/skill intensity	0.461	0.351	0.304	0.426	0.349	0.156
Non-routine cognitive: Analytical	0.477	0.367	0.331	0.423	0.317	0.188
Non-routine cognitive: Interpersonal	-0.065	-0.067	0.118	-0.049	-0.142	0.057
Routine cognitive	0.272	0.018	0.367	0.231	0.063	0.326
Routine manual	0.190	-0.028	0.119	0.156	-0.057	0.043
Non-routine manual physical	0.259	-0.018	-0.053	0.195	-0.262	0.033
%STEM workers	0.461	0.233	0.218	0.555	0.422	0.106

Note: Each cell represents the industry-employment-weighted Pearson correlation between within-industry skill dispersion measures and within-industry productivity dispersion pooling the data across 2000, 2005, 2008, 2011, 2014, and 2017. Industries include all 86 four-digit NAICS manufacturing industries.

Source: Authors' calculations based on the OEWS, the O*NET, and DiSP.

Table 5. Associations between Within-industry Dispersion Measures of Productivity and Skill, Manufacturing Sector (IQR Dispersion Measures)

	Labor Productivity			**Total Factor Productivity**		
	All	**High-tech**	**Non-tech**	**All**	**High-tech**	**Non-tech**
TSB task/skill intensity	2.395***	2.816***	1.366***	1.980***	2.300***	0.769***
	(0.421)	(0.826)	(0.288)	(0.454)	(0.718)	(0.209)
	[0.272]	[0.169]	[0.110]	[0.209]	[0.105]	[0.073]
TSU task/skill intensity	3.246***	3.866***	1.534***	2.295***	2.817**	0.493***
	(0.580)	(1.272)	(0.348)	(0.613)	(1.161)	(0.183)
	[0.174]	[0.111]	[0.099]	[0.098]	[0.059]	[0.023]
Non-routine cognitive:	1.072***	1.242***	0.576***	0.720***	0.688**	0.227***
Analytical	(0.176)	(0.386)	(0.108)	(0.182)	(0.342)	(0.065)
	[0.200]	[0.119]	[0.141]	[0.102]	[0.046]	[0.048]
Non-routine cognitive:	-0.444**	-1.375	-0.111	-0.301	-2.325*	0.136**
Interpersonal	(0.213)	(1.058)	(0.123)	(0.219)	(1.358)	(0.067)
	[0.027]	[0.063]	[0.013]	[0.013]	[0.087]	[0.014]
Routine cognitive	0.296**	-0.609	0.388***	0.209*	-0.674	0.306***
	(0.134)	(0.755)	(0.099)	(0.127)	(0.906)	(0.049)
	[0.026]	[0.049]	[0.102]	[0.013]	[0.034]	[0.130]
Routine manual	0.390***	-0.506	0.143**	0.409***	1.184**	0.032
	(0.112)	(0.445)	(0.073)	(0.125)	(0.496)	(0.043)
	[0.055]	[0.049]	[0.028]	[0.059]	[0.074]	[0.004]
Non-routine manual physical	0.667***	0.518	0.177*	0.573***	0.096	0.171**
	(0.143)	(0.545)	(0.103)	(0.138)	(0.516)	(0.072)
	[0.082]	[0.045]	[0.024]	[0.065]	[0.023]	[0.032]
%STEM	1.339***	0.991*	0.030	1.299***	1.321***	-0.235
	(0.305)	(0.512)	(0.349)	(0.304)	(0.333)	(0.326)
	[0.201]	[0.085]	[0.009]	[0.209]	[0.100]	[0.004]

Note: Each cell contains coefficient estimates from a linear regression of within-industry productivity dispersion on within-industry skill dispersion including a constant and year effects. Industry employment weights are used. Standard errors are in parentheses and R-squared values are in brackets. All regressions are estimated using OLS by pooling data on all 86 four-digit NAICS manufacturing industries across years (2000, 2005, 2008, 2011, 2014, and 2017).

Source: Authors' calculations based on the OEWS, the O*NET, and DiSP.

Table 6. Associations between Within-industry Dispersion Measures of Productivity and Skill, Manufacturing Sector (90–10 Dispersion Measures)

	Labor Productivity			**Total Factor Productivity**		
	All	**High-tech**	**Non-tech**	**All**	**High-tech**	**Non-tech**
TSB task/skill intensity	1.975***	2.609***	1.388***	1.599***	2.558***	0.507***
	(.306)	(0.669)	(0.226)	(0.275)	(0.553)	(0.159)
	[.305]	[0.206]	[0.138]	[0.329]	[0.279]	[0.060]
TSU task/skill intensity	3.166***	3.917***	1.626***	2.326***	3.307***	0.532***
	(0.529)	(1.201)	(0.310)	(0.450)	(0.838)	(0.182)
	[0.214]	[0.127]	[0.093]	[0.191]	[0.133]	[0.035]
Non-routine cognitive: Analytical	1.014***	1.260***	0.565***	0.717***	0.929***	0.200***
	(0.162)	(0.363)	(0.093)	(0.135)	(0.245)	(0.057)
	[0.231]	[0.139]	[0.115]	[0.191]	[0.115]	[0.048]
Non-routine cognitive: Interpersonal	-0.253*	-0.783	-0.130	-0.094	-0.952	0.058
	(0.151)	(0.828)	(0.097)	(0.137)	(0.905)	(0.060)
	[0.014]	[0.025]	[0.014]	[0.007]	[0.048]	[0.011]
Routine cognitive	0.561***	0.090	0.469***	0.375***	0.216	0.246***
	(0.103)	(0.491)	(0.074)	(0.087)	(0.433)	(0.043)
	[0.079]	[0.011]	[0.141]	[0.060]	[0.021]	[0.118]
Routine manual	0.298***	-0.392	0.114*	0.204***	-0.508	0.033
	(0.074)	(0.529)	(0.059)	(0.052)	(0.581)	(0.028)
	[0.041]	[0.015]	[0.022]	[0.033]	[0.028]	[0.011]
Non-routine manual physical	0.521***	-0.150	0.163*	0.323***	-1.103**	0.061
	(0.111)	(0.621)	(0.092)	(0.073)	(0.486)	(0.053)
	[0.074]	[0.012]	[0.021]	[0.049]	[0.091]	[0.013]
%STEM	1.083***	0.853**	0.989***	1.027***	1.378***	0.303*
	(0.218)	(0.422)	(0.352)	(0.171)	(0.248)	(0.185)
	[0.214]	[0.059]	[0.053]	[0.315]	[0.192]	[0.020]

Note: Each cell contains coefficient estimates from a linear regression of within-industry productivity dispersion on within-industry skill dispersion including a constant and year effects. Industry employment weights are used. Standard errors are in parentheses and R-squared values are in brackets. All regressions are estimated using OLS by pooling data on all 86 four-digit NAICS manufacturing industries across years (2000, 2005, 2008, 2011, 2014, and 2017).

Source: Authors' calculations based on the OEWS, the O*NET, and DiSP.

Figure 4. Relationships Between TFP Dispersion and Task/Skill Dispersion Measures

Panel A. TSB Task/Skill Intensity Dispersion

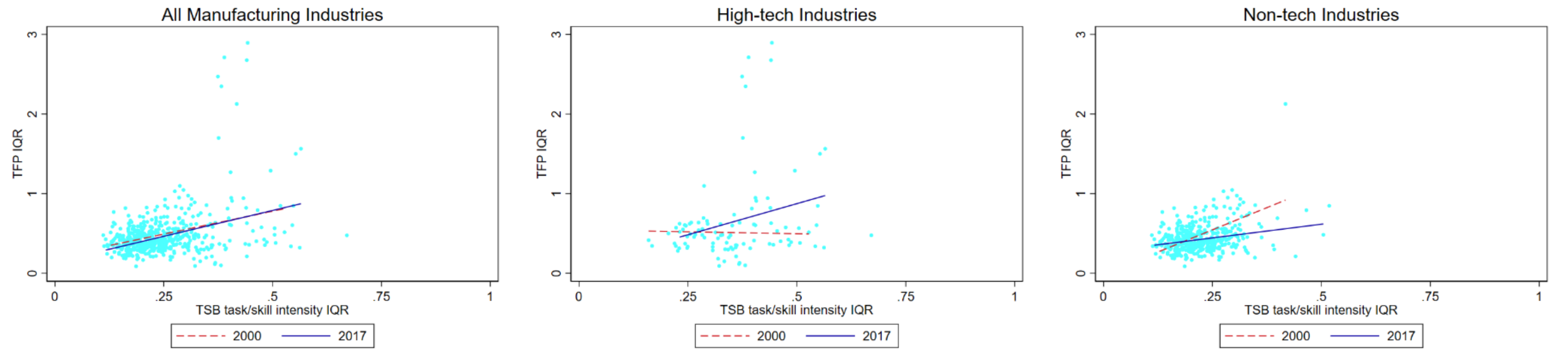

Panel B. Non-routine Cognitive: Analytical Dispersion

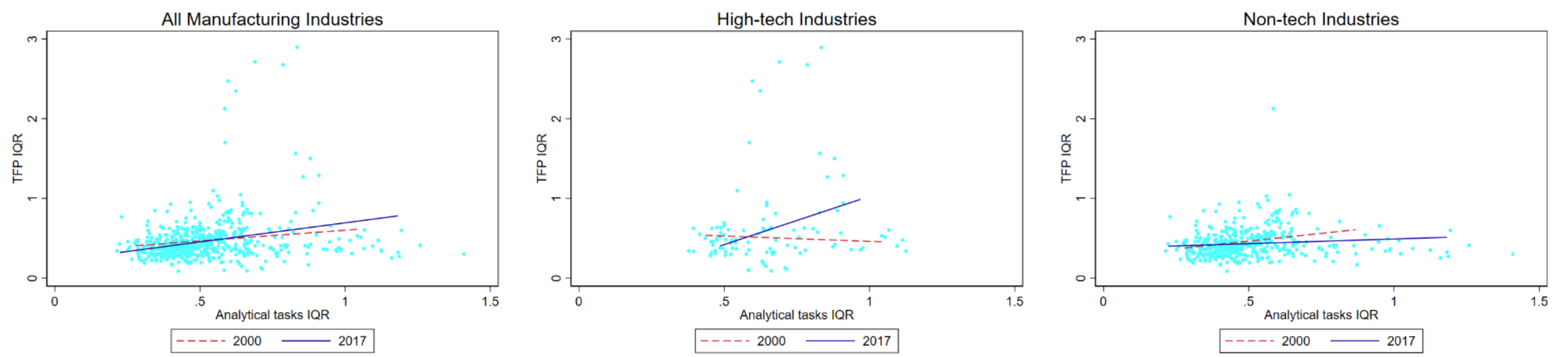

Panel C. Non-routine Manual Physical Dispersion

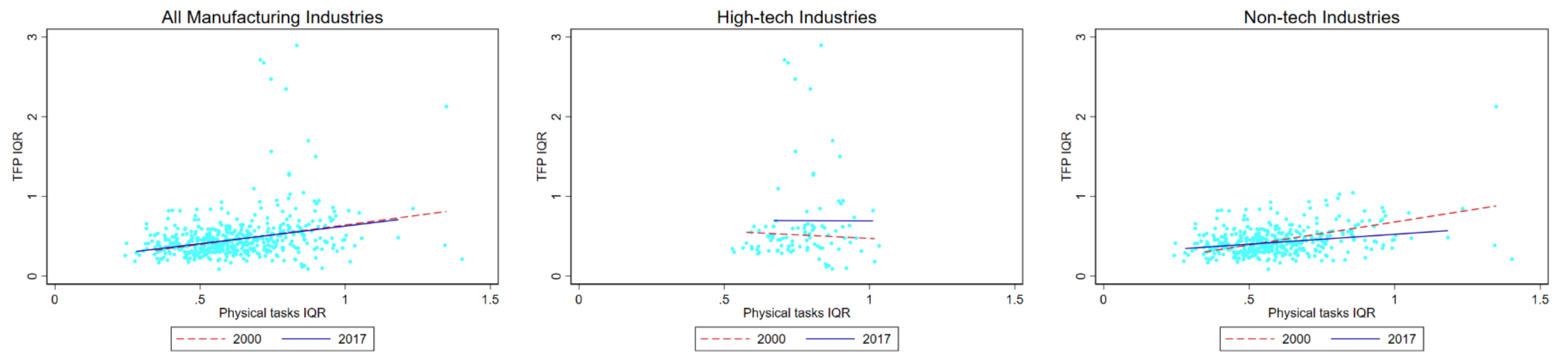

Panel D. % STEM Workers

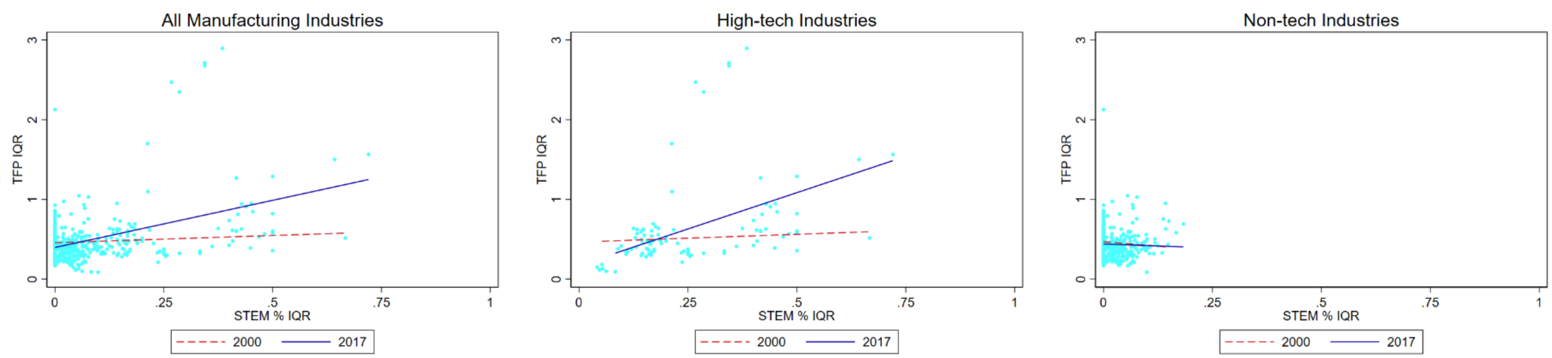

Note: Pooled interquartile range dispersion measures for four-digit NAICS manufacturing industries from 2000, 2005, 2008, 2011, 2014 and 2017 are depicted. Slopes for 2000 and 2017 are also included.

Source: Authors' calculations based on the OEWS, the O*NET, and DiSP.

Figure 5. Distributions of between EIN Share of Total Variation, All Units Imputed or Observed, All EINS

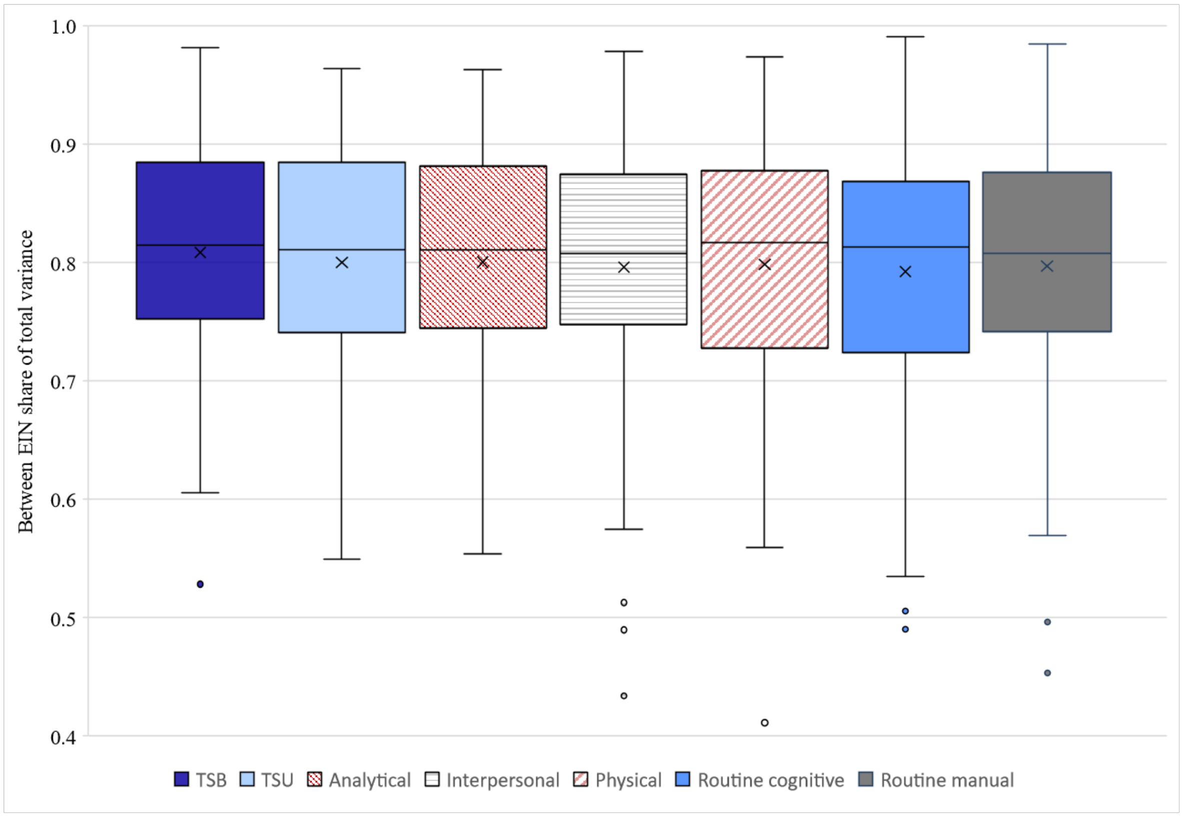

Notes: The boxplots represent the distribution of the between EIN share of total variation for all four-digit NAICS manufacturing industries for each skill/task measure. The line in the middle reports the share for the median industry, while the x reports the mean share. The boxes represent the interquartile range, which bound the industries that lie between the 25th and 75th percentiles, respectively. The upper and lower whiskers span the lowest and highest quartiles within 1.5 IQR of the nearest quartile. The dots represent outliers.

Source: Authors' calculations based on the OEWS and the O*NET.

Figure 6. Distributions of between EIN Share of Total TSB Measure Variation

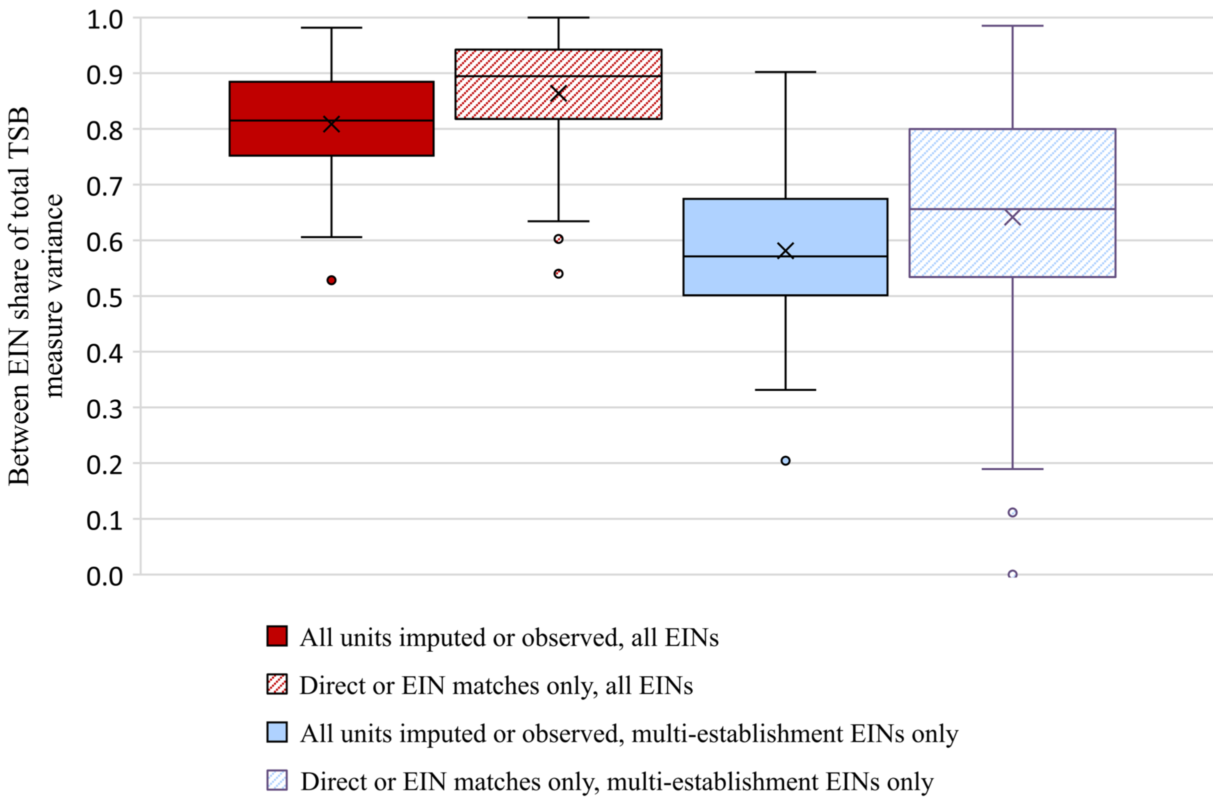

Notes: The boxplots represent the distribution of the between EIN share of total variation for all four-digit NAICS manufacturing industries for each skill/task measure. The line in the middle reports the share for the median industry, while the x reports the mean share. The boxes represent the interquartile range, which bound the industries that lie between the 25th and 75th percentiles, respectively. The upper and lower whiskers span the lowest and highest quartiles within 1.5 IQR of the nearest quartile. The dots represent outliers.

Source: Authors' calculations based on the OEWS.

Figure 7. Between EIN Share of Total TSB Measure Variation, Direct or EIN Matches Only vs. All Units Either Imputed or Observed, All EINs

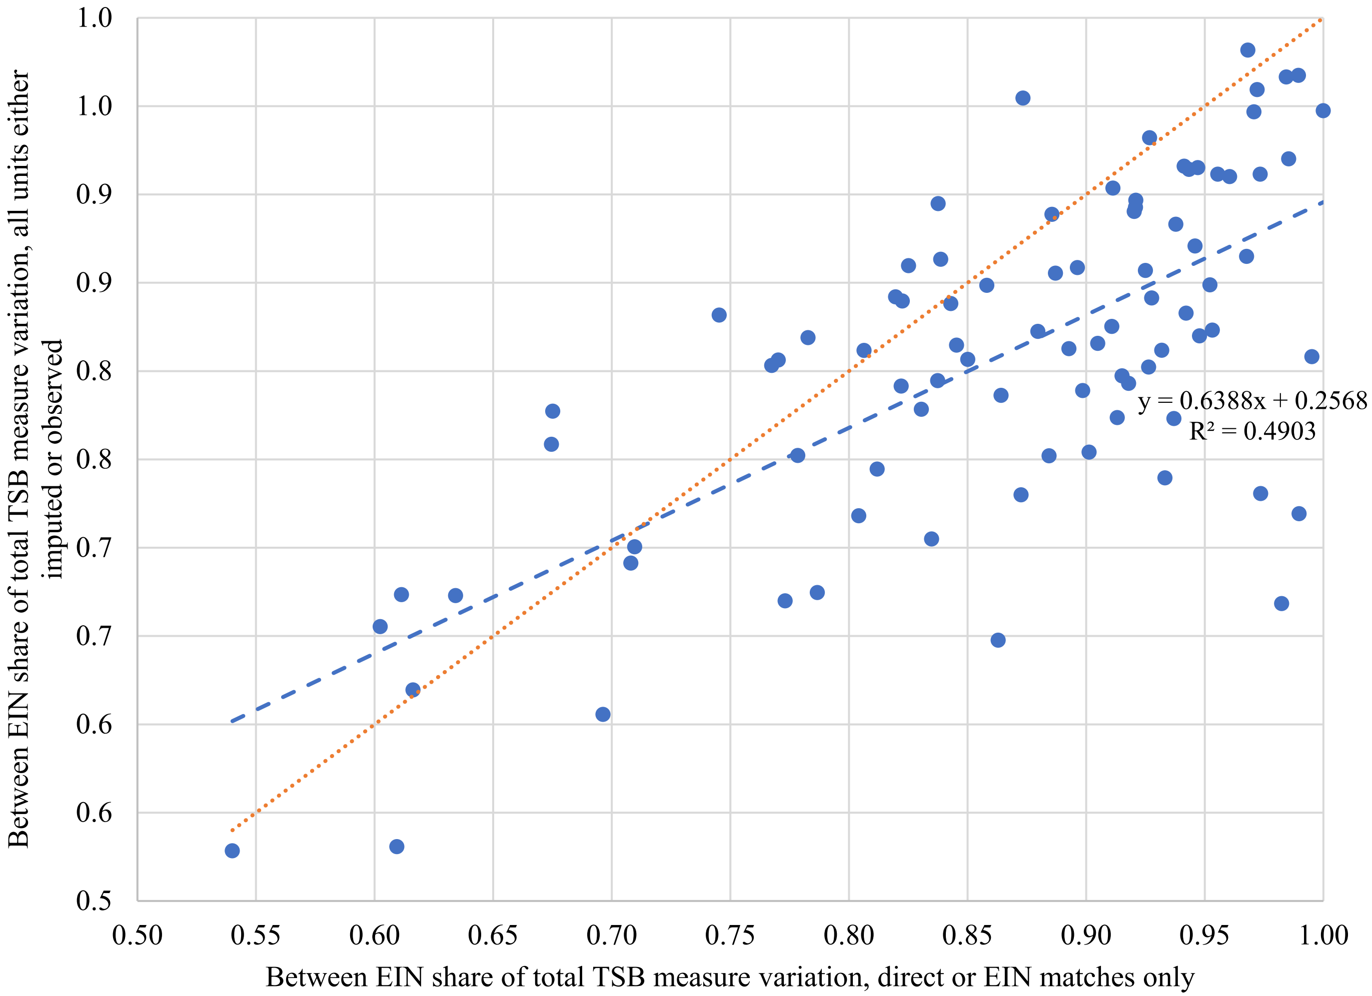

Source: Authors' calculations based on the OEWS.

Figure 8. Between EIN Share of Total TSB Measure Variation, Direct or EIN Matches Only vs. All Units Either Imputed or Observed, Multi-establishment EINs Only

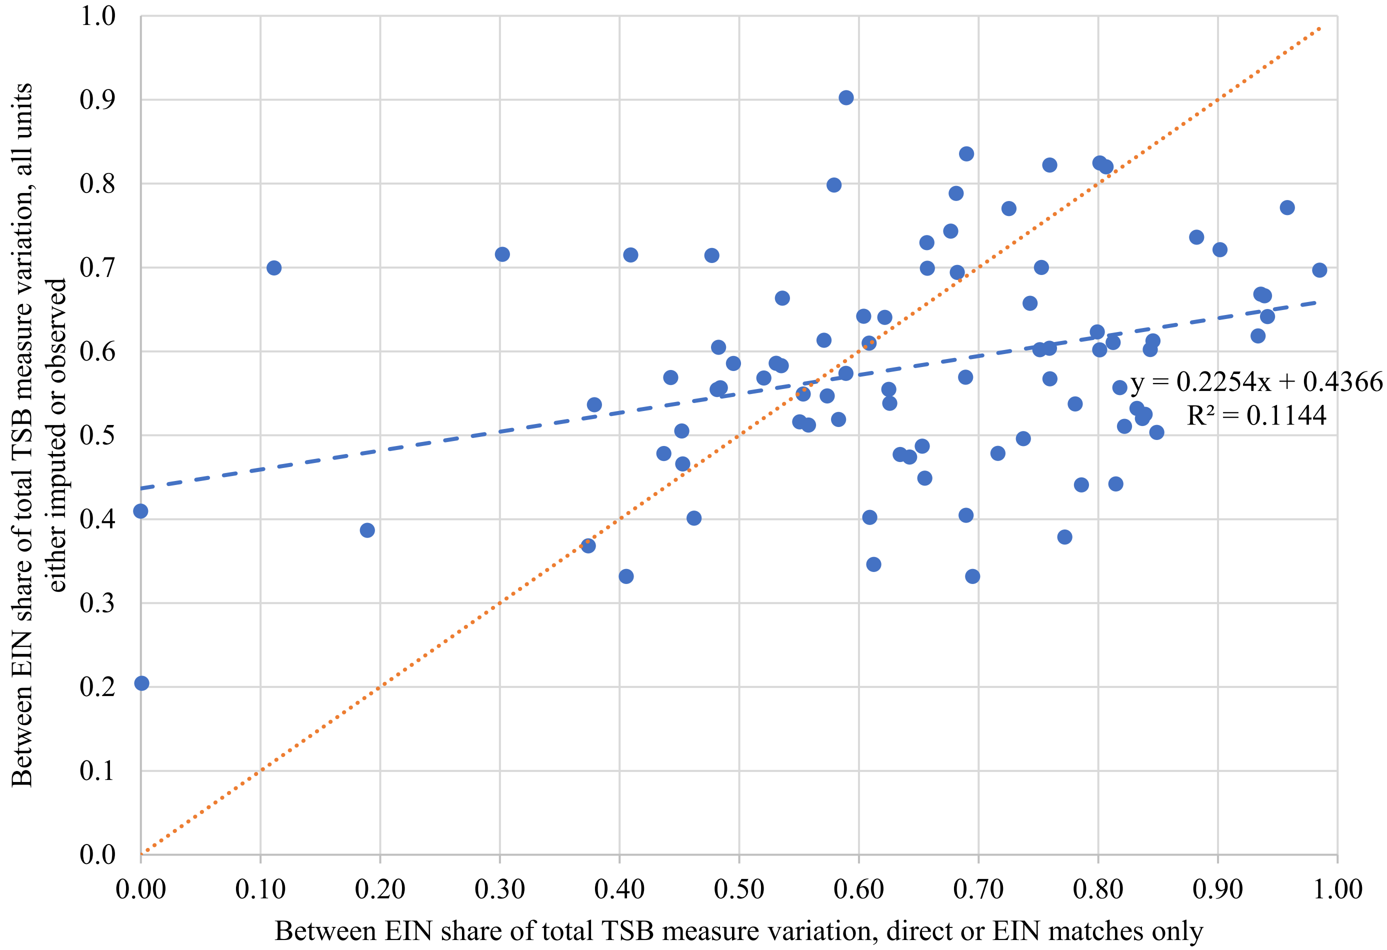

Source: Authors' calculations based on the OEWS.