Estimating the Nature of Technological Change: Exploiting Shifts in Skill Use Within and Between Occupations*

Costas Cavounidis,† Vittoria Dicandia,‡ Kevin Lang,§ and Raghav Malhotra¶

Abstract

We exploit employment trends to uncover changes in skills' productivities. Whereas Autor, Levy, and Murnane (2003) study the *degree* to which routine-intensity can rationalize employment trends, our reverse approach characterizes the *kind* of technological change that best explains shifts. We combine a tractable GE model with three editions of the Dictionary of Occupational Titles, three versions of the O*NET, the 1960, 1970, 1980, 1990, Censuses, 2005-07, 2015-17 American Community Survey, and March Current Population Surveys to estimate changes in the relative productivity of skills. We conclude that routine productivity grew rapidly while abstract-skill productivity lagged, a form of 'skill bias'. Together with substitutability between abstract and routine inputs, these results explain changes in skill use within occupations.

*This research was funded in part under NSF grant SES-1851636. It draws heavily on our earlier working paper, "The Evolution of Skill Use Within and Between Jobs," but extends the analysis by over two decades. We gratefully acknowledge the comments and suggestions on that and/or this paper from participants in the NBER Personnel group, Brigham Young University, Cambridge University, George Washington University, the Gilbert Memorial Lecture at the University of Rochester, Hong Kong University, University of Wisconsin-Madison, the Society of Labor Economists, the Applied Research in Labor Economics and Human Capital workshop in Moscow, and the European Economic Association. We are truly grateful to Pascual Restrepo for excellent suggestions and to David Dorn for two very helpful conversations. Nevertheless, we absolve them from responsibility for any errors remaining in the paper and the opinions expressed herein. The views expressed in this paper are of the authors and are not necessarily those of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System

†C.Cavounidis@warwick.ac.uk (University of Warwick)

‡vittoria.dicandia@clev.frb.org (Federal Reserve Bank of Cleveland)

§lang@bu.edu (Boston University, NBER, and IZA)

¶r.malhotra@leicester.ac.uk (University of Leicester)

1 Introduction

Consider the IBM Selectric, an electronic typewriter introduced in 1961. It replaced the traditional strikebars with a golf-ball-like element and, in later versions, was even 'self-correcting.' The Selectric made typing much more productive. Secretaries and typists could produce many more and more attractive typewritten pages. Typists who quickly caught a mistake could correct it invisibly, while previously, they retyped the page entirely.

One could say this change was biased towards specific tasks in the economy. In particular, it made typists far more productive than they were earlier. However, the interaction with task demand seems ambiguous. The Selectric did not replace typists or substitute for typing outputs, the approach used by the task model (Autor et al. [2003], Acemoglu and Autor [2011], Acemoglu and Restrepo [2018]). If anything, the demand for typed pages increased after the Selectric's introduction.

Neither did the Selectric make high-skilled labor more productive or more in demand, as in the canonical SBTC models (Katz and Murphy [1992], Berman et al. [1994], Berman et al. [1998], Juhn [1999]). Instead, it simply increased the speed at which everyone could type. Thus, the individuals whose productivity increased dramatically were primarily middle-skilled workers in typing-intensive jobs. Further, this technological innovation might have changed the skills used in specific jobs, e.g., by expanding the set of occupations involving some typing.

Thus, the Selectric had two effects within occupations: a primitive effect, which made typing significantly more productive, and an adaptation effect, in which workers respond to the primitive effect by altering their use of typing skills. At first glance, one may expect that the second effect acts as a confounding factor, making it hard to recover the primitive technological change. However, we demonstrate that, in a relatively general model, a first-order approach allows pure employment changes to act as a "sufficient statistic" to identify relative skill productivity changes.

We integrate insights from both task-based and skill-biased approaches in a tractable general equilibrium framework. We model occupations as combining skills, akin to tasks in the tasks model, to produce intermediate goods. Occupations utilize workers' skills in heterogeneous ways, and we impose nearly no structure on their production technologies. Our model allows workers to choose first the skills they develop and then their occupations. The market prices intermediate goods that aggregate to a single final good.

As in the SBTC model, we allow for skill-enhancing technological change. The Selectric made typing - or routine skill - more productive. The number of typed pages produced must increase, but whether workers deepen their typing skills or not depends on the elasticity of substitution between skills in each specific occupation. Routine skill use (typing) can decline in one occupation (secretaries) but increase in another (economics professors). Employment in typing-intensive occupations can increase or decrease. If the elasticity of demand for such an occupation's output is less than one, as we believe plausible, demand for that occupation falls.

We also allow for shifts in product demand (possibly due to trade shocks) or outside competition (possibly due to robots or offshoring) that alter the demand for workers with different skills. The model thus clarifies the distinction between technological changes to the productivity of individual skills and changes to demand for particular kinds of workers.

The model provides us with a simple approach to measuring the relative increase in the technological productivity of different skills while taking demand shifts into account. In effect, we develop a transparent structural model for interpreting within and between-occupation changes in skill use that, for local estimation, relies only on ordinary least squares and weighted means, using readily observable variables.

Notice this is orthogonal to the exercises conducted by Autor et al. [2003] and Goos et al. [2014]. They identify the technological change they believe to be important, namely the routine intensity of occupations, as a measure of vulnerability and study its relationship with employment changes. Our approach instead leverages employment changes to identify what is the relevant technological change. Autor et al. [2003] observe the correlation between computerization and performance of routine tasks and show that this type of technological change can provide relevant insights into changes in employment in the United States. Goos et al. [2014] study the role of routine biased technological change and offshoring in explaining changes in employment in 16 European countries, providing evidence of a much bigger role played by the former.

Kogan et al. [2021], too, adopt an approach orthogonal to ours. They create a measure of the similarity between the technology introduced by patents and the tasks performed in an occupation as a proxy of exposure to technological advancement and use it to study its association with changes in employment and wages over a time span of almost two centuries. Bárány and Siegel [2020] estimate productivity change down to the sector/occupation level, assuming that each occupation uses only a single skill, whereas our occupations mix different skills in different amounts, and we account for sector-level demand. Acemoglu and Restrepo [2019] features the emergence of entirely new occupations; our analysis is based only on the *relative* employment in existing occupations, and is thus uninfluenced by such occurrences.

We depart from the skill-weights approaches of Lazear [2009], Gathmann and Schönberg [2010], and Cavounidis and Lang [2020] by allowing the production function translating skills or tasks into output to be a general constant-returns-to-scale neoclassical production function. The earlier papers assume that output in each occupation is a linear function of skills, with occupation-varying weights. While Yamaguchi [2012] uses a somewhat more general specification for determining wages, it, too, makes wages in each occupation a linear function of the worker's skills. Moreover, Yamaguchi [2012] limits the analysis to cognitive and motor tasks. In addition, these papers focus on mobility across occupations and skill acquisition, either by investment or learning by doing, among individual employed workers. We abstract from the latter and focus on labor market equilibrium.

To estimate the model, we use the skills studied by Autor et al. [2003] and measured in the Dictionary of Occupational Titles, using the third edition for skill use in 1960, the original fourth edition for 1971, and the revised fourth edition for 1983. We combine these measures with data from the Current Population Surveys and Censuses to measure between and within-occupation changes in skill use from 1960 to 1983. We also create [1] and use skill measures from the 3.0, 12.0, and 24.2 versions of the O*NET. We combine these measures with data from the Current Population Surveys, American Community Surveys, and Censuses to measure between and within-occupation changes in skill use from 1995 to 2016. We focus on changes within each of the periods covered by the *DOT* and O*NET. For reasons discussed in the text, we are skeptical of changes occurring during the transition between the two sources.

We find that workers moved into abstract-intensive and away from routine-intensive and manual intensive occupations in all periods for which we have data. We also show that within-occupation shifts can dwarf those due to movement across occupations. In all periods, abstract-skill use grew within occupations while routine-skill use fell. The pattern for manual-skill use is less consistent but fell in most periods. We show how complementarity and substitutability of skills with respect to their own and other skills' growth in productivity explain these patterns.

We are not the first to look at within-occupation changes in skill use. Spitz-Oener [2006] and Black and Spitz-Oener [2010], using German data, and Deming and Noray [2020], using Burning Glass data, track significant within-occupation shifts in skill use, but for a later period. Atalay et al. [2020], using keyword frequencies from three newspapers' job ads over an impressively long period, show that within-occupation changes account for most task variation over time. It is an open question as to how representative these ads are. Consoli et al. [2023] is closest to this paper. It examines within-occupation changes in routine-task intensity (RTI) from 1980-2010. Much of their paper focuses on reconciling the *DOT* and O*NET measures so that they can examine changes between 1990 and 2000 (or more precisely, between the 1991 revision of the 4th edition of the *DOT* and the 2000 O*NET version

[1] with the exception of abstract-interpersonal skills, which are those used by Acemoglu and Autor [2011]

3.0). In our preliminary work, we, like them, found a large shift in within-occupation skill use when shifting between the two sources. We have, therefore, chosen to exclude this period from our analysis. [2] Most significantly, we develop a model to help us interpret the results and apply it to a very long period. Autor and Price [2013] also study a very long period but do not allow for within-occupation changes in skill use. This paper can be read in two ways. Those interested solely in a better accounting of the changes from the 1960s to 2016 can jump to the data section and then examine Tables 1 and 2 and the accompanying text in the results section. We think this analysis is a contribution in its own right. However, we are hopeful that readers will find that the model presents a simple, versatile framework allowing for different kinds of technological shocks and, therefore, assists in thinking about our results and the large literature in this area.

2 A model of skill and job choice in general equilibrium

2.1 Skill acquisition and intermediate good production

Before employment, each worker chooses a vector of skills $S \in \mathbb{R}^n_+$, where each component S_i reflects ability at task i. Once workers have acquired skills, each chooses a job $J \in \mathcal{J}$, where $\mathcal{J}$ is the set of all jobs. If a worker with skills S is employed at job J, she produces a quantity $y((A_iS_i)_{i\leq n}, J)$ of intermediate good J, where each $A_i > 0$ is common to all jobs and is a measure of the general productivity of skill i. Thus, each A_iS_i is the 'effective' amount of input i.[3]

We place as little structure on $\mathcal{J}$ and y as possible. We assume only that $\mathcal{J}$ is a compact subset of a Euclidian space, that $y(\cdot, J)$ is a constant-returns standard neoclassical production function,[4] and that y is continuous.

For simplicity, we assume that workers have a fixed budget for skills, which we normalize to 1, so that for any individual $\Sigma_i S_i = 1$. This assumption captures the idea that a worker can study plumbing or philosophy, but if she chooses to spend more time on philosophy, she must spend less time learning plumbing. We do not allow her to choose to spend more time

[2] Autor et al. [2003] examine the relation between computer use and within-occupation change in task use between the 1977 and 1991 revisions of the *DOT* but do not discuss the magnitudes of these changes.

[3] Thus output y depends on the vector of effective inputs $(A_iS_i)_{i\leq n}$.

[4] $y(\cdot, J)$ is strictly increasing in each A_iS_i on $\mathbb{R}^n_{++}$, is twice continuously differentiable, features a bordered Hessian with non-vanishing determinant on $\mathbb{R}^n_{++}$, is strictly quasiconcave, and $y((A_iS_i)_{i\leq n}, J) = 0$ iff $A_iS_i = 0$ for some i. This will imply that optimal skills are continuously differentiable in A and, more importantly, interior. If skills are quite occupation-specific, e.g., plumbing or surgery skills, this may be a bad assumption; however, the skills used in our empirical section are relatively general. We thus think that excluding corner solutions is unproblematic for our application.

on learning.[5]

A worker who anticipates holding job J will therefore

$$\max_{S\geq 0} y((A_iS_i)_{i\leq n}, J) \tag{1}$$

$$\text{subject to} \sum_i S_i = 1. \tag{2}$$

The optimal $S^*(J)$ and $y^*(J) := y((A_iS_i^*(J))_{i\leq n}, J)$ are given by solving the Lagrangian. The Lagrangian's first order condition at the optimum with respect to any S_i is

$$A_iy_i'((A_iS_i^*(J))_{i\leq n}, J) = \lambda = y^*(J) \tag{3}$$

where the second equality follows straightforwardly from constant returns to scale. We assume that workers always have skills that are optimal for the job they perform. Although this assumption is strong, we maintain that in the sort of timescales our empirics cover, workers will, at the least, endeavor to develop the right skills for the careers they select. Allowing for investment while employed, as in Cavounidis and Lang [2020], would make this a sensible assumption for workers not too far advanced in their work lives.

How do optimal output and skills change with A? From the Envelope Theorem,

$$\frac{\partial y^*(J)}{\partial A_i} = S_i^*(J)y_i'((A_iS_i^*(J))_{i\leq n}, J) \tag{4}$$

so that substituting for y_i' using (3), we get

$$\frac{\partial \ln y^*(J)}{\partial \ln A_i} = S_i^*(J). \tag{5}$$

This is effectively an application of Roy's Identity, with our skill constraint playing the role of the budget constraint in standard utility maximization.

To speak sensibly about the effect of changes in A on $S^*(J)$, we proceed by inspecting $y(\cdot, J)$'s i-j elasticity of substitution for any two inputs at the optimum

$$\sigma_{i,j}((A_iS_i^*(J))_{i\leq n}, J) = \frac{\partial ln\left(\frac{A_iS_i^*(J)}{A_jS_j^*(J)}\right)}{\partial ln\frac{A_i}{A_j}} = 1 + \frac{\partial \ln(S_i^*(J)/S_j^*(J))}{\partial \ln(A_i/A_j)} \tag{6}$$

[5]This is without loss of generality since we can always normalize the time she chooses to spend on learning to 1. This could affect comparative statics on total production through a labor/leisure/learning trade-off. That said, since this only adjusts the effective number of labor units each worker provides, with a constant returns to scale aggregate production function, it will not affect the objects of interest to us.

which we can rearrange as

$$\frac{\partial \ln(S_i^*(J)/S_j^*(J))}{\partial \ln(A_i/A_j)} = \sigma_{i,j}((A_i S_i^*(J))_{i\leq n}, J) - 1. \tag{7}$$

Thus, if inputs i and j are gross substitutes (complements) in job J at the optimal skill bundle, a relative increase in the productivity of skill i will cause workers to acquire relatively more (less) of it. If all inputs are gross substitutes (complements) in job J at the optimal skill bundle, the constraint that $\sum_i S_i^*(J) = 1$ further implies that $\frac{\partial S_i^*(J)}{\partial A_i} > 0$ (< 0).

2.2 Final good production and worker allocation

So far, the model somewhat resembles Cavounidis and Lang (2020) in the sense that workers are aligning their skill choices and occupation choices. We extend it by assuming that instead of goods of intrinsic value, workers produce inputs in a CES final good production function

$$Y(q) = \left[\int_{\mathcal{J}} h(J) q(J)^{\varepsilon} dJ\right]^{\frac{1}{\varepsilon}}. \tag{8}$$

Here, $h(J)$ is the relative importance of input J for final production, and $q(J)$ is the total quantity of intermediate good J used as an input. We assume h is continuous. The economy has workers of total measure 1, and each worker acquires skills, subject to the constraint, and may choose any job in $\mathcal{J}$.

The model satisfies conditions under which the decentralized equilibrium is Pareto efficient. Therefore, we solve for the equilibrium by solving the planner's problem subject to the skill acquisition and worker measure constraints. Efficiency implies that workers producing good J will all be identical and acquire skills $S^*(J)$; therefore, $q(J) = y^*(J)f(J)$, where $f(J)$ is the density of workers assigned to producing intermediate good J.

Therefore, we can write the planner's problem as

$$\max_f \left[\int_{\mathcal{J}} h(J) \left[y^*(J) f(J)\right]^{\varepsilon}\right]^{\frac{1}{\varepsilon}} \tag{9}$$

$$\text{subject to} \quad \int_{\mathcal{J}} f(J) = 1. \tag{10}$$

We can then pointwise differentiate the Lagrangian and obtain

$$h(J) y^*(J)^{\varepsilon} f(J)^{\varepsilon - 1} = h(J') y^*(J')^{\varepsilon} f(J')^{\varepsilon - 1}, \tag{11}$$

which we can write as

$$f(J)h(J')^{\frac{1}{1-\varepsilon}}y^*(J')^{\frac{\varepsilon}{1-\varepsilon}} = f(J')h(J)^{\frac{1}{1-\varepsilon}}y^*(J)^{\frac{\varepsilon}{1-\varepsilon}} \tag{12}$$

so that we can now integrate out J' and using constraint (10) get

$$f(J) = \frac{h(J)^{\frac{1}{1-\varepsilon}}y^*(J)^{\frac{\varepsilon}{1-\varepsilon}}}{\int_{\mathcal{J}} h(J')^{\frac{1}{1-\varepsilon}}y^*(J')^{\frac{\varepsilon}{1-\varepsilon}}}. \tag{13}$$

2.3 Comparative statics

We consider the effect of technological progress that is broadly skill enhancing, as measured by A, and changes in the demand for intermediate goods, as measured by h. The distinction is imperfect. For example, the reduction in transportation costs, at least partly due to technological change, reduced demand for some locally produced intermediate goods that had hitherto been too expensive to import. Still, we think of changes in A as capturing broad-based technological progress, such as electronic calculators rather than adding machines for routine-cognitive skills and electric rather than manual drills for manual skills, and h as capturing the effects of trade and, more recently, robots.

2.3.1 The effect of skill-augmenting technological change

What happens if skill i becomes more productive? Taking the derivative of (13) with respect to A_i gives

$$\frac{\partial f(J)}{\partial A_i} = \frac{\varepsilon}{1-\varepsilon} f(J)\left[\frac{\partial \ln y^*(J)}{\partial A_i} - \int_{\mathcal{J}} \frac{\partial \ln y^*(J')}{\partial A_i} f(J')\right] \tag{14}$$

or simply, using (5),

$$\frac{\partial \ln f(J)}{\partial \ln A_i} = \frac{\varepsilon}{1-\varepsilon}\left[S_i^*(J) - \int_{\mathcal{J}} S_i^*(J')f(J')\right]. \tag{15}$$

In other words, if and only if the elasticity of substitution among intermediate goods $1/(1-\varepsilon)$ is less than 1, will an increase in the productivity of skill i move workers away from jobs where it is used more than average, and towards jobs where it is used less than average. So, for example, if routine skill is a complement to other skills in intermediate good production, and intermediate good demand is inelastic, an increase in A_R (a technological change that makes routine skill more productive) will (a) reduce routine use in all jobs (within) and (b) shift workers to less routine-intensive jobs (across).

The idea that sectors experiencing slower productivity growth also experience faster employment growth is old (Baumol [1967], see also Ngai and Pissarides [2007] and Acemoglu and Guerrieri [2008]). We build on that idea. In our case, jobs that make more use of skills whose productivity grows slowly will experience more employment growth.

2.3.2 The effect of changes in demand for intermediate goods

What about changes in h? In our setup, these will move workers around but have no effect on skill use within a job. A decrease in horseshoe demand merely alters how many people shoe horses, not how they shoe them.

To see the effect of changes in h on employment, we take the log of each side in (13) and totally differentiate to get

$$d\ln f(J) = \frac{1}{1-\varepsilon} d\ln h\,(J) + \frac{\varepsilon}{1-\varepsilon} d\ln y^{*}(J) - d\ln\left(\int_{\mathcal{J}} h(J')^{\frac{1}{1-\varepsilon}} y^{*}(J')^{\frac{\varepsilon}{1-\varepsilon}}\right). \tag{16}$$

For a change in h, the second term in (16) is 0 and the third term does not depend on J. A few manipulations yield

$$d\ln f(J) = \frac{1}{1-\varepsilon}\left[d\ln h\,(J) - \int_{\mathcal{J}} d\ln h(J') f(J')\right]. \tag{17}$$

Thus, the percentage employment growth in job J is proportional to the deviation of the percentage change in $h\,(J)$ from the employment-weighted average.

2.3.3 Putting it all together

Combining (15) and (17), we have

$$\begin{aligned} d\ln f(J) &= \frac{\varepsilon}{1-\varepsilon}\sum_i \left[S_i^{*}(J) - \int_{\mathcal{J}} S_i^{*}(J') f(J')\right] d\ln A_i \\ &\quad + \frac{1}{1-\varepsilon}\left[d\ln h\,(J) - \int_{\mathcal{J}} d\ln h(J') f(J')\right]. \end{aligned} \tag{18}$$

The model distinguishes between changes that replace (or reduce demand for) occupations by automating or offshoring them (a decline in h) as when data input is imported from abroad, and those in which technology makes relevant skills more productive as when key-punch machines are replaced by input at computer terminals. When h declines, the number of workers employed in data entry in the home country falls, but any workers engaged in data input continue to input data using the same skill set. Suppose the productivity A_i of a skill i important to data entry increases. If skill inputs are complements at data entry and intermediate-good demand is inelastic, workers in data entry jobs end up with less of skill i, and fewer workers are hired to input data.

Interpreted within our model, Autor et al. [2003] found that, in a later period, technological innovation increased the productivity of routine skills. Since the demand for these skills was inelastic, the amount of time individual workers spent on them declined as did total employment in routine-intensive occupations. Our interpretation of the period that we study will be that the productivity of abstract skill use did not increase as rapidly as the productivity of other skills, most notably routine skill. This caused a shift towards abstract-skill use because the elasticity of substitution between intermediate goods is less than one, thereby shifting employment to abstract-intensive occupations. Within occupations, declining relative abstract-skill productivity shifted skill use toward greater abstract and less routine-skill use.

We note that our model assumes *ex-ante* identical workers. In a richer model with *ex-ante* heterogeneous workers, demand changes might alter how jobs are done. Intuition suggests that workers "better at routine tasks" do jobs more routinely than other workers. In such a world, a reduction in demand for routine-intensive outputs *would* shift such workers to less-routine jobs who would then perform them *more* routinely than before, which is the reverse of what we observe.

2.4 Implications for empirical work

For empirical analysis, we rewrite (18) as

$$\Delta \ln(emp_{I,J}) = \frac{\varepsilon}{1-\varepsilon}\Sigma_i \left(d\ln A_i \left(S_{i,J} - \overline{S}_i \right) \right) + \gamma_I + \mu_{I,J} \tag{19}$$

where $\Delta \ln(emp_{I,J})$ is the change in the employment level in industry I in occupation J, the empirical counterpart of $f(J)$ and γ_I is the coefficient on an industry that captures demand changes due to shifts in industry demand. We note that this is an imperfect proxy for changes in h. It will capture changes in demand for an occupation resulting from, for example, import competition but not changes due to occupation-specific factors such as robots, although we note that robot penetration is typically measured at the industry or braod geographic level. We measure $S_{i,J}$ by its average in two proximate editions of the *DOT* or O*NET. μ is a mean-zero error term. We estimate (19) separately for each time-period pair.

Since each worker's skills sum to 1; skill use on a job sums to 1, as does mean skill use. Therefore, (19) still applies if we add a constant term to each $d\ln A$, and we can rewrite the equation as

$$\begin{aligned} \Delta \ln(emp_{I,J}) &= \frac{\varepsilon}{1-\varepsilon}\Sigma_i \left(\left(d\ln A_i - d\ln \overline{A} \right) \left(S_{i,J} - \overline{S}_i \right) \right) + \gamma_I + \mu_{I,J} && (20) \\ &= \frac{\varepsilon}{1-\varepsilon}\Sigma_i \left(\left(d\ln A_i - d\ln \overline{A} \right) S_{i,J} \right) + \gamma_I + \mu_{I,J} && (21) \\ &: = \Sigma_i S_{i,J} \beta_i + \gamma_I + \mu_{I,J}. && (22) \end{aligned}$$

Equation (22) describes a regression of the (approximate) percentage change of employment in an occupation/industry cell on the skills used in that occupation and industry dummies. The coefficients show the change in each skill's productivity relative to the average up to a factor of proportionality. This factor is negative if the elasticity of substitution between intermediate goods is less than 1, which we assume. Thus, a negative coefficient means that the productivity of that skill grew faster than the average of the skills.

Although derived quite differently, our final equation is similar to the one in Goos et al. [2014]. Their theoretical model includes wages in the equivalent of (22), which they proxy by industry-year and occupation dummies.[6] Since we first-difference the data and estimate the model separately for each pair of years, we implicitly control for occupation and year while explicitly controlling for industry. They also use an alternative specification in which they explicitly control for wages but do not include it in the main text as there are concerns about their endogeneity. While we agree with such concerns, we perform the same exercise and observe that the inclusion of wages does not alter the outcome of our analysis. [7] The major difference in our specifications is that they include only routine-task intensity and not the other skills but also include a measure of offshorability.

Assuming an elasticity less than 1 seems natural. As Jones [2011] notes in a somewhat different context, intermediate goods are unlikely to be substitutes. As he puts it, computers are close to essential for producing some goods. Consistent with this argument, Goos et al. [2014] estimate that the elasticity of substitution across industry outputs is 0.42. Our case is even stronger; the outputs of secretaries, sales workers, plumbers, and truck drivers cannot easily substitute for each other. Note that this is different from the statement that someone who works as a secretary might be almost as productive if he worked in sales. This is entirely plausible in our model if the required underlying skills are close.

Note that we must drop a skill because the skills sum to 1. Therefore, we can interpret the coefficients as the rate of growth of productivity of each skill relative to the excluded skill, again up to a multiplicative factor. Together with the requirement that the sum of the deviations from average productivity growth equals 0, this fully identifies the relative productivity of all the skills.

Equation (22) addresses only changes in the productivity of skills and not shifts in the demand for occupations except through the inclusion of the two-digit industry dummies, in line with the effect of h in (17). Demand for occupations concentrated in industries facing import competition or declining demand will fall even absent technological change.

[6] We ignore the country component since we study only one country.

[7] We do not model wages in our current framework. A case in favor of their inclusion could be made if we assumed the labor supply to an occupation to be elastic but not infinitely elastic. However, this is unnecessary given that, empirically, wages do not have an effect on our results.

Controlling for industry will capture employment losses due to import competition but not robots or outsourcing of specific occupations to other countries. Fortunately, in the period we study, these sources of employment loss are likely to be modest.

We estimate (22) by ordinary least squares. Consistent with Solon et al. [2015] and Dickens [1990], we experimented with feasible weighted least squares and found no evidence of important heteroskedasticity with respect to occupation size.

3 Data

Following Autor et al. [2003], our skill-use measures for the first part of our time period come from the *Dictionary of Occupational Titles (DOT).* We use the third edition, issued in 1965 but compiled starting sometime after the release of the second edition in 1949, as our measure of skill use in an occupation in 1960, although it may be centered more on the late 1950s. The 1965 *DOT* has not, to the best of our knowledge, been previously used for this type of analysis. We use the fourth edition, published in 1977 and based on data starting in 1965 for job use in 1970-72 ('1971'). Finally, we use the last revision of the fourth edition, based on revisions from 1977 to 1991 for skill use in 1982-84 ('1983'). The files for the 4th and revised 4th versions of the *DOT* come from Autor et al. [2003]. As others have noted, the revised fourth edition is not a 'fifth' edition; many occupations were not revisited between the fourth edition and the revised 1991 edition because the revision addressed only occupations believed to have changed the skills they used. Therefore, we probably underestimate the extent of within-occupation changes in skill use between 1971 and 1983. However, we observe differences between the 4th and revised 4th editions for most of the occupations present in both 1971 and in 1983.[8]

For the second half of our time period, we follow Firpo et al. [2011] in relying on O*NET, which is the successor of the *DOT*, and was first issued in 1998. Since then, there have been multiple revisions of the O*NET. Each revision updates a subsample of occupations. We use the 3.0 version issued in 2000 for skill use in 1994-96 ('1995')[9], the 12.0 version issued in 2007 for skill use in 2005-07 ('2006')[10], and the 24.2 version issued in February 2020 for skill

[8] We do not observe a change in the use of abstract and routine skills for 4 occupations, and of manual skills for 22 occupations. It should be noted that these are changes for Census occupations, while skill use is reported for *DOT* occupations. Therefore, we observe a change for a Census occupation even if only one of the *DOT* occupations it comprises has been updated.

[9] It is not possible to use the first version of the O*NET because it used a classification of occupational titles that can't be aggregated at the Census occupation-level. For those occupational titles that have a one-to-one correspondence between the 1998 and the 2000 classifications, we do not observe any differences in the reported skill levels. Hence, it is probable that the only changes in skills observed between the two versions are due to the different classifications of occupational titles.

[10] This is the latest version of the O*NET issued prior to the Great Recession.

use in 2015-17 ('2016')[11].

The *DOT* identifies aptitudes, temperaments, and abilities used in a job and measures them numerically. The O*NET identifies abilities, skills, work activities, and work contexts. In both datasets, observations are at the occupation-title level.

The 1965 *DOT* includes all of the skill-use (task) measures used in Autor et al. [2003]. With some small caveats discussed below, it recorded them on the same scales as the later edition, allowing us to have consistent skill measures. Of course, we cannot be sure that individuals interpreted the measures in the same way in the 1950s, 60s, and 70s, but we see no reason that this concern should be greater than for many measures used to compare time periods or geographies.

The one small change is that the earlier edition provides a single measure of "General Education Development" while the later releases measure reasoning, mathematical, and language development separately. We experimented with using the average or the maximum of these three to generate a single measure comparable to the 1965 measure and checked whether this affected the correlation between the third and fourth edition measures. The correlations were similar. Looking across groups did not create a strong case for either. We present results using the average of the reasoning, mathematical, and language development measures for General Education Development in the 1977 and 1991 *DOT*s. In addition, the 1965 *DOT* sometimes provides more than one value of an aptitude, temperament, or ability for a single job title. In such cases, we use a simple average of the values reported.

Like Autor et al. [2003], we measure routine-cognitive skill using the variable "adaptability to situations requiring the precise attainment of set limit, tolerances, or standards," routine-manual skill by "finger dexterity," manual skill by "eye-hand-foot coordination", abstract-interactive skill by "adaptability to accepting responsibility for the direction, control, and planning of an activity". For our measure of abstract-cognitive skill, we use "General Education Development" rather than only its mathematical component, to allow for consistency across all *DOTs*.

As previously stated, the O*NET is a successor of the *DOT*, but it does not have the exact same variables as its predecessor, nor does it use the same scales to measure them. We select variables, some of which are shared with Acemoglu and Autor [2011], that are present and measured consistently in all of the used versions of the O*NET.[12] Among the

[11] This version of the O*NET has been chosen because it is the last one issued prior to the COVID-19 pandemic, and thus it maximizes the number of updated occupational titles without the skills data being affected by the pandemic. The period 2005-17 is preferred to later ones, to avoid complications due to the changes in the classification of Census Occupations occurred in 2018.

[12] As with the *DOT*, we can't be sure that the interpretation of the measures was consistent over time, but we don't see this as a reason for concern.

variables used in the *DOT*, “finger dexterity”is the only one present in the O*NET as well. We thus use it as our measure for routine-manual skill. We measure manual skill with “gross body coordination”, in absence of a perfect match between the two datasets. We measure routine-cognitive skill using the average of “importance of being exact or accurate”, “importance of repeating the same tasks”, “number facility”. Like Acemoglu and Autor [2011], we measure abstract-interactive skill by the average of “establishing and maintaining interpersonal relationships”, “guiding, directing, and motivating subordinates”, “coaching and developing others”.

For our measure of abstract-cognitive skill, we use an average of the three components of General Education Development used in the *DOT*. More specifically, we use “mathematical reasoning”for the mathematical component, an average of “oral comprehension”, “oral comprehension", and “written comprehension”for the language component, “critical thinking”, “processing information”, and “making decisions and solving problems”for the reasoning component[13].

Like Autor and Dorn [2013], we then use the average of routine-cognitive and routine-manual skill for our routine skill measure, and the average of abstract-cognitive and abstract-interactive skill for our abstract skill measure in each period. For each census occupation, we use a weighted average (by employment share) of the skill use in the *DOT* or O*NET occupations comprising that census occupation.

For consistency with our theoretical model, we depart from Autor et al. [2003] and Autor and Dorn [2013] in how we use these measures. Autor et al. [2003] use the absolute value of each skill, while Autor and Dorn [2013] focus on routine intensity defined as ($RTI = ln(R) - ln(M) - ln(A)$).[14] Instead, we first scale the absolute level of skill use by where it lies between the maximum and minimum of that skill's use in any occupation over our two sample periods corresponding to the two different datasets. Thus, use of skill i in occupation J at time t is:

$$\widetilde{skill}_{i,J,t} = \frac{skill_{i,J,t} - skill_i^{min}}{skill_i^{max} - skill_i^{min}} \tag{23}$$

where $skill_{i,J,t}$ is the value obtained directly from the *DOT* or O*NET measures aggregated at the occupation level, $skill_i^{min}$ and $skill_i^{max}$ are the minimum and maximum absolute values (at the occupation level) for skill i in any version of the *DOT* or O*NET. Finally, we compute

[13] These variables were chosen on the basis of the General Education Development components description in the 1977 *DOT*.

[14] We, like everyone else in this literature, have to treat the ordinal measures in the *DOT* and the O*NET as measured on an interval scale. We do so with an unusual level of chagrin given that one of us has pointed out (Bond and Lang [2013], Bond and Lang [2019]) that findings can be sensitive to how an ordinal scale is converted to an interval scale. Unfortunately, the approaches in Bond and Lang are not available to us in this setting.

the share of each skill in the overall sum

$$S_{i,J,t} = \frac{\widetilde{skill}_{i,J,t}}{\Sigma_k \widetilde{skill}_{k,J,t}} \tag{24}$$

so that our four skill measures sum to 1.

Census occupations are more highly aggregated than the *DOT's* and the O*NET's job titles. Following Autor et al. [2003], we use the *DOT*-augmented version of the April 1971 Current Population Survey for this aggregation in the first half of our time period since this is the only dataset with both *DOT* and census codes. For the second part of our time period, we use the Occupational Employment and Wage Statistics of the U.S. Bureau of Labor Statistics, which provide annual employment for each occupational title, as well as a crosswalk between occupational titles and Census occupations.

We use the consistent occupation system created by Dorn [2009] and the crosswalk files provided by Autor and Dorn [2013], linking these occupations to previous census classifications. This gives us 204 occupations in the initial period, 259 in the second period, and 323 in the third period, 318 in the fourth period, 320 in the fifth period and 308 in the last period. We create the occupation skill measures using occupation weights from all full-time workers not living in group quarters between age 18 and 64 in the IPUMS 1960 5% sample, in the IPUMS 1970 1% State sample, the IPUMS 1980 5% sample, the IPUMS 2000 5% sample, the 2005-2007 ACS 3-year sample, and the 2015-17 ACS year samples.

Our data on the occupation distribution come from the Census (IPUMS), from March (Annual Social and Economic Supplement) Current Population Surveys (CPS), and from the American Community Surveys (ACS) and are limited to workers age 25-64, but otherwise, our sample restrictions are the same as for the calculation of the skill weights. Since economists know these data well, we do not describe them here. With the exception for the last two periods, for which we always use the 2005-07 ACS 3-years sample and the 2015-2017 ACS year samples, our choice of which sources to use for different purposes reflects an admittedly arbitrary trade-off between sample size and proximity of the employment data to the timing of the *DOTs* or O*NET. Before 1968, the CPS coded occupations in fewer than forty categories and did not use the Census classification. Therefore, we use the 1960 1% Census sample for our initial period. We rely on the 1970, 1980, and 1990 Census samples for the three later periods when we believe greater accuracy in estimating the employment cells is critical. Thus, we use the censuses or the ACS to aggregate from *DOT* and O*NET to census occupations and when using occupation/industry cells as observations in our regressions. Our decomposition of skill use into within and between-occupation changes relies on occupation, not industry, and therefore, uses larger cells. Consequently, we use the current occupation in the 1970-72, 1982-84, and 1994-96 March CPS for this purpose.

4 Results

Table 1 shows how the use of the three skills evolved from 1959 to 2016. Recall, however, that the measures in the first and last three periods are not comparable. Therefore, when discussing trends, we focus on changes within each sub-period. If the changes within each sub-period are directionally similar, it is plausible that the change applies to the entire period, but we cannot be sure that this is correct.

Subject to this caveat, we see a steady increase in the use of abstract skills. Abstract-skill use grew from .34 to .42 between 1959 and 1971 and from .39 to .45 between 1995 and 2016. If we believed that the *DOT* and O*NET measures were comparable, we would conclude that abstract-skill use fell between 1983 and 1995. We find differences between the measures to be a more plausible explanation for the decline.

As abstract-skill use increased, the use of the other skills increased. Note that this must be true in aggregate since skill use sums to 1. However, we observe declines in the use of each of the other two skills over the full period. There are, however, some notable departures from trends. Nonroutine-manual-skill use increased in the 1960s, while the use of routine skill remained steady during the most recent period.

Differences in the standard deviations of skill use across occupations reinforce our concerns about treating the results from the *DOT* and O*NET as a single time-series. Except for manual skill in 1959, the standard deviation of skill use is higher in the *DOT*. We note that our results suggest that routine-skill use declined between 1983 and 1995. Previous work has used the two data sources to document this decline. While we expect that routine-skill use did decline over this period, we urge caution in combining the two data sets.

4.1 Within-occupations changes in skill use are important

Table 2 decomposes skill-use changes into within and across-occupation changes using the following decomposition:

$$Skill_{e+1,t+1} - Skill_{e,t} = \underbrace{(Skill_{e+1,t+1} - Skill_{e+1,t})}_{\Delta \text{ across}} + \underbrace{(Skill_{e+1,t} - Skill_{e,t})}_{\Delta \text{ within}} \tag{25}$$

where e indicates the *DOT* or O*NET edition, and t indicates the period considered. Thus, Δ within shows how skill use would have changed had the occupations in which, for example, people worked been the same in 1959 and 1971. In parallel, Δ across shows how much skill use would have changed had skill use in each occupation remained constant between 1959 and 1971, and only the distribution of workers across occupations shifted. With exceptions discussed in the introduction, this latter measure corresponds to that typically presented in the literature. Therefore, we begin with across-occupation changes. Differences from the prior literature may reflect our use of different editions of the *DOT* and *O*NET* and/or our somewhat different use of the skill measures.

The across-occupation patterns we observe are consistent with the prior literature. In each period, we observe movement towards abstract-intensive occupations. This shift was particularly large in the 1970s. Much of the change over the full period reflects the interpersonal component of abstract (not shown). The movement away from routine-intensive occupations was particularly notable during the *DOT* periods, while the movement away from manual-intensive occupations was concentrated in the 1970s. occupations that make

As with Table 1, we show changes between 1983 and 1995 for completeness despite our concerns about their reliability. Subject to this caveat, we observe the general across-occupation shift from routine and manual skills to abstract skills reported in Consoli et al. [2023].

Perhaps the most important message of Table 2 is that between-occupation shifts miss a great deal of the action. Within-occupation changes are less consistent from period to period but are often large. Between 1959 and 1971, there were large within-occupation increases in abstract and manual skill use and a large decline in the use of routine skill. Within occupation, abstract-skill use rose notably between 1971 and 1983. Over the first two periods, the within-occupation changes roughly doubled the movement away from routine skill and significantly increased the movement toward abstract skill, while partially offsetting the movement away from manual skill.

In contrast, changes between 1995 and 2016 are almost entirely within occupation. While the shift towards more abstract-intensive and less manual-intensive occupations continued in the period during which we rely on O*NET, these movements were much more limited, even allowing for the lower standard deviation of use in the O*NET relative to the *DOT*. Similarly, the continued decline in routine-skill use is almost entirely within occupation. Moreover, the absence of a continued decline in its use, reflects a slight shift towards more use within occupations.

We note that if we look only across occupations, merging the *DOT* and O*NET gives plausible results. However, we must believe that the within-occupation shifts in abstract and manual-skill use were fundamentally different between 1983 and 1995 than during the 24 years before or 21 years after.

In sum, over our entire period, we observe a sharp shift from routine to abstract-skill use both within and between occupations. We observe a similar shift away from manual-skill use except for a surprising increase in its use within occupations in the 1960s.

4.2 Relative skill-productivity growth matters (sometimes)

Recall that estimating (22) and imposing that the coefficients sum to 0 allows us to identify the relative growth of skill productivity.[15] Table 3 shows the results of this exercise (See Table 1A in the Appendix for non-transformed coefficients). The coefficients in the table measure the relative growth rate of the productivity of the skills multiplied by $\varepsilon/\left(1-\varepsilon\right)$. Assuming that the elasticity of substitution is less than one, then $0 > \varepsilon/\left(1-\varepsilon\right) > -1$, and we can bound the difference relative to the average in the annualized rate of growth over some period by the coefficient divided by the period's length.

The two coefficients are jointly significant in each time period except the first.[16] Our earlier work [Cavounidis et al., 2021] showed skill productivity growth among men was similar across skills during the earliest period but that among women, the productivity of finger dexterity grew noticeably faster than the productivity of other skills. We do not replicate those results here because we cannot get gender-specific skill-use changes in the later period.

A clear pattern emerges from the table. In every period, the productivity of routine skill grew more rapidly than the average (the coefficients are all negative). The coefficients are negative in all four periods and statistically significant at the .1 level in the 1970s and at any conventional level in the last two periods. While the coefficient does not reach statistical significance at even the .1 level in the 1960s, its value is similar in the 1960s and 1970s.

Similarly, the productivity of abstract skill grew more slowly than average in each of the four periods. The coefficients are statistically significant at any conventional level except for during the first period. Still the coefficient is only modestly smaller for that period than for the 1970s. It seems plausible that the faster relative productivity growth of routine skills and slower productivity growth of abstract skills began in the 1960s and accelerated thereafter.

Finally, the coefficients on manual skills are generally small and never significant. This is consistent with manual-skill productivity growth falling between the productivity growth of the other two skills.

We provide further evidence of the importance of differences in skill-productivity growth in the line labeled "proportion due to skills." This shows the proportion of the R-squared

[15] To reduce measurement error, we restrict the sample to occupation/industry combinations comprising at least .0001% of employment in each year included in the pair and at least an average of .0002% over the two years. The second requirement ensures that we do not create this bias by dropping observations near the threshold that saw a modest change in employment that caused it to cross the .01% threshold but keep similarly small occupation/industry observations that happen not to cross the threshold. Nevertheless, many of the employment changes we observe remain implausible. Since occupations are coded consistently across periods, we are not concerned that changes in occupation drive these changes. We winsorize the data fairly severely at the 5th and 95th percentiles. Finally, we average our skill-use measures from the two editions (or the revision) of the *DOT*, and the three editions of the O*NET corresponding to the pair of years in our analysis.

[16] Recall that the mean change is normalized to 0, making one coefficient redundant.

allocated to skill using the Shapley-Owen decomposition. Strikingly, this proportion grows over the four periods. The jointly insignificant skill variables play only a small role in the first period. This doubles in the next period and grows further in each of the two last periods. While accounting for only one quarter of the explanatory power may appear modest, recall that the number of industries dwarfs the two skill variables that are included in the regression. We conclude that relative skill-productivity growth matters at least after 1970.

As noted previously, we also perform this exercise by adding the % change in average wages to the controls. Table 2A in the appendix shows that this doesn't have a meaningful impact on our estimation.[17]

4.3 Slow growth of abstract productivity and fast growth routine-cognitive skill (mostly) explains the within shifts

To understand what our model says about within-occupation skill shifts, we take a linear expansion of $S_i(J)$ with respect to relative changes in skill productivities:

$$dS_i(J) = \Sigma_k \frac{\partial S_i(J)}{\partial \ln A_k} d \ln A_k. \tag{26}$$

Now, we multiply by $f(J)$ and integrate over all jobs

$$\int_{\mathcal{J}} dS_i(J) f(J) dJ = \Sigma_k \left(d \ln A_k \int_{\mathcal{J}} \frac{\partial S_i(J)}{\partial \ln A_k} f(J) dJ \right). \tag{27}$$

Now, we use the fact that $\Sigma_k S_k = 1$ to get

$$\Sigma_k \frac{\partial S_k(J)}{\partial \ln A_i} = 0. \tag{28}$$

A short argument based on Slutsky symmetry and the regularity and constant returns

[17]Although not reported, we also performed this analysis using the change in the average log wage, and reached the same conclusion, with point estimates for skills productivity that are even more similar to those in Table 3. Table available upon request.

to scale assumptions on $y(\cdot, J)$ shows that[18]

$$\frac{\partial S_i\left(J\right)}{\partial \ln A_k} = \frac{\partial S_k\left(J\right)}{\partial \ln A_i} \tag{29}$$

so that we can rewrite (28) as

$$\Sigma_k \frac{\partial S_i(J)}{\partial \ln A_k} = 0. \tag{30}$$

Thus, we can normalize (27) with respect to an arbitrary $d\ln A_n$:

$$\int_{\mathcal{J}} dS_i\left(J\right) f(J) dJ = \Sigma_{k\neq n}\left(d\ln A_k - d\ln A_n\right)\int_{\mathcal{J}} \frac{\partial S_i\left(J\right)}{\partial \ln A_k} f(J) dJ. \tag{31}$$

Denoting the integral on the right by $\partial\overline{S_i}/\partial \ln A_k$, and replacing the left-hand-side with the within estimates in Table 2 and the $d\ln A_k$ terms with the estimates in Table 3, we arrive at

$$\widehat{\text{within}}_i = \Sigma_{k\neq n}\left(\widehat{d\ln A_k} - \widehat{d\ln A_n}\right)\frac{\partial \overline{S_i}}{\partial \ln A_k}. \tag{32}$$

These $\partial\overline{S_i}/\partial \ln A_k$ terms represent the average changes in workers' skills brought on by isolated productivity changes, and we are most interested in extracting them. As Section 4.2 suggests, however, more than one A_k changed in each of our periods, making this exercise nontrivial.

Assuming that these derivative terms do not change over time, we have twelve equations and three unknowns after imposing symmetry per (29). We choose the parameter estimates that minimize the sum of the squared differences between the calculated within change and the predicted within change.[19]

The derivatives, $\partial\overline{S_i}/\partial \ln A_k$, capture a concept analogous to p and q complementarity and substitutability. If the derivative is positive, an increase in the productivity of skill k increases

[18]As we have assumed that $y(\cdot, J)$ is a neoclassical production function subject to a linear skill budget constraint, we can turn to standard demand theory. The arguments of $y(\cdot, J)$, $(A_i S_i)_{i\leq n}$, can be thought of as 'effective' skills . Now, $A_i S_i$ is simply the Marshallian demand for effective skill i, where the price of effective skill i is $1/A_i$. We denote by $A_i S_i^{Hicks}$ the Hicksian demand of effective skill i, and by ω the skill budget constraint. The Slutsky equation is $\frac{\partial(A_i S_i)}{\partial \frac{1}{A_k}} + \frac{\partial(A_i S_i)}{\partial\omega} A_k S_k = \frac{\partial(A_i S_i^{Hicks})}{\partial \frac{1}{A_k}}$. From Slutsky symmetry, $\frac{\partial(A_i S_i^{Hicks})}{\partial \frac{1}{A_k}} = \frac{\partial(A_k S_k^{Hicks})}{\partial \frac{1}{A_i}}$, and from constant returns to scale we have symmetric income effects $\frac{\partial(A_i S_i)}{\partial\omega} A_k S_k = A_i S_i A_k S_k = \frac{\partial(A_k S_k)}{\partial\omega} A_i S_k$. Thus, $\frac{\partial(A_i S_i)}{\partial \frac{1}{A_k}} = \frac{\partial(A_k S_k)}{\partial \frac{1}{A_i}}$, so that $-A_i A_k^2 \frac{\partial S_i}{\partial A_k} = -A_k A_i^2 \frac{\partial S_k}{\partial A_i}$ or simply $\frac{\partial S_i}{\partial \ln A_k} = \frac{\partial S_k}{\partial \ln A_i}$ as desired.

[19]One of the equations within each set of three is redundant. We would have eight independent equations if the linear approximation and our data were both perfect. Rather than arbitrarily discarding one equation, we minimize the sum of the squared deviations.

the amount of skill i acquired by workers. We refer to this case as A-complementarity. Note that, unlike p-complementarity, a skill may be A-complementary or A-substitutable with itself.[20]

Recall that in Table 3, we estimate $\varepsilon / (1 - \varepsilon) * d \ln A_i$. So, as ε is unknown, with a change of sign, the coefficients represent lower bounds on the absolute values of the skill-productivity changes. Therefore, using these coefficients yields upper bounds on the derivatives. Consequently, we focus on the signs of the estimated derivatives rather than their precise magnitude and ignore the $\varepsilon / (1 - \varepsilon)$ term other than to assume that it is negative. Thus, in reading Table 4, which displays the results of this exercise, readers can rely on their intuition to divide the estimated derivative by something in the range 1.3 to 1.7.

As seen in Table 4, all skills are, on average, A-substitutes for themselves, although the substitution effects are small for manual and routine skill productivity. Notably, routine and abstract skills are A-complements. Consequently, the slow productivity growth of abstract skill and fast productivity growth of routine skill both contributed to the within-occupation growth in abstract-skill use and decline in routine-skill use. Manual and abstract skills are A-complements, while (nonroutine) manual and routine skills are A-substitutes. Thus, the patterns of skill growth in these two skills contributed to the within-occupation decline in manual skill use.

Table 5 leverages these results to provide more precise estimates of how the change in the productivity of each skill accounts for the overall within-occupation shift in skill use. It also compares the predictions of the model with the data.

Not surprisingly, given the imprecision of the skill-growth estimates for the first period, the model predictions are also imprecise. Within-occupation growth of abstract-skill use is predicted accurately, but we are less successful with the other two skills.

In contrast, the model captures the direction and rough magnitude of the within-occupation shifts between 1971 and 1983. While it somwhat underestimates each of these shifts, it is directionally correct, and the magnitudes are reasonably close.

The 1990-2006 period experienced large within-occupation changes and exhibited large differences in skill growth. Consequently, it is the period for which we had the greatest expectation of matching the within-occupation shifts. In fact, the model does a good job of capturing these large shifts, coming within 20% in each case.

The results for the final period are somewhat disappointing. The model does not do a good job of predicting the small within-occupation change in skill use and over-predicts the change in abstract-skill use, although it is almost perfect regarding manual-skill use.

We note that the predicted relative magnitudes of the within-occupation changes come

[20] In contrast $A_i S_i$, the 'effective' amount of skill i supplied by the worker, must increase with A_i.

from Table 3. The fact that we predict much larger within-occupation changes for the 1990-2006 period than for the other periods, is not trivial and provides support for the model.

5 Summary and conclusion

We make two contributions. First, at a purely empirical level, we provide new evidence on changes in skill use over a very long period. Second, we develop a simple model that reconciles or combines two approaches to technological change, the SBTC and task-based literatures, by modeling technological change as increasing the productivity of individual skills rather than, for example, college-educated workers. While our model also allows us to account for technological change that replaces occupations, we focus on detecting changes in skill productivity; we capture changing demand for occupations only through changes in industry demand.

We use the insights from the model to measure the pattern of skill-productivity growth needed to explain the employment shifts that we observe. For the 1960s, we find that differences in skill productivity growth account for very little of the employment changes that we observe. In contrast, they account for a significant component in the remaining periods, even after controlling for industry.

Our empirical results suggest that if a skill's productivity increases, use of that skill within an occupation generally decreases. Thus, skills generally are A-substitutes for themselves. Abstract and routine skills are A-complements, as are abstract and manual skills. In contrast, routine and manual skills are A-substitutes. We can understand the evolution of skill use both across and within-occupations as heavily influenced by the rapid productivity growth of routine skill and slow productivity growth of abstract skill. These two factors explain not only the observed across and within-occupation shifts in abstract and routine-skill use, but also in manual-skill use.

We hope and believe that we have demonstrated that our simple model provides a useful framework for understanding changes in skill use both between and within occupations. Obviously, readers must make that judgment for themselves.

References

Daron Acemoglu and David H Autor. Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of Labor Economics*, volume 4, pages 1043–1171. Elsevier, 2011.

Daron Acemoglu and Veronica Guerrieri. Capital deepening and nonbalanced economic growth. *Journal of Political Economy*, 116(3):467–498, 2008.

Daron Acemoglu and Pascual Restrepo. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542, 2018.

Daron Acemoglu and Pascual Restrepo. Automation and new tasks: How technology displaces and reinstates labor. *Journal of economic perspectives*, 33(2):3–30, 2019.

Enghin Atalay, Phai Phongthiengtham, Sebastian Sotelo, and Daniel Tannenbaum. The evolution of work in the United States. *American Economic Journal: Applied Economics*, 12(2):1–34, 2020.

David H Autor and David Dorn. The growth of low-skill service jobs and the polarization of the US labor market. *American Economic Review*, 103(5):1553–97, 2013.

David H Autor and Brendan Price. The changing task composition of the US labor market: an update of Autor, Levy, and Murnane (2003). Mimeo, 2013.

David H Autor, Frank Levy, and Richard J Murnane. The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, 118(4):1279–1333, 2003.

Zsófia L Bárány and Christian Siegel. Biased technological change and employment reallocation. *Labour Economics*, 67:101930, 2020.

William J Baumol. Macroeconomics of unbalanced growth: the anatomy of urban crisis. *American Economic Review*, 57(3):415–426, 1967.

Eli Berman, John Bound, and Zvi Griliches. Changes in the demand for skilled labor within US manufacturing: evidence from the annual survey of manufactures. *The Quarterly Journal of Economics*, 109(2):367–397, 1994.

Eli Berman, John Bound, and Stephen Machin. Implications of skill-biased technological change: international evidence. *The Quarterly Journal of Economics*, 113(4):1245–1279, 1998.

Sandra E Black and Alexandra Spitz-Oener. Explaining women's success: technological change and the skill content of women's work. *Review of Economics and Statistics*, 92(1): 187–194, 2010.

Timothy N Bond and Kevin Lang. The evolution of the Black-White test score gap in grades K–3: the fragility of results. *Review of Economics and Statistics*, 95(5):1468–1479, 2013.

Timothy N Bond and Kevin Lang. The sad truth about happiness scales. *Journal of Political Economy*, 127(4):1629–1640, 2019.

Costas Cavounidis and Kevin Lang. Ben-Porath meets Lazear: microfoundations for dynamic skill formation. *Journal of Political Economy*, 128(4):1405–1435, 2020.

Costas Cavounidis, Vittoria Dicandia, Kevin Lang, and Raghav Malhotra. The evolution of skill use within and between jobs. *Working Paper*, No. w29302, 2021.

Davide Consoli, Giovanni Marin, Francesco Rentocchini, and Francesco Vona. Routinization, within-occupation task changes and long-run employment dynamics. *Research Policy*, 52 (1):104658, 2023.

David J Deming and Kadeem Noray. Earnings dynamics, changing job skills, and STEM careers. *The Quarterly Journal of Economics*, 135(4):1965–2005, 2020.

William T Dickens. Error components in grouped data: is it ever worth weighting? *Review of Economics and Statistics*, 72(2):328–333, 1990.

David Dorn. *Essays on inequality, spatial interaction, and the demand for skills*. PhD thesis, University of St. Gallen, 2009.

Sergio Firpo, Nicole M Fortin, and Thomas Lemieux. Occupational tasks and changes in the wage structure. *IZA discussion paper*, page No. 5542, 2011.

Sarah Flood, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, Daniel Backman, Annie Chen, Grace Cooper, Stephanie Richards, Megan Schouweiler, and Michael Westberry. Integrated Public Use Microdata Series, Current Population Survey: Version 11.0 [dataset]. IPUMS, Minneapolis, MN, 2023. DOI : 10.18128/D030.V11.0.

Christina Gathmann and Uta Schönberg. How general is human capital? a task-based approach. *Journal of Labor Economics*, 28(1):1–49, 2010.

Maarten Goos, Alan Manning, and Anna Salomons. Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review*, 104(8):2509–26, 2014.

Charles I Jones. Intermediate goods and weak links in the theory of economic development. *American Economic Journal: Macroeconomics*, 3(2):1–28, 2011.

Chinhui Juhn. Wage inequality and demand for skill: evidence from five decades. *ILR Review*, 52(3):424–443, 1999.

Lawrence F Katz and Kevin M Murphy. Changes in relative wages, 1963–1987: supply and demand factors. *The Quarterly Journal of Economics*, 107(1):35–78, 1992.

Leonid Kogan, Dimitris Papanikolaou, Lawrence D. W Schmidt, and Bryan Seegmiller. Technology, vintage-specific human capital, and labor displacement: Evidence from linking patents with occupations. Working Paper 29552, National Bureau of Economic Research, 2021.

Edward P. Lazear. Firmâspecific human capital: A skillâweights approach. *Journal of Political Economy*, 117(5):914–940, 2009.

National Academy of Sciences. Committee on Occupational Classification and Analysis. Dictionary of Occupational Titles (DOT): Part I - Current Population Survey, April 1971, Augmented With DOT Characteristics and Dictionary of Occupational Titles (DOT): Part II - Fourth Edition Dictionary of DOT Scores for 1970 Census Categories. ICPSR7845 - v2, Inter - University Consortium for Political and Social Research, Ann Arbor , MI, 1977. DOI : 10.3886/ICPSR07845.v2.

L Rachel Ngai and Christopher A Pissarides. Structural change in a multisector model of growth. *American Economic Review*, 97(1):429–443, 2007.

Gary Solon, Steven J Haider, and Jeffrey M Wooldridge. What are we weighting for? *Journal of Human Resources*, 50(2):301–316, 2015.

Alexandra Spitz-Oener. Technical change, job tasks, and rising educational demands: Looking outside the wage structure. *Journal of labor economics*, 24(2):235–270, 2006.

Matthew Sobek Daniel Backman Annie Chen Grace Cooper Stephanie Richards Renae Rodgers Steven Ruggles, Sarah Flood and Megan Schouweiler. IPUMS USA: Version 15.0 [dataset]. IPUMS, Minneapolis, MN, 2024. DOI : 10.18128/D010.V15.0.

United States Department of Labor. *Dictionary of Occupational Titles (DOT): Third Edition*, volume 1 and 2. US Government Printing Office, Washington, DC, 1965.

United States Department of Labor, Employment and Training Administration. *O*NET Database*.

United States Department of Labor. United States Employment Service, and the North Carolina Occupational Analysis Field Center. Dictionary of Occupational Titles (DOT): Revised Fourth Edition. ICPSR6100 - v1, Inter - University Consortium for Political and Social Research, Ann Arbor , MI, 1991. DOI : 10.3886/ICPSR06100.v1.

Shintaro Yamaguchi. Tasks and heterogeneous human capital. *Journal of Labor Economics*, 30(1):1–53, 2012.

Table 1: Skills use levels by year

	1959	1971	1983	1995	2006	2016
Abstract skills						
	0.335	0.365	0.418	0.389	0.434	0.451
	(0.193)	(0.232)	(0.240)	(0.150)	(0.095)	(0.107)
Routine skills						
	0.539	0.494	0.467	0.413	0.397	0.398
	(0.152)	(0.215)	(0.200)	(0.094)	(0.059)	(0.055)
Manual skills						
	0.126	0.141	0.115	0.198	0.169	0.151
	(0.095)	(0.156)	(0.139)	(0.101)	(0.113)	(0.115)

Notes: Estimates use the occupation distributions from the 1960 Census, the March 1970-72, 1982-84, and 1994-96 Current Population Surveys, the 2005-2007 3-Year ACS/PRCS, the 2015-2017 ACS. The skills used in each occupation are taken from the third, fourth, and revised fourth editions of the *Dictionary of Occupational Titles* and fromthe 3.0, 12.0, 24.4 versions of the O*NET. *DOT* occupations are aggregated to census occupations using the April 1971 Current Population Survey. O*NET occupations are aggregated using the Occupational Employment and Wage Statistics (OEWS). The data for 1959-1983 and 1995-2016 are not strictly comparable. Standard deviations across occupations are provided in parentheses.

Table 2: Within- and across-occupation components

	1959-1971	1971-1983	1983-1995	1995-2006	2006-2016
Abstract skills					
Δ within	0.011	0.018	-0.042	0.042	0.013
Δ across	0.019	0.035	0.013	0.003	0.004
Routine skills					
Δ within	-0.028	-0.008	-0.046	-0.015	0.002
Δ across	-0.017	-0.019	-0.008	-0.001	-0.002
Manual skills					
Δ within	0.018	-0.010	0.088	-0.027	-0.015
Δ across	-0.002	-0.016	-0.005	-0.002	-0.002

Notes: This table decomposes the change in the use of each of three skills into the change that would have been observed if the occupation distribution had been the same at the end of the period as at the beginning of the period (Δ within) and what would have been observed if the skill use were always the skill use at the end of the period but the occupation distribution had changed. Estimates use the occupation distributions from the 1960 Census, the March 1970-72, 1982-84, and 1994-96 Current Population Surveys, the 2005-2007 3-Year American Community Survey/PRCS, the 2015-2017 American Community Survey. The skills used in each occupation come from the decennial censuses, and from the American Community Survey samples for the last two periods. DOT occupations are aggregated to census occupations using the April 1971 Current Population Survey. O*NET occupations are aggregated using the Occupational Employment and Wage Statistics (OEWS).

Table 3: Skill Productivity Growth Relative to Average

	(1) 1960-1970	(2) 1970-1980	(3) 1990-2006	(4) 2006-2016
Abstract	0.250	0.353	0.940	0.680
	(0.167)	(0.130)	(0.193)	(0.119)
Routine	-0.279	-0.261	-1.105	-0.767
	(0.225)	(0.148)	(0.324)	(0.163)
Manual	0.029	-0.091	0.164	0.087
	(0.238)	(0.177)	(0.266)	(0.091)
r2	0.17	0.12	0.18	0.15
N	4867	6940	8464	8478
proportion due to skills	0.08	0.15	0.21	0.24
p(all skill coefs=0)	0.231	0.013	0.000	0.000
p(routine=manual)	0.477	0.571	0.024	0.000

Notes: Standard errors in parentheses, clustered at the occupation level. Estimates are transformed from regression of change in log employment in log employment in an occupation/industry cell on average skill (Routine, Manual) use in that cell over the period (equation (22) in the text) and imposing that the mean deviation from mean skill growth for all four skills is 0. Proportion due to skills is the proportion of the R-squared attributable to the three skills in the regression using the Shapley-Owen decomposition.

Table 4: Derivatives of Skill Use with Respect to Skill Productivity

	Skill Used		
$\Delta\ lnA_i$	Abstract	Routine	Manual
Abstract	-0.024		
Routine	0.013	-0.0026	
Manual	0.011	-0.0104	-0.0006

Notes: Each cell shows the derivative of the average use of the column skill with respect to a change in the relative productivity of the row skill. Estimates are up to a factor of proportionality of $\frac{-\varepsilon}{1-\varepsilon}$ (which is strictly between 0 and 1). The estimates are derived from combining changes in skill use across time with estimates of relative productivity growth from Table 3. See equation (32) in the text for the precise formulation. See the text for more detail.

Table 5: Decomposition of Within-Occupation Changes in Skill Use

	1960-1971 Predicted Skill-Use Change		
Source of Change	Abstract	Routine	Manual
Abstract	0.006	-0.003	-0.003
Routine	0.004	-0.001	-0.003
Manual	0.000	0.000	0.000
Total Predicted	0.009	-0.004	-0.006
Data	0.011	-0.028	0.018
	1971-1983 Predicted Skill-Use Change		
Source of Change	Abstract	Routine	Manual
Abstract	0.008	-0.005	-0.004
Routine	0.003	-0.001	-0.003
Manual	0.001	-0.001	0.000
Total Predicted	0.013	-0.006	-0.007
Data	0.018	-0.008	-0.010
	1995-2006 Predicted Skill-Use Change		
Source of Change	Abstract	Routine	Manual
Abstract	0.023	-0.012	-0.010
Routine	0.014	-0.003	-0.011
Manual	-0.002	0.002	0.000
Total Predicted	0.035	-0.013	-0.022
Data	0.042	-0.015	-0.027
	2006-2016 Predicted Skill-Use Change		
Source of Change	Abstract	Routine	Manual
Abstract	0.016	-0.009	-0.007
Routine	0.010	-0.002	-0.008
Manual	0.001	0.001	0.000
Total Predicted	0.027	-0.010	-0.015
Data	0.013	0.002	-0.015

Notes: Each entry is the predicted change in the within-occupation use of the column skill due to changes in the productivity of the row skill according to equation (32) in the text and using the values from Tables 3 and 4. Total predicted is the sum of the four values above. The predictions can be compared with the within changes reported in Table 2 and repeated in the line labelled Data.

Table 1A: Skill Productivity Growth Relative to Average - non transformed coefficients

	(1) 1960-1970	(2) 1970-1980	(3) 1990-2006	(4) 2006-2016
Routine	-0.529	-0.614	-2.045	-1.447
	(0.317)	(0.216)	(0.462)	(0.270)
Manual	-0.221	-0.444	-0.776	-0.593
	(0.344)	(0.272)	(0.333)	(0.134)
r2	0.17	0.12	0.18	0.15
N	4867	6940	8464	8478
proportion due to skills	0.08	0.15	0.21	0.24
p(all skill coefs=0)	0.231	0.013	0.000	0.000
p(routine=manual)	0.477	0.571	0.024	0.000

Notes: Standard errors in parentheses, clustered at the occupation level. Estimates from regression of change in log employment in an occupation/industry cell on average skill (Routine, Manual) use in that cell over the period (equation (22) in the text) before the transformation shown in Table 3.

Table 2A: Skill Productivity Growth Relative to Average - including wages

	(1) 1960-1970	(2) 1970-1980	(3) 1990-2006	(4) 2006-2016
Abstract	0.252	0.358	0.938	0.680
	(0.167)	(0.130)	(0.196)	(0.119)
Routine	-0.287	-0.264	-1.104	-0.768
	(0.226)	(0.148)	(0.323)	(0.163)
Manual	0.036	-0.094	0.166	0.088
	(0.239)	(0.177)	(0.270)	(0.091)
% Change mean wage 60	-0.058			
	(0.037)			
% Change mean wage 70		0.010		
		(0.005)		
% Change mean wage 90			0.007	
			(0.050)	
% Change mean wage 06				0.022
				(0.011)
r2	0.17	0.13	0.18	0.15
N	4867	6910	8464	8478
proportion due to skills	0.08	0.15	0.21	0.23
p(all skill coefs=0)	0.221	0.011	0.000	0.000
p(routine=manual)	0.458	0.572	0.024	0.000

Notes: Standard errors in parentheses, clustered at the occupation level. Nominal wages are converted to 1999 Dollars with CPI99 provided by IPUMS - USA. Estimates are transformed from regression of change in log employment in an occupation/industry cell on average skill (Routine, Manual) use in that cell over the period (equation (22) in the text) and imposing that the mean deviation from mean skill growth for all four skills is 0. Proportion due to skills is the proportion of the R-squared attributable to the three skills in the regression using the Shapley-Owen decomposition.