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*

February 15, 2022

A major gap in studies that examine establishment-level productivity empirically is that the datasets used typically have limited information about the characteristics of establishments' workers or the tasks they perform. 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 with a restricted-access establishment-level productivity dataset created as part of the joint BLS-Census Bureau Collaborative Micro-productivity Project. This paper takes a first step toward this objective by exploring the conceptual, specification, and measurement issues that need to be confronted. In addition, we include suggestive empirical analysis of the relationship between dispersion in within-industry productivity measures and dispersion in within-industry task and skill measures. 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 necessarily 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 and John Eltinge 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. (2021a) found that establishments at the 90th percentile of the total factor productivity (TFP) distribution are about 2.9 times as productive as establishments at the 10th percentile within four-digit NAICS manufacturing industries. In a survey article, Syverson (2011) discusses several possible sources for this dispersion in productivity, 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 a restricted-access dataset and publicly available dispersion statistics (Dispersion Statistics on Productivity or DiSP) posted on the BLS and Census Bureau websites.[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 that the establishment employs and the tasks they perform. Our ultimate goal is to integrate establishment-level data from the Occupation Establishment and Wage Survey (OEWS) with Census Bureau establishment-level business data.

[1] DiSP is available at: https://www.bls.gov/lpc/productivity-dispersion.htm and https://www.census.gov/disp. See Cunningham et al. (2021a) for a detailed description of the development of the datasets.

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

The productivity dispersion statistics come from the DiSP data. We develop several measures of task/skill dispersion from the OEWS data. One measure we explore is a composite index of task/skill intensity that accounts for the pricing of occupations in the labor market. It is related to, but distinct from, the skill-adjusted labor input measure that BLS publishes as part of its official TFP measures.[2] Our primary task/skill intensity index is best thought of as a counterfactual average establishment wage—it is 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, because the pricing of tasks is bundled through the occupations.

We also consider a second task/skill intensity index that is a similar counterfactual wage but based on predicted wages from a linear regression of OEWS national occupational wages on five task indexes—non-routine cognitive (analytical), non-routine cognitive (interpersonal), routine cognitive, routine manual, and non-routine manual physical—constructed from work activities and work context importance scales in the Occupational Information Network (O*NET) as described in Acemoglu and Autor (2011, p. 1163). This counterfactual wage is the

[2] See Changes in the Composition of Labor - BLS Multifactor Productivity Measures 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).

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

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

Beyond these task/skill indexes, we also consider 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. We examine the establishment values of the five O*NET task indexes individually. Finally, we look at the percentage of STEM workers in each establishment, based on the occupational mix.

We compare the within-industry labor productivity (LP) and TFP dispersion measures from DiSP to within-industry dispersion in these OEWS-O*NET task measures. To preview our results, we find that the bundled and unbundled task/skill intensities are highly correlated with each other at the establishment-level within industries. Again, using establishment-level variation within industries, we find that the bundled task/skill index is positively correlated with the analytical task index, the interpersonal task index, and the STEM intensity index but negatively correlated with the non-routine manual physical task index, the routine manual task index, and the routine cognitive task index. These establishment-level correlation patterns suggest that within-industry dispersion of task/skill intensity measures will be positively correlated with the other indexes. This pattern holds even if the establishment-level correlations are negative, indicating that there is a general tendency for dispersion on one dimension of task/skill intensity to be associated with dispersion on other dimensions.

Turning to the relationship between within-industry dispersion in task/skill indexes and productivity dispersion, we find that higher dispersion in within-industry productivity is associated with higher within-industry dispersion of the bundled and unbundled task/skill intensity indexes, the analytical task index, and the STEM intensity index. These patterns differ quantitatively across different groupings of industries and are especially strong among the high-tech industries of the manufacturing sector. For example, the elasticity of the dispersion of within-industry TFP with respect to the bundled task/skill intensity is about three times larger among the high-tech manufacturing industries than among 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 highlights that both the outcomes and the ways of doing business differ, especially among the most innovative sectors of the economy. The high-tech sectors are already well-known to exhibit high productivity growth and have a higher intensity of workers in STEM occupations (see, e.g., Decker et al. 2020). Our findings imply the strong relationship in first moments carries over to the related second moments.

While our results are suggestive, they are promising for the longer-run objective of integrating the OEWS data with Census Bureau business microdata at the establishment-level. It will be in this establishment-level analysis that we can explore related issues such as the relationship between technology adoption and the task/skill mix of businesses. This work is left for the future, but we provide an overview of the potential for such analysis with establishment-level data integration.

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 first the productivity data and its public-data product and then the OEWS data in detail along with the task/skill intensity indexes we construct. Section 4 presents the results relating within-industry productivity and task/skill dispersion. Section 5 provides an overview of next steps. Concluding remarks are provided 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} denotes a Hicks-neutral productivity term and is often interpreted as technical change, and e and t index establishments and time respectively, the way businesses use labor differently may show up as differences in L, $\frac{\partial F}{\partial L}$, or both. Our objective is to take a first step towards exploring the potential heterogeneity in the labor input and its effect on productivity dispersion. An establishment may have higher measured productivity than its competitors due to more efficient use of its labor force because it relies on certain types of labor more heavily than others, 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(.) becomes $Z_{et}\cdot L_{et}$, which accounts for differences across establishments that are 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 (see equation (2) in Acemoglu and Autor (2011)) for understanding skill premia is based on two types of workers, low-skilled and high-skilled:

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

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, respectively. 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 labor 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 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

[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/mfp/mprtech.htm).

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, see Acemoglu (1998), postulates that technological advances may complement or fully replace either labor type:

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

where B_l and B_h denote skill-replacing technologies that are perfect substitutes for the two labor types.[4] The terms that include L and H are aggregated using a harmonic sum with appropriate relative weights based on α, which reflect the distribution of tasks between different types of labor, see 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 were empirically successful in studying the supply and demand for skills, they were less successful in explaining aggregate moments 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 proposed and explored alternative models where the production process amounts to combining different production tasks with the appropriate skills (e.g., Autor, Levy, and Murnane 2003; Acemoglu and Restrepo 2018a, 2019).

A common property of these more recent approaches is to distinguish between tasks and skills, thus expanding the canonical model where the two are the same for analytical tractability. Autor, Levy, and Murnane (2003) show, for example, that information and communication

[4] In the case of skill-complementing technologies, the currently linear terms of the inner brackets in (2) would have the same harmonic sum structure as in (1) with exponents less than zero.
[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.

technologies can substitute for workers that perform routine tasks and at the same time complement workers that perform more complex or non-routine tasks. They also show that 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) conceptualize 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 that 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 circumstance gives rise to the possibility of an endogenous response of technology to market conditions, automation, or endogenous technology choice. See Acemoglu (1998), Acemoglu and Restrepo (2018b, 2018c, 2018d, 2018e, 2019), and Dinlersoz and Wolf (2018).

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

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

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. These tasks can be performed by 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 neutral productivity shock term, $\Pi(I,N)$, is assumed to be a function of innovation and tasks as well. This approach highlights that what is needed are proxies for the number and nature of production tasks as well as the relationship between innovation and tasks. This specification also helps emphasize that there is a potentially endogenous relationship between the choice of innovation, tasks, the factor content of tasks, and 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, ideally, it would be desirable to control for these differences when calculating establishment-level differences in productivity. The first step in this direction is to compare information on occupations and wages from OEWS and O*NET and then explore the relationship between within-industry dispersion in task/skill intensity and within-industry dispersion in productivity, which we do in this paper.

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 (CMP data hereafter) for the manufacturing sector from 1972 to2018.[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 5-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 years, in

[7] CMP stands for Collaborative Micro-productivity Project. These data are available for use by qualified researchers on approved projects in the Federal Statistical Research Data Centers (https://www.census.gov/fsrdc).
[8] ASM panels start in years ending in "4" and "9".

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., 2021a). 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]

3.3. Basic 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 Autor and Acemoglu (2011). *Tasks* are activities that when combined with intermediate goods create a good or service and are the true factors of production that we would like to measure. However, because we do not observe time

[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. See Cunningham et al. (2021a) for more details on the construction of the DiSP.

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" (Autor and Acemoglu 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 Autor and Acemoglu (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. A challenge is that the relationship between skills and tasks can vary over time and across businesses.

3.4. Occupational Employment and Wage Statistics (OEWS)

The occupation data come from the BLS's Occupational Employment and Wage Statistics (OEWS) survey, which is a semi-annual mail survey that samples 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

[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. 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.

each of 12 wage intervals.[11] The OEWS survey uses the Office of Management and Budget's (OMB) 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 most other surveys that include information about occupation.

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.

The OEWS sampling and weighting methods guarantee that 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 that 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 that attempts 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 that was created using a modified version of the imputation approach developed by Dey et al. (2019).

Dey et al. (2019) impute data for the entire QCEW. For each reference year, they use the same dating convention as for the official OEWS release (that is, May of the reference year

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

combined with the five previous panels). For each observation in the QCEW, they identify 5–10 donor observations based on the characteristics of the establishments. The characteristics include employment, industry (6-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 that 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 et al. (2019) approach is designed as a potential replacement for the current methodology for generating official estimates. The main advantage of this approach is that every establishment in the QCEW is represented and has an establishment weight of one. The 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 methodology for our analysis, which focuses on distributions.

The primary modification to this methodology 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 (6-digit NAICS), ownership (private or type of

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

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, where 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.

Second, for purposes of integrating with other data, we 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). 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 that the occupational mix within an establishment is fixed over the three-year interval.

The result is a full data set 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, to be consistent with the DiSP.

[14] The 5-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.

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

Our first index of task/skill intensity is a counterfactual wage that is the average wage paid by the establishment if the establishment paid the national average occupational wage for each worker that it employs in that occupation for each year sample. 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*. 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:

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

where *B* is the set of 12 wage intervals and E_n is the set of all establishments (nationwide, manufacturing, and non-manufacturing). The 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*, L_{ej} is the number of employees in occupation *j* for 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*, $\tilde{w}_e$, is equal to:

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

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) that pay 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 that 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 that have occupation mixes that differ 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

[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 that 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.

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 nearly 1.0. 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. 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 use O*NET data to operationalize the Autor, Levy, and Murnane (2003) taxonomy of tasks (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

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

occupations at establishments.[17] O*NET contain over 275 variables that describe each occupation. Acemoglu and Autor (2011) use 16 of these variables that correspond 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 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[20] and estimate a regression of the national occupational mean log wage for each year on these five task indexes as follows:

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

[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 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 that 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 v. unstructured work (reverse). Routine manual includes pace determined by speed of equipment, controlling machines and processes, and spend time making repetitive motions. Non-routine manual physical includes operating vehicles, mechanized devices, or equipment; spend time using hands to handle, control, or feel objects, tools, or controls; manual dexterity; and spatial orientation. (page 1163 of Acemoglu and Autor (2011))

[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.

where x_{jk} is the O*NET measure of task k for occupation j. The coefficients on the task indexes, β_k, are akin to prices in a hedonic regression. We then predict average establishment wages for each year from this regression, $\widehat{w}_e$, conditional on the occupational distribution in the establishment, as our second measure of skill intensity. It is worth noting that this unbundled task/skill intensity index has the same property as the bundled index, in that there are many possible combinations of tasks that can result in the same value of the index.

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 reflect a bundle of tasks. We discuss these differences 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 task measures based on the average value of the individual O*NET task indexes. For each of the five task indexes, we measure an employment-weighted establishment-level average for task index k as follows:

$$\tau_{ek} = \frac{1}{L_e} \sum_{j \in J_e} \tau_{jk} \times L_{ej} \tag{8}$$

where k = 1, …, 5. Thus, τ_{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.

3.7. STEM Intensity Index

As another alternative task index, we calculate the percentage of STEM workers in each establishment as follows:

$$\%STEM_e = \frac{1}{L_e} \sum_{j \in J_s} L_{ej} \tag{9}$$

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.

3.8. 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 that 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. That is, the TSU measure can be thought of as the average price of tasks performed by employees in the establishment. 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 correlations are the employment-weighted average 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 (ρ = 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 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 that 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 individual task indicators are negatively correlated with TSB. That is, establishments that have a high composite task/skill intensity generally employ fewer occupations that are 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.

3.9 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

[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.

each sample year. To account for industry differences in average skills/tasks so that 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., 2021a, for details). Given that the DiSP measures are only available for manufacturing industries, the summary statistics in Table 2 are only for manufacturing.

Summary statistics for all our within-industry dispersion measures (productivity and task/skill) using the pooled data are shown in Table 2 for 4-digit manufacturing industries. 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 that 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. Not only do we see higher dispersion in the high-tech manufacturing industries than the non-tech manufacturing industries, but there is also more dispersion in dispersion in high-tech manufacturing industries across all our measures. Figures 2 and 3 highlight that 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 the TSB task/skill measure and the TSU task/skill measure, the analytical task measure, and the interpersonal tasks measure. An industry with above-average TSB task/skill dispersion is likely to exhibit above-average dispersion in each of these task/skill measures. Among non-manufacturing industries (last column), these task/skill measures are similarly highly correlated. Among manufacturing industries (second column), the correlations between the TSB task/skill measure and the TSU task/skill measure, the analytical task measure, routine manual, non-routine manual physical, and %STEM all exceed 0.5. The interpersonal and routine cognitive task dispersions are less likely to be important for TSB task/skill intensity 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 task/skill intensity, analytical tasks, non-routine manual physical tasks, and %STEM varies more closely with dispersion in TSB task/skill intensity in the high-tech group, while dispersion in routine manual tasks and interpersonal tasks varies more with TSB task/skill dispersion in the non-tech group, as could be expected.

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 Skills, Tasks, and Productivity

To analyze the link between productivity and tasks/skills, we consider the relationship between within-industry dispersion measures of productivity, skills and tasks. This 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 innovation, tasks, and factor mixes. In considering such innovation, this might be product-quality-enhancing or process-enhancing innovation.

We first calculate Pearson correlations between our dispersion measures for manufacturing industries (all, high-tech, and non-tech). In Table 4, we observe that 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 measures also have strong positive relationships with the productivity dispersion measures. Looking at high-tech and non-tech manufacturing industry groups separately, we see that 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.

In Tables 5 and 6, we 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 and at the same time 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, at the same time, they are informative about the explanatory power of the variation in skills and task-content for productivity dispersion. Again, we see positive and statistically significant association 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, which is not surprising given that LP dispersion is determined by the variation in output and hours while TFP dispersion is determined by variation in output and all other factors of production.

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 indicates that the TSB dispersion accounts for about a quarter of the variation in LP 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 differences. 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.

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 is (evaluated at means) more than three times larger in the high-tech grouping of industries than in non-tech industries. 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 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 in Cunningham et al. (2021b) who find 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 standard establishment characteristics and look at the basic characteristics of workers and tasks.[22] Our exploratory regressions suggest that as much as 20–25 percent of IQR variation and up to 30 percent of 90–10 variation (see Table 6) in differences in within-industry productivity dispersion across industries and time can be accounted for using variation in the within-industry dispersion in the task/skill content of establishments.

Figure 4 illustrates graphically the relationship 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 that 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

[22] Appropriate caution is needed in making these comparisons, because Cunningham et al. (2021b) investigated the relationship between establishment-level productivity and characteristics while here we are relating within-industry dispersion in productivity and task/skill characteristics.

strongly positive in 2017. On the other hand, for the non-tech industries, the relationship weakened between 2000 and 2017. We find similar patterns using the analytical tasks dispersion. Also, looking at high-tech industries, we find that 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 became weaker among the high-tech industries. Among non-tech industries, the relationship is strongly positive, although it was 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 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 that is 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

The next step in the larger project is to integrate the OEWS and Census Bureau microdata on productivity (CMP). In this section, we discuss the challenges that will need to be overcome in this data integration along with plans for the type of analysis that can be conducted with this data.

One challenge is that the OEWS is drawn from the BLS business register, while the underlying source CMP data are drawn from the Census Business Register. The common identifier on both files is the Employer Identification Number (EIN). Studies have shown (Fairman et al. 2008; Haltiwanger et al. 2014) that there is a high match rate of EINs across the registers. 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. A working (and testable) assumption is whether the occupational mix varies much across establishments within the same EIN. Put differently, one important next step will be to examine the extent of between establishment occupational mix variation within versus between EINs. If, as we expect, most of the variation is between EINs, this will facilitate our integration of the occupational mix information with the CMP data.

When matching OEWS data to the CMP data, we will need to consider that both represent samples. For this analysis, we will restrict the OEWS sample to exact matches (that is, non-imputed observations). We will then 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. (2021a). An alternative is to consider integrating the OEWS data that includes imputations as we have used in the current paper. In considering either option, it is useful to note that the CMP data provide essentially universe data for productivity in Economic Census years. Put differently, for Census years we are less constrained by the difference in the sampling of establishments in the OEWS and the ASM.

We plan on both exploratory data exercises to understand the joint distribution of establishment-level productivity and occupational mix as well as more structured analysis following the conceptual framework discussion in section 2. The 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 given that 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, then it is reasonable to calculate dispersion accounting for such differences either by adjusting the labor input measure or using more than one labor elasticity or both. How such clusters should be identified is an open question, but information on occupations from OEWS could provide insights on 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 various task/skill-based labor inputs. 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 to what extent adjusting for task-content differences in this manner can help account for observed differences in measured productivity that lacks such adjustments. 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. A 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, while deviating from this assumption allows the elasticity to be estimated. Existing studies suggest that such generalizations are likely important in this context (see, e.g., Dinlersoz and Wolf 2018; Raval 2019; Oberfield and Raval 2021). The CES technologies in section 2 are a good starting place given the recent insights of Acemoglu and Restrepo (2019) (e.g., see equation 3 borrowed from their work). We anticipate that considering even more flexible functional forms is of interest. Foster et al. (2021) have found that much of the rising mean and variance of measured markups using the DeLoecker et al. (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 interests here because measured markup dispersion is closely related to measured revenue productivity dispersion shown in 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., DeLoecker et al. 2016; DeLoecker et al. 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, the estimation is of the revenue function. Therefore, the endogeneity of prices including markups needs to be accounted for. While a challenge, this issue enhances the interest in helping us understand the causes and consequences of businesses doing business differently.

A related but distinct set of exercises 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 included modules on the adoption of advanced technologies. Analysis of those modules by Zolas et al. (2020) highlights that 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.[23] Using the CES specifications in section 2 combined with Shepherd's lemma will permit an examination of the relationship between skill and task mix and the adoption of technologies.

6. Conclusion

Our analysis has barely opened the black box of the establishment, but our results suggest that we can gain considerable insights about an important potential source of differences in measured productivity across establishments within an industry and about the mix of tasks and skills being used at establishments. Our evidence is only suggestive, as we find that industries with high dispersion in within-industry productivity also exhibit high dispersion in a number of the task/skill intensity measures that we construct. Interestingly, the relationship between dispersion in productivity in the high-tech industries is especially related to dispersion in the analytical task and STEM intensities.

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

[23]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.

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. 2018c. "Robots and Jobs: Evidence from US Labor Markets." NBER Working Paper No. 23285.

Acemoglu, Daron, and Pascual Restrepo. 2018d. "Modeling Automation." *AEA Papers and Proceedings* 108: 48–53.

Acemoglu, Daron, and Pascual Restrepo. 2018e. "Artificial Intelligence, Automation and Work." NBER Working Paper No. 24196.

Acemoglu, Daron and Pascual Restrepo. 2019. "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives* 33 (2): 3–30.

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics* 116 (4): 1279–333.

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):142–172.

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. 2021a. "Dispersion in Dispersion: Measuring Establishment Level Differences in Productivity." IZA Discussion Paper No. 14459. http://ftp.iza.org/dp14459.pdf.

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

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. https://www.bls.gov/opub/mlr/2019/article/model-based-estimates-for-the-occupational-employment-statistics-program.htm.

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.” *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 Troske. 1997. "Technology and Jobs: Secular Changes and Cyclical Dynamics." Carnegie-Rochester Conference Series on Public Policy, 1997, 46(1): 107–178.

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.

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* 49:157–176.

Oberfield, Ezra and Devesh Raval. 2021. “Micro Data and Macro Technology.” *Econometrica* 89 (2): 703–732.

Raval, Devesh. 2019. “The Micro Elasticity of Substitution and Non-Neutral Technology.” *RAND Journal of Economics* 50(1): 147–167.

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): 326–365.

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), https://www.bls.gov/opub/btn/volume-5/mobile/the-high-tech-industry-what-is-it-and-why-it-matters-to-our-economic-future.htm.

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.