Tasks and Technology in the United States 1880-2000*

Guy Michaels
London School of Economics and CEP†
Ferdinand Rauch
University of Oxford and CEP‡
Stephen J. Redding
Princeton University and CEP§

September 13, 2016

Abstract

We provide theory and evidence on changes in task inputs in the United States from 1880-2000. We combine a Roy model of worker selection across occupations with a new methodology for measuring individual production tasks performed by workers within occupations. We show that the recently-documented rise in non-routine tasks and decline in manual tasks extends much further back than hitherto thought to the late-nineteenth century. We reveal substantial heterogeneity within these broad categories of tasks, with those involving the formation of ideas increasing by up to twice the growth for non-routine tasks as a whole, and those involving the manipulation of inorganic matter decreasing by nearly twice the overall decline for manual tasks. We establish that these changes in task inputs are explained by new technologies (in particular office and computing machinery) and are larger in urban than in rural areas (implying a transformation in the nature of agglomeration). We show that changes in the wage premia for tasks can account for a substantial proportion of the decline in wage inequality from 1880-1940, the rise in wage inequality from 1940-2000, and the larger rise in wage inequality in urban areas than in rural areas, even after controlling for observed worker characteristics.

KEYWORDS: Tasks, Technology, Urbanization, Wage Inequality.

JEL CLASSIFICATION: N92, O18, R12

*This is a much revised version of a paper that previously circulated under the title "Task Specialization in U.S. Cities from 1880-2000" (*NBER Working Paper*, 18715). We are grateful to the Centre for Economic Performance and Princeton University for research support. We thank Gilles Duranton, William D. Caughlin, Rowena Gray, George Kupczak and Matt Turner for sharing data. We are grateful to the editor, three anonymous referees, Don Davis, Gilles Duranton, Gene Grossman, Jessie Handbury, Gordon Hanson, Vernon Henderson, Rick Hornbeck, Jeffrey Lin, Marco Manacorda, Alan Manning, Nathan Nunn, Gianmarco Ottaviano, Steve Pischke, Nancy Qian, Esteban Rossi-Hansberg, Yona Rubinstein, Will Strange, Tony Venables and seminar and conference participants at CEPR-CURE, Copenhagen, the London School of Economics, MIT, NARSC, Oxford, Sorbonne, Sussex and UCSD for helpful comments and suggestions. The usual disclaimer applies.

†CEP, LSE, Houghton Street, WC2A 2AE, London, UK. Email: g.michaels@lse.ac.uk. Tel: +44(0)20-7852-3518.

‡Department of Economics, University of Oxford, Manor Road, Oxford, OX1 3UQ. Email: ferdinand.rauch@economics.ox.ac.uk. Tel: +44(0)77214-88037.

§Fisher Hall, Princeton, NJ 08540. Email: reddings@princeton.edu. Tel: +1(609) 258-4016.

1 Introduction

Examples from history abound where new technologies have substituted for the production tasks performed by some workers and complemented those performed by others.[1] Motivated in part by this and other evidence, recent theoretical and empirical research has emphasized a distinction between tasks and skills.[2] A *task* is a unit of work activity that produces output (goods and services), whereas a *skill* is a worker's endowment of capabilities for performing various tasks. New technologies typically complement or substitute for particular tasks in a pattern that can be poorly summarized by aggregate measures of worker skills such as college degree or equivalent. To study these patterns of complementarity and substitutability, we develop a new framework for measuring inputs of *individual* production tasks. We use this framework to provide the first evidence on changes in inputs of these individual tasks in the U.S. economy over more than a century. We find that an early revolution in information and communication technology (ICT) in the late-nineteenth century changed task inputs in a similar way to the computer revolution in the late-twentieth century. We show that electrical machinery, like ICT, is complementary to non-routine tasks, whereas transport machinery is complementary to manual tasks. We find a transformation in the nature of agglomeration, where the physical tasks that used to be concentrated in urban areas have been replaced by analytical and interactive tasks. We show that task inputs and wage premia account for substantial proportions of the decline in wage inequality from 1880-1940, the rise in wage inequality from 1940-2000, and the larger changes in wage inequality in urban areas than in rural areas over time. We show that these findings continue to hold even after controlling for observed worker characteristics (such as education, age, gender and ethnicity), implying that they are not captured by more aggregated measures of human capital or skills.

Our main methodological contribution is to develop an empirical framework for measuring individual production tasks that uses the verbs from occupational descriptions and their meanings. Using this new methodology, we are able to track the production tasks performed by workers at a much higher resolution than has hitherto been possible. We find far larger changes in task inputs using our more disaggregated measures than using the non-routine, routine and manual measures considered in existing research (up to four times larger using comparable percentile scores). Among non-routine tasks, we find that those involving the formation of ideas (e.g. Analyze, Confer, Evaluate) increased the most. Among manual tasks, we show that those involving the manipulation of physical matter (e.g. Deliver, Grind, Weld) decreased the most. Our main substantive contribution is to use this framework to provide the first evidence on

[1] One of the most famous examples from history is associated with the Luddites: 19th-century English textile workers who protested against newly-developed labour-economizing technologies from 1811-6. See for example Mokyr (1992).

[2] See, in particular, the recent surveys by Acemoglu and Autor (2011) and Autor (2013), which contrast a "task-based" approach with the "canonical model" of the labor market in terms of skilled and unskilled labor.

task inputs in the U.S. economy over a long historical time period from 1880-2000. We use this long-term perspective to show that the rise in non-routine tasks and decline in manual tasks documented for recent decades extends much further back than previously thought to the late-nineteenth century.

We examine the role of new technologies in complementing and substituting for tasks by constructing measures of industry technology use from the input-output matrix. We consider four technologies that experienced substantial innovation over our sample period: (i) office and computing machinery; (ii) electrical machinery; (iii) transport machinery; (iv) all machinery (computing, electrical, transport and other machinery). Innovations in office and computing machinery include not only computers and modern communication equipment, but also an earlier revolution in information and communication technology (ICT) around the turn of the twentieth century. These earlier innovations comprised telephones, typewriters, and other improvements in producing, communicating, storing, and retrieving information. Innovations in electrical machinery capture the dissemination of electrically-powered capital goods. Innovations in transport machinery include the automobile and air travel, which substantially reduced physical transportation costs during our sample period (see Glaeser and Kohlhase 2003).

Of these four technologies, we find the greatest effects for office and computing machinery. Industries making more intensive use of this technology experience the largest increases in non-routine tasks and the largest reductions in manual tasks. We find that the positive impact of office and computing machinery on inputs of non-routine tasks starts in the last two decades of the nineteenth century during the early ICT revolution (with the rise of telephones and typewriters), but accelerates in the latter part of the twentieth century (following the diffusion of the computer). These results capture both the direct effect of office and computing machinery (e.g. typewriters substitute for handwriting) as well as its indirect effect in facilitating changes in the organization of production (e.g. typewriters improve the recording and processing of information, which permits larger-scale modes of production activity). We find that the individual tasks most complementary with office and computing machinery are intellectual tasks, such as Analyze, Design, Program and Review, while the tasks for which office and computing machinery substitutes the most are physical tasks such as Assemble, Collect and Ticket. In contrast, we find a quite different pattern of results for transport machinery. Industries that use transport equipment intensively saw shifts in task use towards manual tasks and away from non-routine tasks.

Whereas most prior research on tasks and the labor market has been concerned with the economy as a whole, we make a further substantive contribution in applying this approach to the organization of economic activity in urban versus rural areas. We find a reversal in the nature of agglomeration over time. In 1880, urban workers performed *less* non-routine tasks than rural workers. In contrast, in 2000, urban workers performed *more* non-routine tasks than rural workers. These changes are substantial: we find that the differential increase in task inputs between urban and rural areas is large relative to the overall increase for either urban or rural areas alone. In 1880, the individual tasks most concentrated in urban areas were concerned with the manipulation of physical matter, such as Thread, Stretch, Ravel and Sew. In contrast, by the year 2000, the individual tasks most clustered in urban locations related to analytical and interactive activity, such as Analyze, Confer, Determine and Review. We find these changes in urban-rural task specialization even after controlling for observed worker characteristics (including schooling), confirming that they are not captured by more aggregated measures of human capital or skills.

We show that these changes in task inputs are consequential for understanding the evolution of wage inequality. We combine the measures of labor income by occupation for 1880 from Abramitzky, Boustan, and Eriksson (2012, 2014) and Preston and Haines (1991) with the individual-level data on labor income that is available in the Population Census from 1940 onwards. Between 1880 and 1940, we find a decline in wage inequality across occupations, with a reduction in the mass of workers at high and low wages, and an increase in the mass of workers at medium wages. In contrast, between 1940 and 2000, we find increased polarization and wage inequality, with a reduction in the mass of workers at medium wages, and an increase in the mass of workers at high wages. We show that changes in task premia explain much of the observed changes in wage inequality, even after controlling for observed worker characteristics. To tighten this connection, we use variation between urban and rural areas. We show that the increase in the dispersion of wages between 1940 and 2000 is larger in urban areas than in rural areas and that this larger change in wage inequality is mainly explained by changes in task premia.

To guide our empirical analysis, we develop a Roy model of worker selection across sectors and occupations that require the performance of heterogeneous tasks. The model highlights employment shares and average wages as sufficient statistics for the impact of new technologies on the economy. The model emphasizes two mechanisms through which these new technologies affect the economy: the average effectiveness of workers in performing tasks (worker productivity in an occupation) and/or the rate of return to human capital accumulation in an occupation. We show how the model can be quantified to recover measures of task effectiveness and the rate of return to human capital accumulation from the observed data. We also show how it can be extended to incorporate multiple groups of workers that differ in observed characteristics (e.g. skilled and unskilled) and multiple locations (e.g. urban and rural areas). Guided by these predictions, our empirical work examines the extent to which employment shares and average wages have changed systematically towards occupations performing certain types of tasks; the extent to which these changes are related to direct measures of new technologies; and the extent to which these changes differ between urban and rural areas.

Our paper is related to a number of existing literatures. First, a growing body of research has argued that the canonical model of the labor market in terms of skilled and unskilled labor is not well suited to explaining several contemporary labor market phenomena. An alternative task-based approach has been pioneered by Autor, Levy, and Murnane (2003) (henceforth ALM) using numerical scores from the *Dictionary of Occupational Titles (DOTs)* that summarize job requirements as measured by the Department of Labor. Following ALM, this literature typically distinguishes between non-routine, routine and manual tasks. In some cases, non-routine is further disaggregated into analytic and interactive, and routine is further broken out into cognitive and non-cognitive. Among the recent labor market phenomena explained by this approach are significant declines in real wages of low-skill workers (e.g. Acemoglu and Autor 2011); non-monotone changes in wages at different parts of the earnings distribution during different decades (e.g. Autor, Katz, and Kearney 2008, henceforth AKK and Autor and Dorn 2013); job polarization with broad-based increases in employment in high-skill and low-skill occupations relative to middle-skilled occupations (e.g. Goos and Manning 2007 and Goos, Salomons, and Manning 2014); and rapid diffusion of new technologies that directly substitute capital for labor in tasks previously performed by moderately-skilled workers (e.g. ALM and Autor and Dorn 2013). Our main contribution relative to this literature is to develop a new framework for measuring *individual* tasks at a far higher resolution than existing approaches, to apply this framework over a much longer time period than previously considered, and to use this framework to analyze the changing nature of agglomeration over time.

The canonical model of the labor market in terms of skilled and unskilled labor typically imposes the assumption that technological change is skill-biased. In contrast, a recent theoretical literature following Acemoglu (1998) and Acemoglu (2002) has argued that the direction of technological change is endogenous. Therefore the extent to which new technologies complement or substitute for skills or tasks can change over time. In the labor literature, Autor, Katz, and Krueger (1998) argue that there was an acceleration in the skill-bias of technical change in the 1980s and 1990s. In the historical literature, several studies argue that technical change often replaced–rather than complemented–skilled artisans in the nineteenth-century, including Hounshell (1985), James and Skinner (1985) and Mokyr (1992). However, there remains substantial debate about the extent to which this was the case. In their classic study of the race between technology and skills, Goldin and Katz (2008) present evidence that manufacturing technologies were skill complementary in the early-twentieth century, but may have been skill substituting prior to that time.[3] In subsequent work, Katz and Margo (2014) report some evidence of de-skilling in manufacturing during the nineteenth-century, but find a reallocation of employment towards high-skill jobs for the aggregate economy as a whole.[4] We use our new methodology to provide the first quantitative evidence on pat-

[3]Using data from the early-twentieth century, Gray (2013) finds that electrification led to a polarization of the employment distribution, increasing the demand for non-routine and routine cognitive tasks, while simultaneously reducing relative demand for the non-routine manual jobs which comprised the middle of the skill distribution.

[4]In a study of the merchant shipping industry in the late-nineteenth and early-twentieth centuries, Chin, Juhn, and Thompson

terns of complementary and substitutability between individual tasks and new technologies over our long historical time period.

Our analysis also relates to the theoretical and empirical literature on agglomeration. Traditionally, this literature has emphasized the costs of moving goods and people across space. In the new economic geography literature, agglomeration is explained by consumer love of variety, increasing returns to scale and transport costs (e.g. Fujita, Krugman, and Venables 1999). In canonical models of urban economies, agglomeration is driven by economies of scale in producing a single final good, while commuting costs provide a dispersion force (e.g. Alonso 1964, Muth 1969 and Mills 1967). However, other mechanisms for agglomeration have been considered, including human capital externalities (e.g. Berry and Glaeser 2005, Moretti 2004 and Davis and Dingel 2013);[5] localization versus urbanization externalities (within versus between sector externalities as in Henderson 2003); the costs of exchanging ideas (e.g. Davis and Dingel 2012 and Gaspar and Glaeser 1998); the role of the size of the market for the division of labor (e.g. Duranton and Jayet 2011); structural transformation and the relocation of manufacturing from urban areas (e.g. Desmet and Rossi-Hansberg 2009); and specialization by sector versus by function (e.g. Brunelle 2013, Duranton and Puga 2005 and Rossi-Hansberg, Sarte, and Owens III 2009).[6] We use our new framework to provide the first systematic application of a task-based approach of the labor market to urban versus rural areas. We highlight a transformation in the nature of agglomeration, with a reversal in the types of tasks concentrated in urban versus rural areas over time. We show that changes in task wage premia account for much of the differential changes in wage inequality between urban and rural areas, even after conditioning on worker observables such as human capital or skills.[7]

The remainder of the paper is structured as follows. Section 2 introduces the model. Section 3 discusses the data. Section 4 introduces our new methodology for measuring the individual production tasks performed by workers within each occupation. Section 5 uses our new methodology to provide evidence on changes in task inputs over our long historical sample period and their implications for wage inequality. Section 6 shows that these changes in task inputs are driven in part by new technologies and examines

(2006) find that the adoption of the steam engine raised skill premia. Using data on manufacturing plants in the late-nineteenth century, Atack, Bateman, and Margo (2004) find that plant wages are decreasing in size, but are increasing in the use of steam power, which is consistent with technology-skill complementarity.

[5]Bacolod, Blum, and Strange (2009a,b) distinguish different types of human capital, where soft or tacit skills are argued to be more important in urban areas. Lin (2011) finds that new occupation codes are concentrated in cities. To the extent that new occupations use tacit knowledge intensively, this is consistent with the concentration of such soft skills in cities.

[6]The large literature on human capital in cities includes Combes, Duranton, Gobillon, and Roux (2012), Berry and Glaeser (2005), Glaeser, Ponzetto, and Toblo (2011), Glaeser and Saiz (2003), Glaeser and Resseger (2009), and Hendricks (2011) and is surveyed in Moretti (2004). The distinction between localization and urbanization externalities dates back to Jacobs (1969) and Henderson (1974). Research on the division of labor dates back to Adam Smith and includes Baumgardner (1988), Becker and Murphy (1992) and Duranton (1998). Specialization by sector versus by function (e.g. headquarters versus plants) is also explored in Ota and Fujita (1993), Glaeser and Kahn (2001), Helsley and Strange (2007) and Fujita and Tabuchi (1997).

[7]For evidence on more recent changes in wage inequality in cities, see also Baum-Snow and Pavan (2012, 2013) and Eeckhout, Pinheiro, and Schmidheiny (2014).

the individual production tasks that these new technologies complement or substitute for. Section 7 shows that these changes in task inputs differ between urban and rural areas and explain a substantial proportion of the differences in the evolution of wage inequality between urban and rural areas. Section 8 concludes.

2 Model

To guide our empirical analysis, we develop a simple Roy model, in which workers endogenously sort across occupations and sectors based on their comparative advantage.[8] We begin by developing a baseline version of the model in which workers are *ex ante* identical and the economy consists of a single location (e.g. the U.S.). We next demonstrate the robustness of our results to two extensions, one to allow for multiple worker types that differ *ex ante* in observable characteristics (e.g. gender, age and general schooling), and the other to allow for multiple locations (e.g. urban and rural areas), as in our data.[9]

We use the model to show how changes in technology affect employment shares and average wages in a setting where workers self-select across many sectors and occupations based on idiosyncratic realizations for ability. We highlight two mechanisms through which technology affects employment shares and wages: the average effectiveness of workers in performing tasks and/or the rate of return to human capital accumulation. Under our assumption of a Fréchet distribution for idiosyncratic ability, changes in the average rate of return to human capital accumulation only affect average wages, whereas changes in average task effectiveness affect both employment shares and average wages. Guided by these predictions, our empirical work examines the extent to which employment shares and average wages have changed systematically towards occupations performing certain types of tasks (measured using both numerical scores and our new methodology); the extent to which these changes are related to direct measures of new technologies (e.g. information and communication technologies); and the extent to which these changes differ between urban and rural areas.

The economy consists of a continuum of people ($\bar{L}$) who can choose to work in O possible occupations. Human capital for each occupation depends on raw worker ability and investments in human capital accumulation for that occupation. People choose an occupation based on the wage and cost of investing in occupational human capital. Occupational human capital is used to produce final goods in S sectors. We allow some occupations (e.g. managers) to be employed in most sectors, while other occupations (e.g. lathe operators) may be employed in only a few sectors.

[8]We extend the version of the Roy model in Hsieh, Hurst, Jones, and Klenow (2013) to incorporate multiple sectors and locations, as observed in our data. The classic treatments of the Roy model are Roy (1951) and Heckman and Honore (1990). See also Lagakos and Waugh (2013), Ahlfeldt, Redding, Sturm, and Wolf (2015) and Burstein, Morales, and Vogel (2015).

[9]A more detailed discussion of the model and the technical derivations of all expressions and results reported in this section are contained in the web appendix.

2.1 Preferences and Technology

A person i with consumption C_i and leisure time $1-\ell_i$ obtains utility:

$$U_i = C_i^{\beta}\left(1-\ell_i\right), \qquad \beta > 0, \tag{1}$$

where C_i is a consumption index; ℓ_i represents investments in human capital accumulation; and β parameterizes the tradeoff between consumption and the accumulation of human capital. The consumption index (C_i) is itself a Cobb-Douglas function of consumption of tradeable goods (C_{Mi}) and a non-tradeable good (C_{Ni}) that we interpret as housing:[10]

$$C_i = \left(\frac{C_{Mi}}{\alpha}\right)^{\alpha}\left(\frac{C_{Ni}}{1-\alpha}\right)^{1-\alpha}, \qquad 0 < \alpha < 1, \tag{2}$$

where housing is assumed to be in inelastic supply $\bar{N}$ and the presence of this non-traded good ensures a non-degenerate distribution of economic activity in the multi-region version of the model below.

The tradeables consumption index (C_{Mi}) is a constant elasticity of substitution (CES) function of consumption of a number of sectors (C_{Mis}) indexed by $s \in \{0, \ldots, S\}$:

$$C_{Mi} = \left[\sum_{s=1}^{S}\left(\zeta_s C_{Mis}\right)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}, \qquad \sum_{s=1}^{S}\zeta_s = 1, \tag{3}$$

where ζ_s controls the strength of relative preferences for sector s and σ is the elasticity of substitution between sectors. Output in each tradeable sector (Y_{Ms}) is a constant elasticity of substitution (CES) function of the human capital of workers from each occupation within that sector (H_{so}):

$$Y_{Ms} = \left[\sum_{o=1}^{O}\left(\xi_{so} H_{so}\right)^{\frac{\kappa-1}{\kappa}}\right]^{\frac{\kappa}{\kappa-1}}, \qquad \sum_{o=1}^{O}\xi_{so} = 1 \;\forall\; s, \tag{4}$$

where ξ_{so} controls the relative productivity of occupation o in sector s; this occupation o is not employed in sector s if $\xi_{so} = 0$; κ is the elasticity of substitution between occupations within sectors; and goods market clearing requires that output of each good equals the sum of all individuals' consumption of that good: $Y_{Ms} = \sum_i C_{Mis}$.[11]

Workers choose an occupation and acquire human capital for that occupation, such that each occupation corresponds to a separate labor market. Workers within each occupation o are mobile across sectors, which implies that the wage per effective unit of labor within that occupation (w_o) is the same across

[10] For empirical evidence in support of the constant housing expenditure share implied by this Cobb-Douglas functional form, see Davis and Ortalo-Magné (2011).

[11] We assume for simplicity that labor is the sole factor of production, but the analysis can be extended to incorporate other factors of production such as capital or land. We also assume for simplicity that κ takes the same value across sectors, but it is straightforward to allow this elasticity to differ across sectors.

sectors. Each worker i's choice of sector s within occupation o is determined by their idiosyncratic realizations for effective units of labor (ability z_{iso}) for each sector and occupation. Each worker i's choice of occupation o depends on these realizations for idiosyncratic ability, the wage per effective unit of labor for each occupation, and the rate of return to human capital investments for each occupation.

Each person i works one unit of time in her chosen occupation $o \in \{1, \ldots, O\}$. Another unit of time is divided between leisure $(1 - \ell_{io})$ and human capital accumulation (ℓ_{io}). The production function for human capital in occupation o is:

$$h_{io}(\ell_{io}) = \bar{h}_o \ell_{io}^{\phi_o}, \tag{5}$$

where $\bar{h}_o > 0$ captures the productivity of human capital investments; $\phi_o > 0$ determines the rate of return to human capital accumulation; and both parameters can differ across occupations o.

Human capital for each sector and occupation (H_{so}) equals the fraction of agents who choose that sector and occupation (λ_{so}) times average human capital conditional on choosing that sector and occupation times the measure of agents in the economy $(\bar{L})$:

$$H_{so} = \lambda_{so}\, \mathbb{E}\left[h_o z \mid \text{Person chooses } s \text{ and } o\right] \bar{L}. \tag{6}$$

where average human capital depends on both human capital accumulation (h_o) and ability (z).

Each person's income depends on the wage per effective unit of labor for her chosen occupation (w_o), her accumulated human capital for that occupation (h_{io}) and her idiosyncratic ability (z_{iso}) for her chosen sector and occupation:

$$\Omega_i = w_o h_{io} z_{iso} = w_o \bar{h}_o \ell_{io}^{\phi_o} z_{iso}. \tag{7}$$

The timing of decisions is as follows. First, each person i observes her realizations of idiosyncratic ability (z_{iso}) and chooses a sector s and occupation o, taking occupational wages (w_o) as given. Second, she chooses her optimal human capital investment in her chosen occupation (ℓ_{io}), given the trade-off between goods consumption and human capital accumulation in utility (1) and the technology for accumulating human capital (5). Third, she makes her optimal choices for overall goods consumption (C_{Mi}), consumption of housing (C_{Ni}), and goods consumption for each sector (C_{Mis}), given observed prices (P_{Ms}, P_N) and her income (Ω_i) in her chosen sector and occupation (7), as determined by wages, human capital investments and idiosyncratic ability.

An equilibrium in this economy is a set of allocations of consumption, production, human capital investments and choices of sector and occupation $\{C_i, C_{Mi}, C_{Ni}, C_{Mis}, Y_{Ms}, \ell_{io}, \lambda_{so}\}$ and a set of prices $\{P_M, P_N, P_{Ms}, w_o\}$, such that individuals choose consumption, human capital investments, sector and occupation to maximize utility; firms choose inputs of human capital to maximize profits; zero profits are made if a good is produced; and the markets for goods, labor and housing clear. We use the timing of decisions and structure of the model to solve for equilibrium recursively. First, we characterize equilibrium consumption and production as a function of human capital investments and choice of sector and occupation. Second, we determine optimal human capital investments as a function of choice of sector and occupation. Third, we solve for the optimal choice of sector and occupation.

2.2 Consumption Decisions

Given an individual's human capital investments and her choice of sector and occupation, the characterization of consumption decisions is straightforward. The Cobb-Douglas functional form (2) implies that each person allocates constant shares of income to consumption of goods and housing: $C_{Mi} = \alpha\Omega_i/P_M$ and $C_{Ni} = (1-\alpha)\Omega_i/P_N$. Using these results in (1), the utility function can be written in terms of income, the prices of goods consumption and housing, and investment in human capital accumulation:

$$U_i = \left(\frac{w_o \bar{h}_o \ell_{io}^{\phi_o} z_{iso}}{P_M^{\alpha} P_N^{1-\alpha}}\right)^{\beta} (1-\ell_{io}) . \tag{8}$$

Using the CES functional form of goods consumption (3) and the CES production technology (4), we can solve for the share of expenditure on goods consumption allocated to each sector, the share of the sectoral wage bill allocated to each occupation, and the occupation wage per effective unit of labor (w_o), as shown in the web appendix. We can also solve for the price of housing (P_N) and the price index for each sector (P_{Ms}), which determines the overall price index for goods consumption (P_M).

2.3 Human Capital Investments

Given a choice of sector s and occupation o, the occupational wage (w_o), realizations of idiosyncratic ability (z_{iso}), and prices of tradeable (P_M) and non-tradeable (P_N) goods, each individual i chooses her human capital investment (ℓ_{io}) to maximize her utility (8):

$$\max_{\ell_{io}} \left\{ \left(\frac{w_o \bar{h}_o \ell_{io}^{\phi_o} z_{iso}}{P_M^{\alpha} P_N^{1-\alpha}}\right)^{\beta} (1-\ell_{io}) \right\} . \tag{9}$$

The first-order condition to this problem yields the equilibrium human capital investment:

$$\ell_o^* = \frac{1}{1+\frac{1}{\beta\phi_o}}, \tag{10}$$

which only varies across occupations o and not across individuals i within occupations. Hence, from now onwards, we suppress the individual subscript i, unless otherwise indicated.

Equilibrium human capital investments depend solely on β (the tradeoff between consumption and the accumulation of human capital) and ϕ_o (the productivity of human capital accumulation in the worker's chosen occupation). Other forces do not affect these investments, because they have the same effect on the return and opportunity cost to human capital accumulation. The expression (10) highlights that a first key mechanism through which technological change can affect the economy is that it can change the productivity of human capital investments (ϕ_o) in some occupations relative to others. For example, computers may increase the ease of acquiring analytical and technical skills used in engineering occupations relative to manual skills used in laboring occupations.

2.4 Sector and Occupation Choice

Given occupational wages (w_o), realizations of idiosyncratic ability in each sector and occupation (z_{iso}), and prices of tradeable (P_M) and non-tradeable (P_N) goods, each individual chooses her sector and occupation to maximize her utility. We model individual ability following McFadden (1974) and Eaton and Kortum (2002). Each individual i draws ability for each sector s and occupation o (z_{iso}) from an independent Fréchet distribution:

$$F_{so}(z) = e^{-T_{so} z^{-\theta}}, \qquad \theta > 1. \tag{11}$$

The Fréchet scale parameter T_{so} determines average effective units of labor for workers in sector s and occupation o, which we refer to as the average effectiveness of workers in performing tasks in that sector and occupation. The Fréchet shape parameter θ determines the dispersion of effective units of labor across sectors and occupations. A reduction in θ corresponds to an increased dispersion of effective units of labor and greater scope for worker specialization according to comparative advantage across sectors and occupations.[12] In this specification (11), another mechanism through which technological change can affect the economy is that it can raise the effectiveness of workers in performing tasks in some sectors and occupations (T_{so}) relative to others. For example, on the one hand, computers may complement workers in performing design and simulation tasks in engineering occupations. On the other hand, computers may substitute for workers in performing routine calculations in clerical occupations. In our empirical work, we use the structure of the model to estimate the extent to which new technologies complement or substitute for different occupations by changing the average effectiveness of workers in performing different tasks (e.g. the formation of ideas versus the manipulation of the physical world).[13]

[12] Although we assume that ability is drawn independently for each sector and occupation, the parameter T_{so} induces a correlation in ability among workers within the same sector and occupation. While we focus on the independent Fréchet distribution for simplicity, it is straightforward to instead consider the multivariate Fréchet distribution, which allows for correlation in the ability draws of individual workers across sectors and occupations.

[13] Our theoretical framework models technological change as determining the relative effectiveness of workers in performing tasks in different sectors and occupations. Alternatively, technological change could be modeled as embodied in physical capital and machines. In both cases, new technologies can either complement or substitute for workers in particular occupations. Our approach enables us to tractably model the effects of technological change in a setting with many sectors and occupations in the context of a standard Roy model.

From utility (8), the following transformation of utility is linear in worker ability:

$$v = U^{1/\beta} = \bar{w}_o z, \qquad \bar{w}_o = \frac{w_o \left(1 - \ell_o\right)^{1/\beta} \ell_o^{\phi_o} \bar{h}_o}{P_M^{\alpha} P_N^{1-\alpha}}. \tag{12}$$

Using this monotonic relationship between utility and worker ability, the distribution of utility across workers within each sector s and occupation o inherits a Fréchet distribution:

$$F_{so}\left(v\right) = e^{-\Phi_{so} v^{-\theta}}, \qquad \Phi_{so} = T_{so} \bar{w}_o^{\theta}. \tag{13}$$

Each worker chooses their occupation and sector to maximize their utility. Note that the maximum of Fréchet distributed random variables also has a Fréchet distribution. Therefore the distribution of utility across all sectors and occupations is given by:

$$F\left(v\right) = e^{-\Phi v^{-\theta}}, \qquad \Phi = \sum_{s=1}^{S} \sum_{o=1}^{O} T_{so} \bar{w}_o^{\theta}, \tag{14}$$

where the Fréchet functional form implies that the distribution of utility conditional on choosing a sector and occupation is the same for each sector and occupation pair and equal to the distribution of utility across all sectors and occupations (14).

Aggregating optimal choices of sector and occupation across people, we arrive at our first key result for equilibrium worker sorting across sectors and occupations.

Proposition 1 ***(Sector and Occupation Choice)*** *Let λ_{so} denote the fraction of people who choose to work in sector s and occupation o. Let λ_o denote the fraction of people who choose to work in occupation o. Aggregating across people, the model yields the following sufficient statistics for the fractions of people choosing to work in each sector and occupation (Ψ_{so}) and in each occupation (Ψ_o):*

$$\lambda_{so} = \frac{\Psi_{so}}{\Psi}, \qquad \lambda_o = \frac{\Psi_o}{\Psi}, \tag{15}$$

$$\Psi_o = \sum_{s=1}^{S} \Psi_{so}, \qquad \Psi = \sum_{s=1}^{S} \sum_{o=1}^{O} \Psi_{so}, \qquad \Psi_{so} = T_{so} w_o^{\theta} \left(1 - \ell_o\right)^{\theta/\beta} \ell_o^{\theta\phi_o} \bar{h}_o^{\theta}.$$

Proof. See Appendix. ■

Therefore the fraction of workers choosing to work in each sector and occupation depends on average effective units of labor (as determined by T_{so}), the wage in each occupation (w_o), and equilibrium human capital investments (as determined by ℓ_o, which in turn depends solely on β and ϕ_o). Note that the terms in the tradeables consumption price (P_M) and the non-tradeables price (P_N) in Φ_{so} in (12)-(13) have cancelled from the choice probabilities λ_{so} in (15), because they are common across sectors and occupations, and hence do not affect the choice of sector and occupation.

Thus the model suggests that employment shares are one of the key endogenous variables of interest in our empirical work. The model also highlights the role of new technologies in influencing these employment shares through changing the effectiveness of workers in performing the tasks in different sectors and occupations. Changes in technology have both direct effects on employment shares (though T_{so}) and indirect effects (through general equilibrium effects via occupation wages w_o). Totally differentiating the sector-occupation choice probabilities (15), holding constant the rate of return to human capital investments (ϕ_o and hence ℓ_o), we have:

$$\frac{d\lambda_{so}}{\lambda_{so}} = \frac{dT_{so}}{T_{so}} - \sum_{k=1}^{S}\sum_{m=1}^{O}\frac{dT_{km}}{T_{km}}\lambda_{km} + \theta\frac{dw_o}{w_o} - \sum_{k=1}^{S}\sum_{m=1}^{O}\theta\frac{dw_m}{w_m}\lambda_{km}. \tag{16}$$

Evaluating these total derivatives holding occupational wages constant at their values in the initial equilibrium (setting $dw_o/w_o = 0$ for all occupations o), the direct effect of an increase in the effectiveness of workers in performing the tasks in a sector and occupation (T_{so}) is to raise the share of employment (λ_{so}) in that sector and occupation (since $0 < \lambda_{so} < 1$) and reduce the share of employment in all other sectors and occupations. Therefore, to the extent that new technologies complement workers in performing tasks in a sector and occupation (higher T_{so}), we would expect them to increase employment shares in that sector and occupation (higher λ_{so}), other things equal. In contrast, to the extent that new technologies substitute for workers in performing tasks in a sector and occupation (lower T_{so}), we would expect them to decrease employment shares in that sector and occupation (lower λ_{so}), other things equal. We provide evidence below that the impact of new technologies on the employment shares of different occupations is related to the production tasks performed by workers within those occupations.

Our second key result for equilibrium worker sorting is for average earnings in each occupation.

Proposition 2 ***(Occupational Average Earnings)*** *The model's sufficient statistic for average earnings in occupation o ($\overline{wage}_o$), including both human capital and ability, is:*

$$\overline{wage}_o = \mathbb{E}\left[w_o h_o z\right] = \gamma\left(1-\ell_o\right)^{-1/\beta}\left(P_M^{\alpha}P_N^{1-\alpha}\right)\Phi^{1/\theta}, \tag{17}$$

where $\gamma = \Gamma\left(\frac{\theta-1}{\theta}\right)$ and $\Gamma\left(\cdot\right)$ is the Gamma function.

Proof. See Appendix. ■

Therefore differences in average earnings ($\overline{wage}_o$) across occupations o are explained in the model by differences in human capital investments ($\left(1-\ell_o\right)^{-1/\beta}$). Occupations in which human capital investments are more productive (higher ϕ_o) have higher human capital investments (ℓ_o) and higher average earnings ($\overline{wage}_o$). In contrast, average earnings are no higher in occupations that have higher average effective units of labor (higher T_{so}) or higher wages per effective unit of labor (higher w_o). The reason is a selection effect. On the one hand, higher T_{so} and w_o directly *increase* average wages for a given fraction of workers choosing to enter an occupation. On the other hand, higher T_{so} and w_o induce a higher fraction of workers to choose an occupation, which indirectly *reduces* average wages through a composition effect of a higher fraction of workers with lower draws for effective units of labor. With a Fréchet distribution for worker ability, these two effects exactly offset one another, leaving average earnings unchanged. Although this exact offset is a feature of the Fréchet distribution, more generally, these two effects work in different directions and dampen the impact of T_{so} and w_o on average earnings ($\overline{wage}_o$).

Thus the model highlights average earnings as the second key endogenous variable of interest in our empirical work. New technologies affect average earnings through changing the return to human capital accumulation in different occupations. Totally differentiating average earnings (17), changes in the relative average earnings ($\omega_{om} = \overline{wage}_o / \overline{wage}_m$) of two occupations o and m depend solely on changes in human capital investments:

$$\frac{d\omega_{om}}{\omega_{om}} = \frac{1}{\beta}\left[\frac{\ell_o}{1-\ell_o}\frac{d\ell_o}{\ell_o} - \frac{\ell_m}{1-\ell_m}\frac{d\ell_m}{\ell_m}\right], \tag{18}$$

where changes in these human capital investments depend solely on changes in the rate of return to these investments (ϕ_o):

$$\frac{d\ell_o}{\ell_o} = \frac{1}{1+\beta\phi_o}\frac{d\phi_o}{\phi_o}, \qquad\qquad \frac{d\ell_m}{\ell_m} = \frac{1}{1+\beta\phi_m}\frac{d\phi_m}{\phi_m}. \tag{19}$$

Therefore, to the extent that new technologies are complementary to human capital investments within an occupation (higher ϕ_o), we would expect them to increase occupation average earnings (higher $\overline{wage}_o$). In contrast, to the extent that new technologies substitute for human capital investments within an occupation (lower ϕ_o), we would expect them to decrease occupation average earnings (lower $\overline{wage}_o$).

2.5 Quantification

We now show how observed values of the two key endogenous variables in the model, employment shares and average earnings, can be used to solve for unobserved values of the rate of return to human capital accumulation and an adjusted measure of the average effectiveness of workers of performing tasks in each sector and occupation. We first assume central values for the model's parameters from the existing empirical literature. We follow Hsieh, Hurst, Jones, and Klenow (2013) in assuming a value for the Fréchet shape parameter determining worker comparative advantage of $\theta = 3.44$ and a value for the parameter governing the tradeoff between consumption and human capital accumulation of $\beta = 0.693$. Using these assumed parameters and the expressions for equilibrium human capital investments (10) and average earnings (17), we can recover the return to human capital accumulation from observed average earnings in each occupation relative to the geometric mean of average earnings across occupations:

$$\frac{\overline{wage}_o}{\prod_{o=1}^{O}\left[\overline{wage}_o\right]^{\frac{1}{O}}} = \frac{(1-\ell_o)^{-1/\beta}}{\prod_{o=1}^{O}\left[(1-\ell_o)^{-1/\beta}\right]^{\frac{1}{O}}} = \frac{(1+\beta\phi_o)^{-1/\beta}}{\prod_{o=1}^{O}\left[(1+\beta\phi_o)^{-1/\beta}\right]^{\frac{1}{O}}}. \tag{20}$$

Given observed occupation average earnings ($\overline{wage}_o$) and an assumed value for β, this provides a system of O equations that can be solved for unique values of the O unobserved human capital returns (ϕ_o).

Using these solutions for human capital returns (ϕ_o), the assumed values of β and θ, and the expressions for equilibrium human capital investments (10) and the choice probabilities (15), we can recover an adjusted measure of the relative effectiveness of workers in performing tasks for each sector and occupation ($\mathbb{A}_{so}$) from observed employment shares (λ_{so}) relative to their geometric mean:

$$\frac{\mathbb{A}_{so}}{\left[\prod_{s=1}^{S}\prod_{o=1}^{O}\mathbb{A}_{so}\right]^{\frac{1}{O}}} = \frac{\lambda_{so}/\left[\prod_{s=1}^{S}\prod_{o=1}^{O}\lambda_{so}\right]^{\frac{1}{O}}}{\mathbb{B}_{so}/\left[\prod_{s=1}^{S}\prod_{o=1}^{O}\mathbb{B}_{so}\right]^{\frac{1}{O}}}, \tag{21}$$

where $\mathbb{B}_{so} = (1-\ell_o)^{\theta/\beta}\,\ell_o^{\theta\phi_o}$ captures the contribution of human capital investments to employment shares and $\mathbb{A}_{so} = T_{so}w_o^{\theta}\bar{h}_o^{\theta}$ captures the average effectiveness of workers in performing tasks within each sector and occupation (T_{so}), human capital productivity ($\bar{h}_o$), and the wage per effective unit of labor (w_o). We use the solutions for unobserved human capital returns (ϕ_o) from (20) and adjusted task effectiveness ($\mathbb{A}_{so}$) from (21) to quantify the relative importance of these two different mechanisms for explaining the observed changes in employment shares and average earnings in the data.

2.6 Extensions

In the web appendix, we show how the above baseline model can be extended to incorporate multiple types of workers with different *ex ante* characteristics (e.g. age, gender and general schooling), as observed in the data. Output (Y_{Ms}) in each sector s is assumed to be a constant elasticity of substitution (CES) function of labor inputs (H_s^x) of each worker type x, which are in turn a constant elasticity of substitution function of the human capital (H_{so}^x) of occupations o within that sector s for worker type x. Workers of each type choose an occupation and sector as above. The model again yields two sufficient statistics in the form of employment shares (λ_{so}^x) for each sector, occupation and worker type and average wages ($\overline{wage}_o^x$) for each occupation and worker type. Differences across worker types in the average effectiveness of performing tasks within each sector and occupation (T_{so}^x) generate variation across worker types in employment shares (λ_{so}^x). In contrast, the return to human capital investments in each occupation (ϕ_o^x) affects both employment shares (λ_{so}^x) and average occupational earnings ($\overline{wage}_o^x$).

In the web appendix, we also extend the above baseline model to incorporate multiple locations indexed by $r = 1, \ldots, R$ (e.g. urban versus rural areas), as again observed in the data. We assume that each worker draws effective units of labor (z) for each occupation, sector and region from an independent Fréchet distribution. Given these realizations for effective units of labor, workers choose an occupation, sector and region. The model again yields two sufficient statistics in the form of employment shares (λ_{rso}) and average wages ($\overline{wage}_{ro}$). Differences across regions in the average effectiveness of workers in performing tasks within each sector and occupation (because of regional differences in technology T_{rso}) generate variation across regions in employment shares (λ_{rso}). In contrast, relative average earnings ($\overline{wage}_{ro}$) across occupations within the same region depends only on relative returns to human capital investments (ϕ_o), because the cost of living within regions is the same across occupations.

Therefore augmenting the baseline model to incorporate multiple types of workers with different observed characteristics or multiple regions preserves the model's key predictions for the determinants of employment shares and average earnings. We now provide evidence on the extent to which employment shares and average wages have changed systematically towards occupations performing particular types of tasks and the extent to which these changes are driven by new technologies. Before doing so, we first discuss our data, and next introduce our new empirical methodology for measuring the individual tasks performed by workers within each occupation.

3 Data Description

In this section, we discuss our data sources for the key objects of interest in the model (employment shares and wages) and for our measures of the production tasks undertaken within each occupation. Our main data source on employment and worker characteristics is the individual-level records from Integrated Public Use Microdata Series (IPUMS) from 1880-2000: see Ruggles, Alexander, Genadek, Goeken, Schroeder, and Sobek (2010). We construct our datasets to make maximum use of available sample sizes, wage and education data, and geographic identifiers. We use the 100 percent samples for 1880 and 1940 and the largest available sample size for all other years (typically 5 percent).[14] Wages and education are reported from 1940 onwards. We also use the estimates of wages by occupation in 1880 from Preston and Haines (1991), as used in Abramitzky, Boustan, and Eriksson (2012, 2014).

To make maximum use of the available information, we create two datasets that take different stands on the trade-off between variable availability and time period. Our first dataset consists of three cross-sections, including the 100 percent Census samples for 1880 and 1940 and the 5 percent sample for 2000. Focusing on these three cross-sections enables us to include employment and wage data for all three periods, and education data for 1940 and 2000. Our second dataset consists of cross-sections for 1860

[14] In robustness tests, we also report some results using the 1860 data, in which the number of occupations reported is substantially smaller than after 1880.

and for twenty-year time intervals from 1880-2000, using either 100 percent census samples (for 1880 and 1940) or the largest available sample for the other years. Including this higher frequency information limits the data available for all years to employment, with wages and education reported from 1940-2000. When we use samples of less than 100 percent, the sampling probabilities vary with worker characteristics, such as ethnicity. Therefore we weight individuals by their sampling weights to ensure that the data are representative for the United States as a whole and each sector and occupation.

We use the standardized 1950 occupation classification from IPUMS, which distinguishes eleven two-digit occupations (e.g. "Clerical and Kindred") and 281 three-digit occupations (e.g. "Opticians and Lens Grinders and Polishers"). We also use the standardized 1950 sector classification from IPUMS, which distinguishes twelve two-digit sectors (e.g. "Finance, Insurance and Real Estate") and 158 three-digit sectors (e.g. "Motor Vehicles and Motor Vehicle Equipment").[15] Since we are concerned with employment structure, we omit workers who do not report an occupation or a sector (e.g. because they are unemployed or out of the labor force). We also exclude workers in agricultural occupations or sectors, because we compare the task content of employment in urban and rural areas over time, and agriculture is unsurprisingly overwhelmingly located in rural areas.[16] We define urban and rural areas based on time-varying metropolitan boundaries to ensure that urban areas correspond to meaningful economic units in each year. But we also report a robustness test in which we define urban areas based on the boundaries of administrative cities that are more stable over time.

We develop a new methodology for measuring the individual production tasks undertaken by workers within occupations using the detailed occupational descriptions in the Dictionary of Occupational Titles (DOTs). Previous research using the DOTs has focused on the numerical scores that summarize job requirements, as constructed by the U.S. Department of Labor. ALM distinguishes between non-routine interactive tasks ("Direction, Control and Planning (DCP)"), non-routine analytic tasks ("Quantitative Reasoning (GEDMATH)"); routine cognitive tasks ("Set Limits, Tolerances and Standards (STS)"); routine manual tasks ("Finger Dexterity (FINGER)"); and non-routine manual tasks ("Eye-Hand-Foot Coordination (EYEHAND)"). Subsequent research following AKK has aggregated the first two categories into "routine" ((DCP+GEDMATH)/2), "non-routine" ((STS+FINGER)/2) and "manual" (EYEHAND). We compare our measures of individual production tasks with these more aggregated numerical scores from existing research, including modern values (from U. S. Department of Labor 1991) and historical values (from U. S.

[15] See IPUMS for the full concordance between two-digit and three-digit occupations and sectors. While both occupation and sector classifications are standardized by IPUMS, there are a small number of occupations and sectors that enter and exit the sample over time. Our results are robust to restricting attention to occupations and sectors that are present in all years.

[16] Our key findings, however, are robust to the inclusion of these agricultural workers. When we examine task inputs in urban versus rural areas in Section 7, we include sector fixed effects in our regressions for each year, which controls for the effect of changes over time in the aggregate share of employment in agriculture on each sector in each year. For an analysis of urbanization and structural transformation away from the agricultural sector, see Michaels, Rauch, and Redding (2012).

Department of Labor 1949, which was the first edition of the DOTs to report these aggregate measures). We follow existing research in converting each numerical score into percentiles of its distribution across occupations, since these numerical scores do not necessarily have a common cardinal scale.[17]

In contrast to this previous research, our new methodology uses the verbs from the detailed occupational descriptions in DOTs to measure individual production tasks. This approach enables us to measure the task content of employment at much higher resolution than previous research and to examine the individual production tasks included within more aggregated measures based on numerical scores. We use a comprehensive list of over 3,000 English verbs from "Writing English," a company that offers English language consulting.[18] We search for appearances of each of these verbs in the occupational descriptions of the DOTs. We quantify the nature of these production tasks using the meanings of the verbs from Roget's Thesaurus, which is the standard reference for word usage in English.[19] In our baseline specification, we use occupational descriptions from the digital edition of the 1991 DOTs.[20] But our use of occupational descriptions also enables us to undertake a robustness test using the first edition of the DOTs in 1939 (U. S. Department of Labor 1939), which did not report numerical scores. Comparing the occupational descriptions from 1991 and 1939 DOTs, we can examine changes in the relative importance of individual production tasks within occupations and the extent to which the ranking of occupations in terms of the frequency with which they use different types of tasks is stable over time.

Our main data source for industry technology use is the 1947 Bureau of Economic Analysis Input-Output Table (Bureau of Economic Analysis 2015). This is the earliest available input-output dataset for the U.S., and its date falls roughly in the middle of our period of analysis. Specifically, we measure an industry's use of a technology by the share of its inputs that it purchases from the industry that produces the technology. This approach has two main advantages. First, we can consider a number of alternative technologies that experienced substantial innovation over our long historical time period. We distinguish (i) Office and computing machines (Office, Computing and Accounting Machines [51]); (ii) Electrical machinery (Electric Industrial Equipment and Apparatus [53], Household appliances [54], Electric Lighting and Wiring Equipment [55], Radio, Television and Communication Equipment [56], Electronic Components and Accessories [57] and Miscellaneous Electrical Machinery, Equipment and Supplies [58]); (iii) Transport machinery (Motor Vehicles and Equipment [59], Aircraft and Parts [60], and Other Transportation Equipment [61]); (iv) All machinery (Industries [43], [46]-[49], and [51]-[63], which includes categories (i)-(iii) above). Second, we can construct these measures of industry technology use with input-output ta-

[17]We find similar results if we use the raw numerical scores instead of their percentiles.

[18]See http://www.writingenglish.com/englishverbs.htm.

[19]We use the online computer-searchable edition of Roget (1911): http://machaut.uchicago.edu/rogets.

[20]Following the first edition of the DOTs in 1939, there were major revisions in 1949, 1965, 1977, and 1991. Each revision updated the occupational descriptions with the objective of reflecting the contemporary nature of work in each occupation.

bles from different years during our long historical time period. While our baseline specification uses the first input-output table from 1947, we also undertake a robustness test using a contemporary input-output table from 2002.

4 Measuring Production Tasks

In this section, we discuss in more detail our methodology for measuring individual production tasks using the verbs from the around 12,000 occupational descriptions in the DOTs. We start with the comprehensive list of over 3,000 English verbs from the language consulting company "Writing English" discussed above. Using this list of verbs, we search each occupational description in the DOTs for occurrences of each verb in the first-person singular (e.g. (I) talk), third-person singular (e.g. (she) talks) or present participle (e.g. (he is) talking). In each case, the verbs capture an action (bring, read, walk), an occurrence (happen, become), or a state of being (exist, stand), and hence capture the tasks performed within an occupation. To take an example from our own experience, the occupational description for an economist is:

> "ECONOMIST: *Plans*, *designs*, and *conducts research* to aid in interpretation of economic relationships and in solution of problems *arising* from production and distribution of goods and services: *Studies* economic and statistical data in area of specialization, such as *finance*, *labor*, or agriculture. *Devises* methods and procedures for *collecting* and *processing* data, *utilizing* knowledge of available sources of data and various econometric and *sampling* techniques. *Compiles* data *relating* to *research* area, such as employment, productivity, and *wages* and hours. *Reviews* and *analyzes* economic data in *order* to *prepare reports detailing results* of investigation, and to *stay* abreast of economic *changes* ...,"

where the words detected by our procedure as capturing the tasks performed by an economist are italicized.[21] Note that sometimes the first-person singular, third-person singular or present participle forms of a verb have the same spelling as the corresponding adjectives and nouns (e.g. "prepare *reports*"). In this case, our procedure treats these adjectives and nouns as verbs. To the extent that the use of the same word as an adjective or noun is closely related to its use as a verb, both uses are likely to capture the tasks performed within an occupation.

We measure the importance of a production task for an occupation, using the frequency with which a verb v appears for an occupation o relative to all appearances of verbs for that occupation:

$$\text{VerbFreq}_{vo} = \frac{\text{Appearances of verb } v \text{ matched to } o}{\text{Appearances of all verbs matched to } o}. \tag{22}$$

[21]As an indication of the wide coverage of our list of over 3,000 verbs, only 1,830 of these verbs appear in the 1991 DOTs occupational descriptions.

We focus on the frequency rather than the number of occurrences of a verb to capture the relative importance of a task for an occupation and to control for variation across occupations in the length of occupational descriptions. As for the numerical scores discussed in the previous section, we convert VerbFreq_{vo} into percentiles of its distribution across occupations.[22]

To quantify the nature of these production tasks, we use the meanings of the verbs. In particular, we use an online computer-searchable version of Roget's Thesaurus (Roget 1911), which explicitly classifies words according to their underlying concepts and meanings. Roget's classification was inspired by natural history, with its hierarchy of Phyla, Classes, Orders and Families. Therefore words are grouped according to progressively more disaggregated classifications that capture ever more subtle variations in meaning. A key advantage of the thesaurus classification is that it explicitly takes into account that words can have different meanings depending on context, by allowing the same word to appear more than once and including extensive cross references to link related groups of words.[23]

Roget's Thesaurus is organized into six "Classes" that are further disaggregated into the progressively finer subdivisions of "Divisions," "Sections" and "Categories." The first three classes cover the *external world*: Class 1 (Abstract Relations) deals with ideas such as number, order and time; Class 2 (Space) is concerned with movement, shapes and sizes; and Class 3 (Matter) covers the physical world and humankind's perception of it by means of the five senses. The last three classes relate to the *internal world* of human beings: the human mind (Class 4, Intellect), the human will (Class 5, Volition) and the human heart and soul (Class 6, Emotion).

To characterize the meaning of each verb v, we use the frequency with which it appears in each subdivision k of Roget's Thesaurus:

$$\text{ThesMean}_{vk} = \frac{\text{Appearances of verb } v \text{ in subdivision } k \text{ of thesaurus}}{\text{Total appearances of verb } v \text{ in thesaurus}}, \tag{23}$$

where our use of a frequency explicitly takes into account that each verb can have multiple meanings and provides a measure of the relative importance of each meaning. In counting the appearances of verbs we make use of the thesaurus's structure, in which words with similar meanings appear under each thesaurus Category in a list separated by commas or semi-colons. Based on this structure, we count appearances of a verb that are followed by a comma or semi-colon, which enables us to abstract from appearances of a word in idioms that do not reflect its common usage.[24]

For our baseline measures of task input, we use time-invariant occupational descriptions from the

[22] VerbFreq_{vo} has a natural interpretation as a frequency and we find similar results using the raw measure.

[23] For further discussion of the genesis of Roget's Thesaurus, see for example Hüllen (2003).

[24] For example, the verb "consult" appears in six thesaurus Categories. The entry followed by a comma is 695 Advice, which captures the word's meaning. Entries not followed by a comma correspond to idiomatic uses not closely related to the word's meaning: 133 Lateness ("consult one's pillow"); 463 Experiment ("consult the barometer"); 707 Aid ("consult the wishes of"); 943 Selfishness ("consult one's own pleasure"); 968 Lawyer ("juris consult [Latin]").

1991 DOTs for all years in our sample, which ensures that our results are not driven by changes in the language used in the occupational descriptions over time (because these occupational descriptions are held fixed for all years in our sample). This use of time-invariant occupational descriptions implies that changes in task input over time are driven solely by changes in employment shares across occupations with different (time-invariant) task inputs. As Roget's Thesaurus was compiled in a different year (1911) from the occupational descriptions (1991), there remains the concern that the meaning of verbs could have changed over time. However, Roget's thesaurus is still the seminal reference for English language use today. Furthermore, we report a robustness test in which we use occupational descriptions from the 1939 DOTs. This is the earliest DOTs, which dates from around the middle of our 1880-2000 sample period, and enables us to look at historical classifications of occupations. We find a similar pattern of results using occupational descriptions for 1939 and 1991, even though they are more than fifty years apart, which suggests that the difference in dates between the thesaurus and occupational descriptions is not consequential for our findings. As a further check on our results, we show that aggregating our measures of individual production tasks to the six thesaurus Classes (using either the 1991 or 1939 occupational descriptions) yields a similar pattern of results to the aggregate numerical scores used in previous research with the DOTs.[25] Additionally, we use the 1939 and 1991 occupational descriptions (as well as numerical score measures for 1949 and 1991) to examine changes in task input within occupations.[26]

Combining the frequency with which a verb appears in each occupation's description (VerbFreq_{vo}) and the frequency with which the verb appears in each subdivision of the thesaurus (ThesFreq_{vk}), we obtain a quantitative measure of the extent to which the tasks performed in an occupation involve the concepts from each thesaurus subdivision:

$$\text{ThesFreq}_{ko} = \sum_{v \in V} \text{VerbFreq}_{vo} \times \text{ThesMean}_{vk}, \tag{24}$$

where we again convert ThesFreq_{ko} into percentiles of its distribution across occupations.[27]

We use VerbFreq_{vo} from (22) and ThesFreq_{ko} from (24) as our two key empirical measures of task inputs. Our use of verbs from occupational descriptions (VerbFreq_{vo}) makes it possible for the first time to measure the importance of inputs of individual production tasks. Our use of the meanings of these verbs based on thesaurus subdivisions (ThesFreq_{ko}) enables us to isolate the common characteristics of those individual production tasks that have become more or less important over time, without imposing prior

[25] We find that our measures of task inputs for each occupation are also correlated in the expected way with separate measures of job requirements from the O*NET database, as discussed below and reported in section A9 of the web appendix.

[26] See Spitz-Oener (2006) for evidence using German survey data on within-occupation changes in task input. Correlating the 1949 and 1991 percentile numerical scores across the occupations in the data in 2000, we find the following correlations: non-routine (0.799), routine (0.770) and manual (0.700), suggesting a high correlation in the ranking of occupations in terms of task input over time.

[27] ThesFreq_{ko} has a natural interpretation as the product of two frequencies and we find similar results using the raw measure.

structure from the numerical scores about what types of tasks are or are not important.

5 Trends in Task Input 1880-2000

In this section, we provide evidence on the model's two key endogenous outcomes of interest, employment shares and average wages. We demonstrate systematic changes in employment shares and average wages across occupations that are related to the nature of the production tasks undertaken within those occupations, as measured by our new methodology. We find the largest increases for tasks involving the internal world, including the formation of ideas (e.g. Analyze, Confer, Evaluate), and the largest decreases for those involving the external world, including the manipulation of physical matter (e.g. Deliver, Grind and Weld). We use the structure of the model to recover the changes in the average effectiveness of workers in performing tasks and in the rate of return to human capital investments underlying these changes in employment shares and average wages. In the next section, we examine the extent to which these observed changes in task inputs are related to new technologies, as suggested by the model.

5.1 Employment Shares and Aggregate Task Inputs

We begin by presenting results for employment shares using the more aggregated numerical scores considered in previous research, but for our much longer historical time period. Figure 1 shows the employment share-weighted average of each numerical score (non-routine interactive, non-routine analytic, routine analytic, routine cognitive and routine non-manual) over time. Numerical scores are expressed as percentiles and weighted by occupational employment shares in each year. Each series is expressed as an index relative to its value in 1880 (so that each series takes the value one in 1880).

Figure 1 is analogous to Figure I in ALM, but for the period 1880-2000 instead of 1960-2000. Comparing these two figures, we confirm existing findings for the 1960-2000 period. First, we find a sharp increase in the share of the labor force employed in occupations that make intensive use of non-routine interactive and non-routine analytic tasks throughout this period. Second, employment in occupations that use non-routine manual tasks intensively declines throughout this period. Third, we find a decline in the share of the labor force employed in occupations that use routine cognitive and routine manual tasks intensively from 1960 onwards. This period from 1960 onwards coincides with the dissemination of computers. Therefore the decline in these two routine categories and the acceleration after 1960 in the rise of the two non-routine categories are consistent with the view that computers complemented non-routine tasks and substituted for routine tasks.

Two other features are apparent from Figure 1. First, notwithstanding the important changes after 1960, there is an increase in the non-routine interactive inputs and a decrease in the non-routine manual inputs from the late-nineteenth century onwards.[28] Second, these longer-term trends are interrupted during the period 1920-60, when non-routine interactive and analytic inputs are relatively flat, and there is a slower decline in non-routine manual inputs from 1920-1940. In the web appendix, we show that this long-term rise in non-routine tasks is not apparent in the earlier 1860-1880 period (see Figure A1).[29] We also show that the late-nineteenth and early-twentieth centuries were periods of rapid diffusion of a cluster of information technologies centered on the typewriter (which reduced the cost of producing written information) and the telephone (which facilitated communication at a distance), as shown in Figure A2 in the web appendix.[30] Improved filing techniques and especially the invention of vertical filing around the turn of the twentieth century allowed for much easier storage and retrieval of information (Ellen-Poe 2014, Michaels 2007). This timing suggests that the growth in non-routine inputs around the turn of the twentieth century could be related to this earlier revolution in information technology, an idea that we explore further below. The interruption of these longer term trends during the period 1920-1960 coincides with a compression in wage inequality in the decades surrounding the Second World War.[31] We provide further evidence on the relationship between these changes in task inputs and wage inequality below.

5.2 Employment Shares and Individual Production Tasks

We now present employment share results for individual production tasks using our new methodology based on occupational descriptions from Section 4 above. We begin by relating the verbs capturing tasks to the numerical scores used in the previous section. We next examine the meaning of these verbs using the sections from Roget's thesaurus. Finally, we use the meaning of these verbs to provide finer resolution evidence on changes in task inputs over time.

Table 1 examines the individual production tasks captured by each of the five numerical scores (non-routine interactive, non-routine analytic, routine analytic, routine manual and routine non-manual). We correlate each of these numerical scores with each verb across occupations and report the top twenty verbs most correlated with each numerical score.[32] Three main features are apparent from the table. First,

[28]Consistent with these results, Katz and Margo (2014) finds a rise in white-collar employment in the late-nineteenth century, so that by 1900 one in 15 workers were white collar. In contrast to that study, we adopt a task-based approach that uses variation in the disaggregated tasks performed by workers within broad occupational categories such as white collar.

[29]We report these results for 1860-1880 as a robustness test rather than as part of our main specification, because the number of occupations reported in the data is substantially smaller for years before 1880.

[30]For a discussion of this earlier revolution in information technology, see for example Hunt and Hunt (1986) and Phister (1979). The first commercially-successful typewriter was invented in 1868 by Christopher L. Sholes, Carlos Glidden and Samuel W. Soule in Milwaukee, Wisconsin. In 1876, Alexander Graham Bell was the first to be granted a United States patent for the telephone. Both technologies diffused particularly rapidly in the closing two decades of the nineteenth century. The electrical telegraph was developed and patented in the United States by Samuel Morse somewhat earlier in 1837.

[31]See, in particular, Piketty and Saiz (2003) and Piketty, Saiz, and Zucman (2014).

[32]While IPUMS uses consistent (1950) occupational descriptions over time, some occupations do not exist in some years. Unless otherwise indicated, all reported correlations below are across the sample of occupations in 2000. We find the same pattern of results both qualitatively and quantitatively using the sample of occupations from other years.

the numerical scores generally capture an intuitive pattern of tasks. For example, the top five tasks most correlated with non-routine analytic are analyze, develop, conduct, plan and direct. In contrast, the top five tasks most correlated with non-routine manual are: prevent, move, line, climb and wall. Second, there is some similarity in the tasks captured by non-routine interactive and non-routine analytical. Ten of the twenty verbs most strongly correlated with these two categories are common to both. Even the verbs that differ are telling. Non-routine interactive is highly correlated with managerial verbs, such as direct, prepare, project, and budget. In contrast, non-routine analytical is more strongly correlated with quantitative tasks, such as test, result, interpret, and formulate. Third, there is also some similarity in the tasks captured by routine cognitive and routine manual. Indeed, many of the tasks most correlated with routine cognitive are manual rather than cognitive (e.g. cut and screw), which suggests either that routine cognitive and manual tasks are highly correlated or that this category provides an imperfect measure of cognitive tasks. Taken together, this pattern of results provides support for more recent research following AKK that has aggregated non-routine interactive and non-routine analytic into a single category ("non-routine"), aggregated routine cognitive and routine manual into a single category ("routine"), and retained non-routine manual as a separate "manual" category. From now onwards, we follow AKK in aggregating the numerical scores into the three categories of non-routine, routine and manual.

Having related our measures to the numerical scores used in previous research, we now examine the meanings of these verbs corresponding to individual production tasks. For each of the 39 sections of Roget's Thesaurus, Table 2 reports the top five verbs whose usage is most concentrated in that thesaurus section (the verbs with the five highest values of ThesMean_{vk} from equation (23)). If two or more verbs have the same values of ThesMean_{vk}, we rank them by their number of occurrences in the thesaurus, so as to give more weight to verbs that are more prevalent in language use. Two main features are again apparent from the table. First, thesaurus sections successfully capture the meaning of verbs. For example, the top five tasks concentrated in Section 3.1 (Matter in General) are: weigh, float, swim, balloon and pound. In contrast, the top five tasks concentrated in Section 4.2.1 (Nature of Ideas Communicated) are: decipher, annotate, interpret, fudge and clarify. Second, thesaurus sections provide a substantially finer resolution on the production tasks performed within occupations than the numerical scores discussed above. For example, the verbs analyze, develop and conduct appear as the top three verbs most correlated with the Non-routine Interactive numerical score in Table 1. However, these verbs have quite different meanings and appear in quite different sections of the thesaurus. Although not shown in Table 2, analyze appears in Sections 1.3 (Abstract Relations: Quantity) and 4.1.2 (Formation of Ideas: Precursory Conditions and Operations); develop occurs in Sections 1.8 (Abstract Relations: Causation), 2.2 (Words Relating to Space: Dimensions) and 2.4 (Words Relating to Space: Motion); and conduct is found in Sections 2.4 (Words

Relating to Space: Motion) and Sections 5.1.3 (Individual Volition: Voluntary Action).

We now use the meanings of these verbs to provide finer resolution evidence on the evolution of task inputs over time. Figures 2-3 are analogous to Figure 1 above, but use thesaurus sections rather than numerical scores to establish the following results. First, we find substantially larger changes in task inputs over time using our measures of production tasks (based on verbs and their meanings in the thesaurus) than using the numerical scores. In Figures 2-3, the scale of the vertical axis ranges from 0.5-2.5, compared to 0.6-1.4 in Figure 1. Although the use of more disaggregated categories is likely to generate greater variation, this confirms the ability of the meanings of verbs to capture changes in task inputs over time. Second, the divisions of the thesaurus with the largest increases in task inputs over time are Class 1: Abstract Relations (top left in Figure 2), Class 4.1: Formation of Ideas (bottom right in Figure 2), and Class 4.2: Communication of Ideas (top left in Figure 2). In contrast, the divisions of the thesaurus with the smallest increases in task inputs over time are Class 2: Space (top right in Figure 2) and Class 3: Matter (bottom left in Figure 2). This pattern of results is consistent with a reduction in the share of the labor force employed in tasks relating to the manipulation of the physical world (space and matter) and an increase in the share of the labor force employed in analytical and interactive tasks (abstract thought and the formation and communication of ideas). It is also noteworthy that the use of tasks from thesaurus division 6 (religion, morality, and emotion) has not changed much over the period we analyze. In other words, even in the realm intellectual tasks, the use of those related to emotions has changed much less than those related to abstract thinking and interaction.[33] Finally, even within the thesaurus classes that involve abstract thinking and interaction there is substantial heterogeneity, highlighting the additional insights from more disaggregated task measures. In particular, some of the largest increases in task inputs are observed for Sections 1.5 (Abstract Relations: Number), 1.6 (Abstract Relations: Time), 4.1.1 (Formation of Ideas in General), and 4.2.3 (Means of Communicating Ideas).

5.3 Tasks and Wage Inequality

While the previous two sections have presented results for employment shares, we now incorporate information on average wages, as the second key endogenous outcome in the model, and examine the relationship between tasks and wage inequality. We compute total employment and average wages for each occupation in 1880, 1940 and 2000. In Figure 4, we display the cumulative distribution of employment shares across percentiles of the occupation wage distributions. On the horizontal axis, occupations are sorted according to their percentiles of the occupation wage distribution in a given year. On the verti-

[33]In a robustness check, we find that our measures of abstract thinking and interaction (Classes 4 and 5 of the thesaurus) are strongly correlated with measures of the interactiveness of occupations based on employee and employer surveys from the O*NET database, as discussed in Section A9.2 of the web appendix.

cal axis, we display the cumulative sum of employment shares across the percentiles of the occupation wage distribution in that given year. As the employment share distributions are cumulative, they necessarily add up to one for each year. Furthermore, the slope of each cumulative distribution corresponds to employment at that percentile of the wage distribution.

Comparing the cumulative distributions for 1880 and 1940, we find that there is greater mass of workers at low wages in 1880 than in 1940 (the solid black line is *above* the solid gray line at low wages) and a greater mass of workers at high wages in 1880 than in 1940 (the solid black line is *below* the solid gray line at higher wages).[34] This implies a decline in wage inequality across occupations during the first half of our sample period.[35] Comparing the cumulative distributions for 1940 and 2000, we find that the two distributions track one another relatively closely at low wages, but there is a smaller mass of workers at intermediate wages in 2000 than in 1940 (the solid gray line has a steeper slope at intermediate wages than the black dashed line) and a greater mass of workers at high wages in 2000 than in 1940 (the black dashed line has a steeper slope than the solid gray line at high wages). Therefore the second half of our sample period is characterized by increased wage inequality across occupations and a polarization of wages towards the top of wage distribution at the expense of the middle of the wage distribution.[36]

We now relate these changes in the distribution of employment across percentiles of the occupation wage distribution to the production tasks undertaken within occupations. In Figure 5, we begin by showing the cumulative distribution of employment shares across percentiles of the task distribution over time. The top left panel sorts occupations based on percentiles of the non-routine task distribution; the top right panel sorts them based on percentiles of the routine task distribution; and the bottom left panel sorts them based on percentiles of the manual task distribution. From the bottom left panel, the largest change in the employment distribution from 1880-1940 was a shift in employment towards low levels of manual tasks (the solid gray line is substantially above the solid black line at low levels of manual tasks). In contrast, from the top-left panel, the largest changes in the employment distribution from 1940-2000 were from low to intermediate levels of non-routine tasks (the dashed black line is below the solid gray line below the median) and from intermediate to high levels of non-routine tasks (the dashed black line is below the solid gray line above the median). Therefore the large reallocations of employment across percentiles of the wage distribution in Figure 4 involve the large reallocations of employment across percentiles of the task distribution shown in Figure 5.

[34] Since the solid black line is below the solid gray line at higher wages, it must have a greater slope at the highest wages in order for the cumulative distribution to add up to one, which implies greater employment at the highest wages.

[35] This decline in wage inequality across occupations from 1880-1940 is consistent with the decline in income inequality across individuals found using personal income taxation data in Piketty and Saiz (2003) and Piketty, Saiz, and Zucman (2014) and with the evidence in Goldin and Katz (2008).

[36] For further evidence of wage polarization in the closing decades of the twentieth century, see Autor and Dorn (2013).

To tighten this connection, Figure 6 shows the cumulative distribution of tasks across percentiles of the occupation wage distribution in each year.[37] The top-left panel shows these distributions for 1880; the top-right panel shows them for 1940; and the bottom-left panel shows them for 2000. In all three years, we find that non-routine tasks are concentrated at higher percentiles of the wage distribution than routine tasks, and routine tasks are concentrated at higher percentiles of the wage distribution than manual tasks. But there are substantial changes in the extent to which this is the case over time. From 1880-1940, we find a convergence in the distribution of tasks across percentiles of the wage distribution (all three lines move closer together). In contrast, from 1940-2000, we find a sharp increase in the extent to which non-routine tasks are concentrated at higher percentiles of the wage distribution. Together Figures 4-6 paint a consistent picture of wage convergence from 1880-1940 and wage divergence from 1940-2000, with the increase in wage inequality from 1940-2000 driven by a reallocation of employment towards non-routine tasks and and an increase in the wage premium for these tasks.

In the model, the increase in average wages in non-routine occupations is driven by changes in the rate of return to human capital investments in these occupations, whereas the reallocation of employment towards these occupations reflects changes in both the rate of return to human capital investments and the average effectiveness of workers in performing tasks in these occupations. In section A6.3 of the web appendix, we use the model to solve for the implied values of the relative rate of return to human capital investment (ϕ_o) and adjusted task effectiveness ($\mathbb{A}_{so}$), as discussed in section 2.5 above. In general, we find larger changes in relative task effectiveness than in relative returns to human capital accumulation. For both 1880-1940 and 1940-2000, we find an increase in the rate of return to human capital investments for non-routine occupations, which is largest in the first sub-period for the most non-routine occupations, before becoming greater in the second sub-period for occupations with intermediate-high levels of non-routine tasks (see Figure A3). From 1880-1940, we find an increase in relative task effectiveness for workers in the most manual occupations; little change in the relative task effectiveness across occupations with different levels of routine tasks; and an increase in task effectiveness for both the least and most non-routine occupations (see Figure A4). From 1940-2000, we find a decline in the relative task effectiveness of workers in more manual occupations; a secular fall in the relative task effectiveness of workers performing more routine tasks relative to those performing less routine tasks; and a secular rise in the relative task effectiveness of more non-routine occupations. The results suggest that the relative productivity of tasks changed somewhat differently from the late-nineteenth to the early-twentieth centuries versus from the early to the late-twentieth century. Whereas the first half of our sample period saw an increase in the

[37] Specifically, each occupation has a task measure (percentile score) and a wage percentile in a given year. We sort occupations by their wage percentile in a given year (horizontal axis). We then cumulate the task measure across these percentiles of the occupation wage distribution, scaling by the sum of the task measure to obtain a cumulative distribution (vertical axis).

relative productivity of the most manual tasks, the second half of our sample period was characterized by a more pervasive increase in the relative productivity of non-routine tasks.

In the model, workers are assumed to be *ex ante* identical and only to differ in terms of their realizations for idiosyncratic ability in each sector and occupation. However, in the data, workers differ in terms of a number of observed characteristics. To examine whether our findings for the change in the wage distribution above could be fully explained within the canonical model of the labor market in terms of skilled and unskilled workers, we now report results controlling for observable worker characteristics. Using our individual-level Census data on annual wages (available at twenty-year intervals from 1940-2000), we estimate the following Mincer regression across workers i in each year t separately:

$$\ln w_{it} = X_{it}\nu_t + u_{it} \tag{25}$$

where w_{it} is the annual wage; X_{it} are observable worker characteristics (education, gender, age and race); ν_t are coefficients that we allow to differ across years to capture changes in premia to these characteristics; and u_{it} is a stochastic error. Following Juhn, Murphy, and Pierce (1993) and AKK, we use the estimated residual ($\hat{u}_{it}$) as a measure of residual wage inequality after controlling for worker observables. In Figures A5-A6 of the web appendix, we show that we find the same pattern of results for 1940-2000 as in Figures 4-6 if we use the average residual wage for each occupation instead of the average wage. Therefore our findings using the relative wages of occupations in Figures 4-6 are not driven by changes in the distribution of observed worker characteristics or the premia to these observed characteristics.

To provide further evidence on how relative wages for different production tasks have changed over time after controlling for worker observables, we augment the Mincer regression (25) with measures of the production tasks undertaken within each occupation:

$$\ln w_{it} = X_{it}\nu_t + \mathbb{T}_{o(i)t}\zeta_t + u_{it} \tag{26}$$

where $\mathbb{T}_{o(i)t}$ are measures of the tasks undertaken by worker i within her occupation $o(i)$ at time t (either from the DOTs numerical scores or using the verbs from occupational descriptions); and ζ_t are wage premia for each task that we again allow to change over time.

In Panel A of Table 4, we estimate the Mincer regression (26) using the three AKK categories of non-routine, routine and manual tasks. In 1940, the first year for which we have the individual-level data on wages, we find a positive and statistically significant premium for non-routine tasks, a positive but insignificant premium for routine tasks, and a negative and statistically significant premium for manual tasks, consistent with the wage distribution results above. Over the period from 1940-2000, the premium for non-routine tasks increases substantially (by over 30 percent); the premium for manual tasks becomes statistically insignificant; and the premium for routine tasks rises from 1940-1960, before declining from 1960 onwards following the dissemination of the computer.

In Panel B of Table 4, we estimate the same specification using the eight divisions of the thesaurus from the panels of Figures 2-3 as our task measures. In 1940, we find the largest positive and statistically significant wage premium for the formation of ideas and the largest (in absolute value) negative and statistically significant wage premium for the manipulation of physical matter. From 1940-2000, the wage premium for the formation of ideas increases the most; the premium for individual volition also increases albeit from a lower base; and the premium for the manipulation of physical matter decreases the most. Combining these results with those above, we find that tasks involving the manipulation of physical matter experience a decline in both employment shares and relative wages, while those involving analytical and interactive tasks experience a rise in both employment shares and relative wages.

Finally, we use these regression estimates to examine the counterfactual implications of changes in task wage premia for wage inequality across occupations. We begin by using our estimates of the Mincer regression (26) with thesaurus task measures to generate fitted values for worker wages in 1940 and 2000.[38] We compute the fitted average wage for each occupation as the average of the fitted wages across all workers within that occupation. In the left panel of Figure 7, we show the cumulative distribution of employment shares across percentiles of the fitted occupation wage distribution in each year. Comparing this left panel with Figure 4, the fitted wage distributions for 1940 and 2000 are relatively successful in capturing the shift in the actual wage distribution between these two years. We next use our estimates to generate counterfactual values for worker wages in which we hold all worker characteristics and coefficients (education, gender, age and ethnicity) constant at their values in 2000 except for the task premia, which we set equal to their 1940 values. We compute the counterfactual wage for each occupation as the average of the counterfactual wages across all workers within that occupation. In the right panel of Figure 7, we show the cumulative distribution of employment shares across percentiles of the counterfactual occupation wage distribution (dashed black line), as well as the two fitted distributions for 1940 (solid black line) and 2000 (solid gray line) from the left panel. We find that the counterfactual distribution lies substantially closer to the 1940 fitted distribution than the 2000 fitted distribution, particularly at intermediate to high values for wages where the biggest shifts in the wage distribution occur. Therefore, after controlling for education and other observable worker characteristics, the changes in the wage premia for production tasks from 1940-2000 make a substantial contribution towards explaining the observed increase in wage inequality across occupations and the polarization of the occupational wage distribution from intermediate to high values of wages.

[38] We find similar results using other task measures, but focus on the thesaurus task measures because the fitted and counterfactual wage distributions provide a better fit to the data than those using the numerical scores.

6 Tasks and Technology

We now examine the extent to which the changes in task inputs documented in the previous section are related to the development of new technologies, as suggested by the model. Our identification strategy follows ALM in exploiting variation over time in the arrival of new technologies and variation across industries in the extent to which they use these new technologies. Our contribution is to use our much longer historical time period (1880-2000 compared to 1960-2000), to present results for a wider range of technologies (not only computers), and to use our new methodology for measuring individual tasks. In particular, we consider the following technologies that experienced substantial innovation during our sample period: (i) Office and computing machines (Computing), (ii) Electrical machinery (Electrical), (iii) Transport machinery (Transport), (iv) All machinery (Computing, Electrical, Transport and Other Machinery).

We examine whether industries that used these new technologies more intensively experienced larger changes in task inputs. Industry task intensity is measured as the employment-share-weighted average of the task intensity of occupations within that industry. Industry technology use is measured using time-invariant shares of inputs purchased from other industries, which in our baseline specification are 1947 input use shares. We estimate the following regression using observations on sectors s and years t from 1880-2000 for the within-industry relationship between inputs of task k and use of technology m:

$$\mathbb{T}_{skt} = \beta_{tkm} \left(\mathbb{S}_{sm} \times \mathbb{I}_{t} \right) + \eta_{s} + d_{t} + u_{st} \tag{27}$$

where $\mathbb{T}_{skt}$ is input of task k in sector s at time t; $\mathbb{S}_{sm}$ is the share of sector s's inputs that originate from industries that produce technology m; $\mathbb{I}_{t}$ is an indicator variable for year t; η_{s} are sector fixed effects; d_{t} are year dummies; and u_{st} is a stochastic error. We cluster the standard errors by sector to allow for serial correlation in the error term over time. We report standardized beta coefficients (scaled by variable standard deviations) for comparability across the different task and technology measures.

We estimate the regression specification (27) separately for a number of different task measures k, including both the numerical scores (Abstract, Routine and Manual) and our measures of individual production tasks. For each task measure k, we consider each of the four technologies m specified in (i)-(iv) above. The inclusion of sector and time fixed effects implies that this specification has a "differences-in-differences" interpretation. The sector fixed effects control for time-invariant heterogeneity across sectors (including the main effect of technology use $\mathbb{S}_{sm}$). The time fixed effects control from common changes over time in task input across all sectors. The key coefficient of interest is β_{tkm}, which captures the extent to which a sector that uses a technology m intensively experiences a differential change in inputs of task k over time t relative to other sectors. Since sector technology use $\mathbb{S}_{sm}$ is time invariant, the evolution of β_{tkm} over time reflects differential changes in the premium to this time-invariant sector characteristic.

The excluded year is 1880, so that the estimated β_{tkm} have the interpretation of changes relative to 1880.

To connect with existing research for recent decades, we begin by estimating (27) using numerical scores as our measures of tasks (non-routine, routine and manual) and office and computing machines as our measure of technology. While this existing research has largely focused on the post-1960 period, we report results from 1880-2000. In Figure 8, we display the resulting estimated beta coefficients $\hat{\beta}_{tkm}$ and 95 percent confidence intervals (clustered by sector) for the three numerical score measures. In line with existing research, we find that sectors that use office and computing machines intensively experience larger increases in inputs of non-routine tasks than other sectors in the second half of the twentieth century (top left panel). However, consideration of our longer time period and the comparison with other technologies yields a number of additional insights.

In particular, we find that industries that use office and computing machines intensively experience a smaller, but still discernible and statistically significant, increase in inputs of non-routine tasks relative to other industries in the late-nineteenth century. As noted earlier, the timing of the overall increase in inputs of non-routine tasks from 1880 onwards (Figures 1-3 and A1) aligns closely with the period of rapid diffusion of a cluster of information technologies centered on the typewriter and telephone (Figure A2). In Figure 8, we find that this increase in inputs of non-routine tasks from 1880 onwards is larger in industries that make intensive use of office and computing machines, tightening the connection with the diffusion of these information technologies. As shown in the web appendix, this differential trend in inputs of non-routine tasks in these industries is not present in the earlier 1860-1880 period (Figure A7). Together these findings provide empirical support for a historical literature that has argued that this earlier revolution in information technology played a central role in facilitating the development of systematic methods of management, managerial hierarchies and large corporations (e.g. Chandler 1977).[39]

These estimated changes in task inputs in response to improvements in information and communication technology include both the direct effects of these innovations (e.g. typewriters substitute for handwriting) and their indirect effects in facilitating a broader change in the organization of production (e.g. typewriters improve the recording and processing of information, which permits larger scale more capital-intensive modes of production). As shown in bottom left panel of Figure 8, we find that industries that make intensive use of office and computing machines experience a larger decline in inputs of manual tasks than other industries from the late-nineteenth century onwards. In contrast, as shown in the top

[39]Particular emphasis is placed on information and communication technologies in Yates (1989), including the upward flow of reports to inform executive decisions, the downward flow of orders to implement those decisions, and the lateral flow of information between the different divisions of the modern industrial corporation. These improvements in information and communication in turn facilitated the development of systematic methods of scientific management towards the end of the nineteenth and beginning of the twentieth centuries, including cost accounting, job cards, time clocks, inventory control, centralized purchasing and incentives wages, as argued in Nelson (1980, 1995).

right panel, the coefficient on inputs of routine tasks is relatively flat throughout the sample period as a whole, rising gradually until 1960, and then declining thereafter.

In Figure 9, we examine whether it is only office and computing technology that matters or whether other technologies also could have been important for task inputs. We display the estimated beta coefficients β_{tkm} from estimating (27) for all three numerical score measures and all four technology measures. As apparent from the figure, some of the largest estimated effects are for office and computing machinery. For non-routine tasks in the top left panel, we find positive estimated effects for office and computing machinery and electrical machinery, and negative estimated effects for transport machinery. This pattern of results suggests that office and computing machinery and electrical machinery are complementary to non-routine tasks, whereas transport machinery substitutes for non-routine tasks. For routine tasks in the top-right panel, we find the largest negative effects from electrical machinery and transport machinery (and the aggregate category of all machinery), consistent with the idea that these types of machinery substitute for routine tasks (as does office and computing machinery from 1960 onwards). In contrast, for manual tasks in the bottom left panel, we find large negative effects for office and computing machinery and positive effects for electrical and transport machinery. These results imply that office and computing machinery substitutes for manual tasks, whereas these two other categories of machinery are complementary towards manual tasks. We find the largest effects of transport equipment during the early twentieth century (with the dissemination of the automobile and early highway construction) and after 1960 (with the construction of the Interstate Highway System, as discussed in Lewis 1997 and Fernald 1999).

To provide finer resolution evidence on the impact of new technologies on task inputs, we re-estimate our regression specification (27) using our new methodology for measuring individual tasks, where verbs are our measure of task input $\mathbb{T}_{skt}$. Industry verb intensity is measured as the employment-share weighted average of the verb intensity of occupations within that industry. In the left panel of Table 3, we report the top twenty verbs with the greatest increases in task input from 1880-2000 in industries that use office and computing machines intensively (the top twenty values of β_{2000km} for verb k and office and computing technology m). In the right panel, we present the bottom twenty verbs with the greatest decreases in task input from 1880-2000. Although any one measure of production tasks is inevitably imperfect, we find an intuitive pattern. The verbs with the five greatest increases in task input in industries that use computers intensively are program, direct, test, use and engineer, which accord closely with priors about tasks for which computers are complementary. The verbs with the five greatest decreases in task input in computer-intensive industries are truck, serve, clean, pump and cook, which correspond to tasks less obviously connected to computer use. This pattern of results suggests that our findings using numerical scores above are indeed capturing individual production tasks that are closely related to computer use.

7 Tasks and Cities

While most existing research on tasks and the labor market considers the economy as a whole, we now show that the changes in task inputs established above differ systematically between urban and rural areas. In particular, we find a reversal in the types of tasks most concentrated in urban areas between 1880 and 2000. We show that this transformation in the nature of agglomeration helps to account for the larger changes in wage inequality observed in urban areas relative to rural areas over time. We show that these findings cannot be fully explained by the canonical model of the labor market in terms of skilled and unskilled labor, but instead remain even after controlling for worker observables. These results provide further evidence that the changes in wage inequality established above are indeed driven by changes in task premia, by exploiting a different source of variation between urban and rural areas.

7.1 Evidence from Aggregate Task Inputs

In the extension of the model to incorporate multiple locations, urban and rural areas can differ in employment shares in each sector and occupation, because of differences in the productivity of workers in performing tasks within each sector and occupation. We begin by providing evidence on the change in task inputs in urban areas relative to rural areas using the three numerical scores (non-routine, routine and manual). Figure 10 is analogous to Figure 1, but shows the employment-share-weighted average of the task measures for metro and non-metro areas separately (and uses the three AKK aggregations instead of the five ALM measures). As shown in the top left panel, we find a substantially larger increase in inputs of non-routine tasks over time in metro areas than in non-metro areas.[40] Whereas, in 1880, metro areas performed *less* non-routine tasks than non-metro areas, by 2000, this pattern is reversed and they undertook *more* non-routine tasks than non-metro areas. This differential increase between metro and non-metro areas in Figure 10 is 0.074, which corresponds to more than two thirds of the increase for non-metro areas of 0.100 and just under half of the increase for metro areas of 0.174. Therefore the urban-rural difference is large relative to the overall increase in non-routine inputs for the economy as a whole, which has been the subject of much research for the post-1960 period. We find that this urban-rural difference extends back to 1880 and is not only economically large but also statistically significant. Taking the difference in occupation employment shares between metro and non-metro areas, and regressing this difference from 1880-2000 on the three occupation numerical scores, we find positive and negative coefficients for non-routine and routine respectively, which are significant at the 5 percent level.[41]

[40] Recomputing the results in Figure 10 using the five ALM measures, we find an increasing concentration of non-routine inputs in metro areas for both non-routine analytic and non-routine interactive inputs.

[41] In Section A9.1 of the web appendix, we provide further evidence of this transformation in the nature of agglomeration, by using a different source of variation across metro areas of different population densities. Consistent with our results for metro versus non-metro areas, we find that employment in non-routine tasks has become increasingly concentrated in more densely-

As shown in the top right panel, inputs of routine tasks are relatively flat in both metro and non-metro areas over the sample period as a whole. However, whereas metro areas performed substantially *more* routine tasks than non-metro areas in 1880, this pattern is reversed by the end of the sample period and they undertook slightly *less* routine tasks than non-metro areas. Again the urban-rural difference is large relative to the overall change in task inputs for each group of locations. Both metro and non-metro areas experience substantial declines in inputs of routine tasks from 1980 onwards. Finally, as shown in the bottom left panel, the decline in inputs of manual tasks from 1880-2000 is larger in metro areas than in non-metro areas. In the first 20 years of the sample period, both sets of locations experience declining inputs of manual tasks at about the same rate, whereas from 1920 onwards, inputs of non-manual tasks continue to decline rapidly in metro areas, but decline less rapidly in non-metro areas.

While the model assumes that workers are *ex ante* identical, they differ in their observed characteristics in the data. To demonstrate the robustness of these findings to controlling for observed worker characteristics, we use our individual-level Census data on education (available from 1940 onwards) to estimate a linear probability model for the probability that individual i is located in a metro area in year t:

$$\mathbb{I}_{it}^{M} = X_{it}\iota_t + \mathbb{T}_{o(i)t}\varsigma_t + \eta_{s(i)} + \chi_{it}, \tag{28}$$

where $\mathbb{I}_{it}^{M}$ is an indicator variable that is equal to one if individual i is in a metro area; X_{it} are observable worker characteristics (education, gender, age and ethnicity); ι_t are coefficients that we allow to differ across years to reflect changes in the premia to these observed characteristics; $\mathbb{T}_{o(i)t}$ are the numerical score measures of the tasks undertaken by worker i within her occupation $o(i)$ at time t; ς_t are task premia that we again allow to change over time; $\eta_{s(i)}$ is a fixed effect for worker i's sector $s(i)$; and χ_{it} is a stochastic error. Although we focus on a linear probability model to facilitate the inclusion of sector fixed effects, we find a similar pattern of results using a Probit specification.

In Table A1 of the web appendix, we report the results of estimating (28) for each twenty-year period from 1940 onwards with different sets of controls (X_{it}). We begin by examining whether the changes in task inputs in Figure 10 simply reflect an increasing concentration of skilled workers in cities, as suggested in the literature on human capital externalities, including Berry and Glaeser (2005), Moretti (2004) and Davis and Dingel (2013). A closely-related hypothesis is that these findings capture a change in demographic composition , including for example an increased concentration of young "power couples" in cities, as suggested in Costa and Kahn (2000). Therefore Column (1) of Table A1 estimates (28) including our controls for observable worker characteristics (education, gender, age and ethnicity) and the three task measures (non-routine, routine and manual). We find an increase in the estimated non-routine coefficient

populated metro areas over time: there is little relationship between inputs of these tasks and population density in 1880 but a strong, positive and statistically significant relationship in 2000 (see Figure A9).

and a decline in the estimated routine coefficient over time, which implies that the observed changes in task inputs cannot be fully explained by changes in educational attainment or demographic composition.

We next examine whether the changes in task inputs in Figure 10 are purely attributable to a change in industry composition, with manufacturing moving out of urban areas towards lower-density locations, as argued in Desmet and Rossi-Hansberg (2009). Column (2) of Table A1 augments the specification from Column (1) with a full set of industry fixed effects. We continue to find an increased concentration of non-routine tasks in metro areas and a reduced concentration of routine tasks in metro areas over time, confirming that our findings are not driven by a change in industry composition. Finally, we consider the extent to which the changes in task inputs in Figure 10 can be explained purely by a shift from sectoral to functional specialization, with headquarters increasingly concentrated in urban areas and production plants dispersing to rural areas, as suggested in Duranton and Puga (2005), Rossi-Hansberg, Sarte, and Owens III (2009) and Ota and Fujita (1993). Column (3) of Table A1 further augments the specification from Column (2) with an indicator variable for occupations typically undertaken in headquarters.[42] Even in this specification including the full set of controls, we continue to find a similar pattern of results. Whereas non-routine tasks were statistically significantly *less* likely to be performed in metro areas in 1940, they were statistically significantly *more* likely to be undertaken in metro areas in 2000, confirming that we find a reversal in the pattern of task specialization in urban and rural areas over time, even after controlling for observed worker characteristics and sectoral and functional specialization.

7.2 Evidence from Individual Production Tasks

We next provide finer resolution evidence on the change in task inputs in urban areas relative to rural areas, using our new methodology for measuring individual tasks, as introduced in Section 4 above. We examine which verbs are most concentrated in metro areas by regressing the share of employment that is located in metro areas within a sector and occupation on the frequency with which a verb is used for that occupation. In particular, for each verb v and year t from 1880-2000, we estimate the following regression using observations across occupations o and sectors s for a given verb and year:

$$\text{MetroShare}_{ost} = \alpha_{vt}\text{VerbFreq}_{vo} + \eta_{vst} + \varepsilon_{ost}, \tag{29}$$

where MetroShare_{ost} is the share of employment within occupation o and sector s that is located in metro areas in year t; VerbFreq_{vo} is defined above in equation (22) for verb v and occupation o; η_{vst} are sector fixed effects for verb v and year t; and ε_{ost} is a stochastic error.

The coefficient of interest α_{vt} captures a conditional correlation: the correlation between occupations' shares of employment in metro areas and their frequency of use of verb v. The sector fixed effects for each

[42] See Section A5 of the web appendix for the list of occupations typically undertaken in headquarters.

verb v and year t (η_{vst}) control for differences across sectors in the frequency of verb use and for differences across sectors and over time in the concentration of employment in metro areas. Since VerbFreq_{vo} is time invariant, a rise in α_{vt} over time implies that employment in occupations using that verb is increasingly concentrating in metro areas within sectors over time.

In Panel A of Table 5, we report the top ten verbs with the highest standardized coefficient α_{vt} (the estimated coefficient scaled by the standard deviation of VerbFreq_{vo}) for each year.[43] In 1880, the verbs with the highest metro employment shares typically involve physical tasks such as "Ravel," "Sew," "Stretch" and "Thread." In contrast, by 1920, the top ten verbs include an increased number of clerical tasks, such as "Bill," "File," "Document," and "Record." By 1980 and 2000, there is a further change in the verbs most concentrated in metro areas towards analytical and interactive tasks, such as "Analyze," "Advise," "Confer" and "Report." These results for individual production tasks confirm that our findings above using numerical scores are indeed capturing a transformation in the individual production tasks most concentrated in urban areas.[44] While the typical urbanite in 1880 was likely to be employed in a manual task rearranging the physical world, their counterpart in 1940 was most frequently engaged in recording and processing information, and the modern city dweller typically performs tasks involving ideas, initiative and interaction. These findings highlight a transformation in the nature of agglomeration at the task level. In Panel B of Table 5, we report for comparison the bottom ten verbs with the lowest standardized coefficient α_{vt}. While we also find evidence of changes in the tasks least concentrated in metro areas (e.g. "Tread" appears from 1880-1960 and "Turn" appears from 1960-2000), these changes are typically smaller than for the tasks most concentrated in urban areas.

To quantify the shared characteristics of the tasks most concentrated in urban areas over time, we regress the share of employment that is located in metro areas within a sector and occupation on the frequency with which an occupation uses verbs from each thesaurus subdivision. In particular, for each thesaurus subdivision k and year t from 1880-2000, we estimate the following regression using observations across occupations o and sectors s for a given thesaurus subdivision and year:

$$\text{MetroShare}_{ost} = \beta_{kt}\text{ThesFreq}_{ko} + \eta_{kst} + \varepsilon_{ost}, \tag{30}$$

where MetroShare_{ost} is the share of employment in metro areas in occupation o, sector s and year t; ThesFreq_{ko} is defined above in equation (24) for thesaurus subdivision k and occupation o; η_{kst} are sector fixed effects for each thesaurus subdivision k and year t; and ε_{ost} is a stochastic error.

The coefficient of interest β_{kt} again captures a conditional correlation: the correlation between occu-

[43] We find a similar pattern of results just using the estimated coefficients instead of the estimated coefficients scaled by the standard deviation of VerbFreq_{vo}.

[44] In Section A8.2 of the web appendix, we show that we find a similar pattern of results using 1939 instead of 1991 occupational descriptions (see Table A2).

pations' shares of employment in metro areas and their frequency of use of verbs in thesaurus subdivision k. The sector fixed effects for each thesaurus subdivision k and year t (η_{kst}) control for differences across sectors in the frequency of use of thesaurus subdivisions and differences across sectors and over time in the concentration of employment in metro areas. Since ThesFreq_{ko} is time invariant, a rise in β_{kt} over time implies that employment in occupations using that subdivision of the thesaurus is increasingly concentrating in metro areas within sectors over time.

In Table 6, we report the estimation results for the thirty-nine Sections of the thesaurus. We calculate the standardized coefficient for each Section of the thesaurus (the estimated coefficient β_{kt} scaled by the variable's standard deviation) and report the ranking of these standardized coefficients in 1880 and 2000 as well the difference in rankings between these two years (1880 minus 2000).[45] Since the thesaurus Section with the highest standardized coefficient is assigned a rank of one, positive differences in rankings correspond to thesaurus categories that are becoming more concentrated in metro areas within sectors over time, whereas negative differences in rankings correspond to those that are becoming less concentrated in metro areas within sectors over time.

As shown in the table, we find a sharp change the relative ranking of thesaurus Sections involving the external world (Classes 1-3) and those involving the internal world of human beings (Classes 4-6). In 1880, four of the top five thesaurus sections most concentrated in metro areas involved the external world: Abstract Relations: Quantity (1.3), Space in General (2.1), Inorganic Matter (3.2) and Organic Matter (3.3). In contrast, in 2000, all of the top five thesaurus sections most concentrated in metro areas involved the internal world: Materials for Reasoning (4.1.3), Means of Communicating Ideas (4.2.3), Volition in General (5.1.1), Voluntary Action (5.1.3) and Possessive Relations (5.2.4). Therefore the transformation in the nature of agglomeration is away from tasks involving the manipulation of the physical world (space and matter) and towards analytical and interactive tasks (abstract thought and the communication of ideas). Consistent with a reversal in the types of tasks most concentrated in urban areas over time, we find a negative and statistically significant correlation between the concentration of thesaurus sections in metro areas in 1880 and 2000 (-0.50).

As with the results using numerical scores above, our findings for verbs and thesaurus sections are robust to estimating a linear probability model and including controls for observable worker characteristics (available from 1940 onwards), sector fixed effects and an indicator variable for headquarters occupations. Therefore our findings of a transformation in the most agglomerated tasks over time cannot be simply explained within the canonical model of the labor market in terms of the distinction between skilled and unskilled labor or by other observed worker characteristics.

[45] Again we find a similar pattern of results using just the estimated coefficient instead of the estimated coefficient scaled by the standard deviation of ThesFreq_{ko}.

7.3 Tasks and Wage Inequality

We now provide further evidence to tighten the connection between wage inequality and changes in task premia using variation between urban and rural areas. We first show that there is a greater change in wage inequality in urban areas than in rural areas in the second half of our sample period. We next show that this greater change in urban wage inequality can be largely explained by changes in task premia.

In Figure 11, we display the cumulative distribution of employment shares across percentiles of the occupation wage distribution in metro and non-metro areas separately. For 1880, we have a single measure of wages for each occupation, which we use for both metro and non-metro areas. For 1940 and 2000, we compute separate measures of average occupation wages for metro and non-metro areas using our individual-level Census micro data. As shown in the left panel, we find that the change in the distribution of employment across percentiles of the wage distribution is of approximately the same magnitude in metro and non-metro areas from 1880-2000. In contrast, we find much larger changes in the distribution of employment across percentiles of the wage distribution in metro areas than in non-metro areas from 1940-2000. Therefore, comparing Figures 4 and 11, most of the increase in overall wage inequality and the polarization of the overall wage distribution towards higher wages from 1940-2000 is driven by the change in the distribution of wages in metro areas.

To examine whether these patterns can be explained within the canonical model of the labor market in terms of skilled and unskilled workers, we again estimate the Mincer regression (25) using our individual Census data on wages from 1940 onwards, and construct a measure of residual wage inequality after controlling for worker observables (including education). In Figure A8 of the web appendix, we show that we find the same pattern of results for 1940-2000 as in Figure 11 if we use the average residual wage for each occupation instead of the average wage. Therefore we continue to find that most of the increase in overall wage inequality and the polarization of the overall wage distribution towards higher wages for the second half of our sample period is driven by the change in the distribution of wages in metro areas, even after controlling for observed skills and other worker characteristics.

To examine the contribution of task premia to the observed changes in wage inequality, we again augment the Mincer regression with measures of the production tasks undertaken within each occupation (as in equation (26)). We first use these estimates to construct fitted measures of average wages for each occupation for metro and non-metro areas separately. The top left panel of Figure 12 shows the cumulative distribution of employment shares across percentiles of the fitted occupation wage distribution for metro and non-metro areas in 1940 and 2000, using our thesaurus-based measures of tasks. As apparent from the figure, the fitted wage distributions are relatively successful in capturing the larger shift in the actual wage distribution in metro areas than in non-metro areas from 1940-2000.

We next use our estimates to generate counterfactual wages for each worker, in which we hold all worker characteristics and coefficients (education, gender, age and ethnicity) constant at their values in 2000 except for the task premia, which we set equal to their 1940 values. We compute the counterfactual wage for each occupation as the average of the counterfactual wages across all workers with that occupation for metro and non-metro areas separately. In the top right panel of Figure 12, we show the cumulative distribution of employment shares across percentiles of the counterfactual occupation wage distribution for metro areas, as well as the fitted distributions for metro areas for 1940 and 2000. In the bottom left panel of Figure 12, we show the cumulative distribution of employment shares across percentiles of the counterfactual occupation wage distribution for non-metro areas, as well as the fitted distributions for non-metro areas for 1940 and 2000. As apparent from the two panels, we find that the counterfactual distribution lies substantially closer to the 1940 fitted distribution than the 2000 fitted distribution for both metro and non-metro areas. Therefore, after controlling for education and other observable worker characteristics, we find that changes in task premia between 1940 and 2000 make a substantial contribution towards explaining the differential changes in wage inequality between metro and non-metro areas.

8 Conclusions

New technologies complement or substitute for particular tasks in ways that can be poorly summarized by aggregate measures of human capital or skills. We develop a new methodology for measuring these patterns of complementarity or substitutability at the level of *individual* production tasks. We use a Roy model of worker selection across sectors and occupations that require the performance of heterogeneous tasks to determine the effect of new technologies on employment shares and average wages through task effectiveness (task productivity) and the rate of return to human capital accumulation. We use our new methodology to provide the first evidence of changes in inputs of individual tasks in the U.S. economy over more than a century and the first application of a task-based approach to the labor market to the organization of economic activity in urban versus rural areas.

We show that the rise in inputs of non-routine tasks documented in recent research for the post-1960 period extends much further back in time than previously thought to an earlier information revolution in the late-nineteenth century, but accelerates following the dissemination of the computer in the late-twentieth century. Similarly, the decline in inputs of manual tasks established for recent decades is the continuation of a longer-term trend that stretches back to the late-nineteenth century. In contrast, inputs of routine tasks are relatively constant over the 1880-2000 period as a whole, but decline sharply after 1960 following the dissemination of the computer. Using our new methodology, we reveal substantial heterogeneity within these broad categories of tasks. We show that the individual production tasks that experienced the largest increases over our sample period are related to abstract thought and the formation of ideas (with the formation of ideas increasing by almost twice the growth for non-routine tasks as a whole). We demonstrate that the individual tasks that experienced the largest decreases over our sample period are related to the manipulation of physical matter (with inorganic matter falling by nearly twice the overall decline for manual tasks).

We establish that these changes in task inputs are explained by new technologies by exploiting variation over time in the arrival of new technologies and variation across industries in the extent to which they use these new technologies. We find the largest effects for office and commuting machinery, which has a positive impact on inputs of non-routine tasks that starts in the last two decades of the nineteenth century during the early information revolution (with the rise of telephones and typewriters), but accelerates in the latter part of the twentieth century (following the diffusion of the computer). We show that these new technologies have rich patterns of complementarity and substitutability with individual production tasks. The tasks most complementary to office and computing machines are again related to abstract thought and the communication of ideas (e.g. Program, Direct, Analyze, Design, Report). The tasks for which office and computing machines substituted the most again involved the manipulation of the physical world (e.g. Serve, Clean, Ticket, Machine, Deliver and Collect).

We demonstrate that changes in task wage premia can account for the decline in wage inequality from 1880-1940 and the rise in wage inequality from 1940-2000, and that these effects remain even after controlling for observed worker characteristics (such as education). To further strengthen the relationship between changes in wage inequality and changes in task premia, we use variation between urban and rural areas. We show that changes in task inputs over time are larger in urban areas than in rural areas, leading to a transformation in the nature of agglomeration. Whereas in 1880 the tasks most concentrated in urban areas involved the manipulation of the physical world (Thread, Stretch, Ravel, Sew), by 2000 they involved analytical and interactive tasks (Develop, Determine, Analyze, Review). We find that changes in task wage premia can account for a substantial proportion of the larger changes in wage inequality in urban areas than in rural areas, even after controlling for observed worker characteristics such as education.

While we concentrate on providing long-term evidence for the U.S. economy and in contrasting urban and rural areas, our framework for quantifying task complementarity and substitutability at a far higher resolution than has hitherto been possible lends itself to a rich range of further applications.

References

Abramitzky, R., L. P. Boustan, and K. Eriksson (2012): "Europe's Tired, Poor, Huddled Masses: Self-selection and Economic Outcomes in the Age of Mass Migration," *American Economic Review*, 102(5), 1832–1856.

——— (2014): "A Nation of Immigrants: Assimilation and Economic Outcomes in the Age of Mass Migration," *Journal of Political Economy*, 122(3), 467–506.

Acemoglu, D. (1998): "Why do New Technologies Complement Skills? Directed Technical Change and Wage Inequality," *Quarterly Journal of Economics*, 113, 1055–1090.

——— (2002): "Directed technical change," *Review of Economic Studies*, 69, 781–810.

Acemoglu, D., and D. Autor (2011): "Skills, Tasks and Technologies: Implications for Employment and Earnings," in *Handbook of Labor Economics*, ed. by O. Ashenfelter, and D. Card, vol. 4b, chap. 12, pp. 1043–1171. Elsevier, Amsterdam.

Ahlfeldt, G., S. Redding, D. Sturm, and N. Wolf (2015): "The Economics of Density: Evidence from the Berlin Wall," *Econometrica*, 83(6), 2127–2189.

Alonso, W. (1964): *Location and Land Use*. Harvard, Cambridge MA.

Atack, J., F. Bateman, and R. A. Margo (2004): "Skill Intensity and Rising Wage Dispersion in Nineteenth-Century American Manufacturing," *Journal of Economic History*, 64(1), 172–192.

Autor, D., L. Katz, and M. S. Kearney (2008): "Trends in U.S. Wage Inequality: Revising the Revisionists," *Review of Economics and Statistics*, 90(2), 300–323.

Autor, D., L. Katz, and A. Krueger (1998): "Computing Inequality: Have Computers Changed the Labor Market?," *Quarterly Journal of Economics*, 113, 1169–1214.

Autor, D. H. (2013): "The "Task Approach" to Labor Markets: An Overview," *Journal for Labour Market Research*, pp. 1–15.

Autor, D. H., and D. Dorn (2013): "The Growth of Low Skill Service Jobs and the Polarization of the U.S. Labor Market," *American Economic Review*, 103(5), 1553–1597.

Autor, D. H., F. Levy, and R. J. Murnane (2003): "The Skill Content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics*, pp. 1279–1333.

Bacolod, M., B. S. Blum, and W. Strange (2009a): "Skills in the City," *Journal of Urban Economics*, 65, 136–153.

——— (2009b): "Urban Interactions: Soft Skills Versus Specialization," *Journal of Economic Geography*, 9, 227–262.

Baum-Snow, N., and R. Pavan (2012): "Understanding the City Size Wage Gap," *Review of Economic Studies*, 79(1), 88–127.

——— (2013): "Inequality and City Size," *Review of Economics and Statistics*, 95(5), 1535–1548.

Baumgardner, J. (1988): "The Division of Labor, Local Markets, and Worker Organization," *Journal of Political Economy*, 96, 509–527.

Becker, G. S., and K. M. Murphy (1992): "The Division of Labor, Coordination Costs, and Knowledge," *Quarterly Journal of Economics*, 107(4), 1137–1160.

Berry, C., and E. L. Glaeser (2005): "The Divergence of Human Capital Levels across Cities," *Regional Science*, 84(3), 407–444.

Brunelle, C. (2013): "The Growing Economic Specialization of Cities: Disentangling Industrial and Functional Dimensions in the Canadian Urban System, 1971-2006," *Growth and Change*, 44(3), 443–473.

Bureau of Economic Analysis (2015): "Benchmark Input-Output Data: Historical SIC Data," Washington DC.

Burstein, A., E. Morales, and J. Vogel (2015): "Accounting for Changes in Between-Group Inequality," *NBER Working Paper*, 20855.

Chandler, A. D. J. (1977): *The Visible Hand: The Managerial Revolution in American Business*. Belknap Press, Cambridge.

Chin, A., C. Juhn, and P. Thompson (2006): "Technical Change and the Demand for Skills during the Second Industrial Revolution: Evidence from the Merchant Marine, 1891-1912," *Review of Economics and Statistics*, 88(3), 572–578.

Combes, P.-P., G. Duranton, L. Gobillon, and S. Roux (2012): "Sorting and Local Wage and Skill Distributions in France," *Regional Science and Urban Economics*, 42(6), 913–930.

Costa, D. L., and M. E. Kahn (2000): "Power Couples: Changes In The Locational Choice Of The College Educated, 1940-1990," *Quarterly Journal of Economics*, 115(4), 1287–1315.

DAVIS, D. R., AND J. I. DINGEL (2012): "A Spatial Knowledge Economy," *NBER Working Paper*, 18188.

——— (2013): "The Comparative Advantage of Cities," *NBER Working Paper*, 20602.

DAVIS, M. A., AND F. ORTALO-MAGNÉ (2011): "Housing Expenditures, Wages, Rents," *Review of Economic Dynamics*, 14(2), 248–261.

DESMET, K., AND E. ROSSI-HANSBERG (2009): "Spatial Growth and Industry Age," *Journal of Economic Theory*, 144(6), 2477–2502.

DURANTON, G. (1998): "Labor Specialization, Transport Costs and City Size," *Journal of Regional Science*, 38(4), 553–573.

DURANTON, G., AND H. JAYET (2011): "Is the Division of Labour Limited by the Extent of the Market? Evidence from French Cities," *Journal of Urban Economics*, 69(1), 56–71.

DURANTON, G., AND D. PUGA (2005): "From Sectoral to Functional Urban Specialisation," *Journal of Urban Economics*, 57(2), 343–370.

EATON, J., AND S. KORTUM (2002): "Technology, Geography, and Trade," *Econometrica*, 70(5), 1741–1779.

EECKHOUT, J., R. PINHEIRO, AND K. SCHMIDHEINY (2014): "Spatial Sorting," *Journal of Political Economy*, 122(3), 554–620.

ELLEN-POE, E. (2014): *Get Organized at Home!: How to Painlessly Find What You Are Looking For*. Sam Biondo, New York.

FERNALD, J. G. (1999): "Roads to Prosperity? Assessing the Link between Public Capital and Productivity," *American Economic Review*, 89(3), 619–638.

FUJITA, M., P. KRUGMAN, AND A. J. VENABLES (1999): *The Spatial Economy: Cities, Regions and International Trade*. MIT Press, Cambridge.

FUJITA, M., AND T. TABUCHI (1997): "Regional Growth in Postwar Japan," *Regional Science and Urban Economics*, 27, 643–670.

GASPAR, J., AND E. L. GLAESER (1998): "Information Technology and the Future of Cities," *Journal of Urban Economics*, 43(1), 136–156.

GLAESER, E. L., AND M. E. KAHN (2001): "Decentralized Employment and the Transformation of the American City," *NBER Working Paper*, 8117.

Glaeser, E. L., and J. E. Kohlhase (2003): "Cities, Regions and the Decline of Transport Costs," *Papers in Regional Science*, 83(1), 197–228.

Glaeser, E. L., A. M. Ponzetto, and K. Toblo (2011): "Cities, Skills and Regional Change," *NBER Working Paper*, 16934.

Glaeser, E. L., and M. G. Resseger (2009): "The Complementarity Between Cities and Skills," *NBER Working Paper*, 15103.

Glaeser, E. L., and A. Saiz (2003): "The Rise of the Skilled City," *NBER Working Paper*, 10191.

Goldin, C., and L. Katz (2008): *The Race between Education and Technology*. Harvard University Press, Cambridge MA.

Goos, M., and A. Manning (2007): "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain," *Review of Economics and Statistics*, 89(1), 118–133.

Goos, M., A. Salomons, and A. Manning (2014): "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring," *American Economic Review*, 104(8), 2509–26.

Gray, R. (2013): "Taking Technology to Task: The Skill Content of Technological Change in Early Twentieth Century United States," *Explorations in Economic History*, 50(3), 351–367.

Heckman, J. J., and B. E. Honore (1990): "The Empirical Content of the Roy Model," *Econometrica*, 58(5), 1121–1149.

Helsley, R. W., and W. C. Strange (2007): "Agglomeration, Opportunism and the Organization of Production," *Journal of Urban Economics*, 62, 55–75.

Henderson, J. V. (1974): "The Sizes and Types of Cities," *American Economic Review*, 64(4), 640–656.

——— (2003): "Marshall's Scale Economies.," *Journal of Urban Economics*, 55, 1–28.

Hendricks, L. (2011): "The Skill Composition Of U.S. Cities," *International Economic Review*, 52(1), 1–32.

Hounshell, D. A. (1985): *From the American System to Mass Production, 1800-1932: The Development of Manufacturing Technology in the United States*. Johns Hopkins University Press, Baltimore.

Hsieh, C.-T., E. Hurst, C. I. Jones, and P. J. Klenow (2013): "The Allocation of Talent and U.S. Economic Growth," *NBER Working Paper*, 18693.

Hüllen, W. (2003): *A History of Roget's Thesaurus: Origins, Development, and Design.* Oxford University Press, Oxford.

Hunt, H. A., and T. L. Hunt (1986): *Clerical Employment and Technological Change.* W. E. Upjohn Institute for Employment Research.

Jacobs, J. (1969): *The Economy of Cities.* Random House.

James, J. A., and J. S. Skinner (1985): "The Resolution of the Labor-scarcity Paradox," *Journal of Economic History*, 45(3), 513–540.

Juhn, C., K. M. Murphy, and B. Pierce (1993): "Wage Inequality and the Rise in Returns to Skill," *Journal of Political Economy*, 101(3), 410–442.

Katz, L. F., and R. A. Margo (2014): "Technical Change and the Relative Demand for Skilled Labor: The United States in Historical Perspective," in *Human Capital in History: The American Record*, ed. by L. P. Boustan, C. Frydman, and R. A. Margo, pp. 15–58. University of Chicago Press, Chicago.

Lagakos, D., and M. E. Waugh (2013): "Selection, Agriculture, and Cross-Country Productivity Differences," *American Economic Review*, 103(2), 948–80.

Lewis, T. (1997): *Divided Highways: Building the Interstate Highways, Transforming American Life.* Viking Penguin, New York.

Lin, J. (2011): "Technological Adaptation, Cities and New Work," *Review of Economics and Statistics*, 93(2), 554–574.

McFadden, D. (1974): "The Measurement of Urban Travel Demand," *Journal of Public Economics*, 3(4), 303–328.

Michaels, G. (2007): "The Division of Labor, Coordination and the Demand for Information Processing," London School of Economics, mimeograph.

Michaels, G., F. Rauch, and S. Redding (2012): "Urbanization and Structural Transformation," *Quarterly Journal of Economics*, 127, 535–586.

Mills, E. S. (1967): "An Aggregative Model of Resource Allocation in a Metropolitan Centre," *American Economic Review*, 57(2), 197–210.

Mokyr, J. (1992): *The Lever of Riches: Technological Creativity and Economic Progress.* Oxford University Press, Oxford.

MORETTI, E. (2004): "Human Capital Externalities in Cities," in *Handbook of Regional and Urban Economics*, ed. by J. V. Henderson, and J. Thisse, vol. 4, pp. 2244–2289. Elsevier North Holland, Amsterdam.

MUTH, R. (1969): *Cities and Housing*. University of Chicago Press, Chicago.

NELSON, D. (1980): *Frederick W. Taylor and the Rise of Scientific Management*. University of Wisconsin Press, Madison.

——— (1995): *Managers and Workers: Origins of the Twentieth-Century Factory System in the United States, 1880-1920*. University of Wisconsin Press, Madison.

OTA, M., AND M. FUJITA (1993): "Communication Technologies and Spatial Organization of Multi-unit Firms in Metropolitan Areas," *Regional Science and Urban Economics*, 23, 695–729.

PHISTER, M. J. (1979): *Data Processing Technology and Economics*. Digital Press.

PIKETTY, T., AND E. SAIZ (2003): "Income Inequality in the United States, 1913-1998," *Quarterly Journal of Economics*, 118(1), 1–39.

PIKETTY, T., E. SAIZ, AND G. ZUCMAN (2014): "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data," *NBER Working Paper*, 20625.

PRESTON, S. H., AND M. R. HAINES (1991): *Fatal Years: Child Mortality in Late Nineteenth-Century America*. Princeton University Press, Princeton, NJ.

ROGET, P. M. (1911): *Thesaurus of English Words and Phrases*. Evergreen Review, Inc., Computer Readable Version of 1911 Edition.

ROSSI-HANSBERG, E., P.-D. SARTE, AND R. OWENS III (2009): "Firm Fragmentation and Urban Patterns," *International Economic Review*, 50(1), 143–186.

ROY, R. D. (1951): "Some Thoughts on the Distribution of Earnings," *Oxford Economic Papers*, 3, 135–146.

RUGGLES, STEVEN, J., T. ALEXANDER, K. GENADEK, R. GOEKEN, M. B. SCHROEDER, AND M. SOBEK (2010): "Integrated Public Use Microdata Series," University of Minnesota, Minneapolis.

SPITZ-OENER, A. (2006): "Technical Change, Job Tasks, and Rising Educational Demands: Looking outside the Wage Structure," *Journal of Labor Economics*, 24(2), 235–270.

U. S. DEPARTMENT OF LABOR (1939): *Dictionary of Occupational Titles*. U.S. Department of Labor, Washington D.C.

——— (1949): *Dictionary of Occupational Titles*. U.S. Department of Labor, Washington D.C.

——— (1991): *Dictionary of Occupational Titles*. U.S. Department of Labor, Washington D.C.

Yates, J. (1989): *Control Through Communication: The Rise of System in American Management*. Johns Hopkins Press, Baltimore.

Figure 1: Task Input by Numerical Score over Time

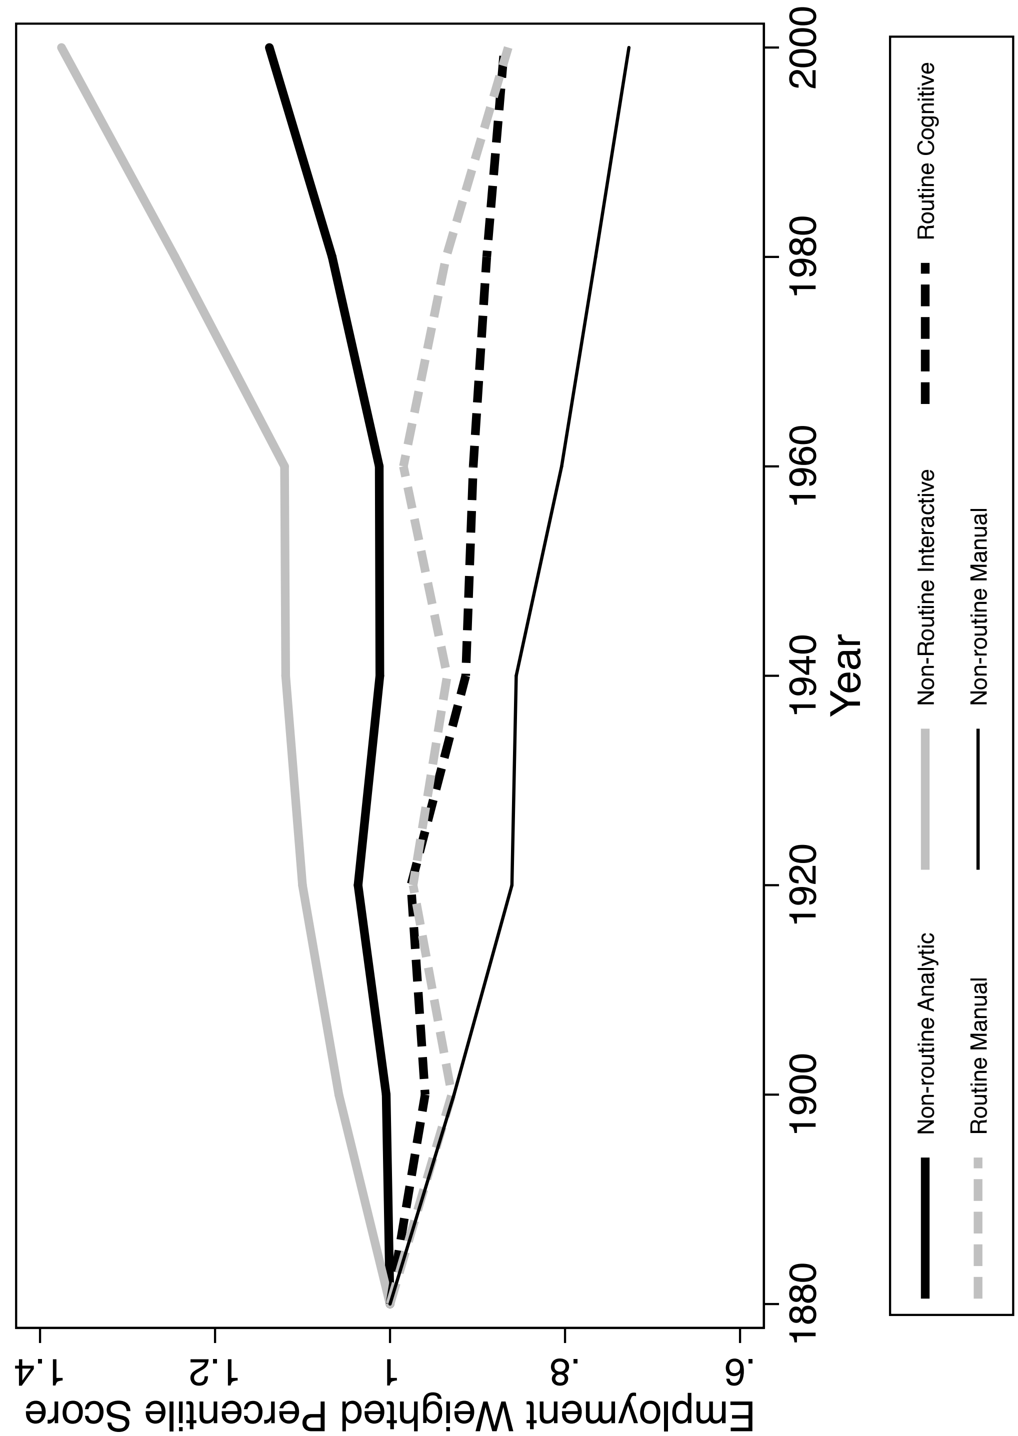

Note: Employment-weighted average of occupation numerical scores summarizing job requirements (non-routine analytic, non-routine interactive, routine cognitive, routine manual and non-routine manual) for each year, expressed as an index that equals one in 1880. Occupation numerical scores from the Dictionary of Occupational Titles (DOTs) for 1991. The time-invariant numerical score for each occupation is converted into the percentile of its distribution across occupations. Employment in each occupation and year is measured using IPUMS population census data for each twenty-year interval from 1880-2000.

Figure 2: Task Input by Thesaurus Section Over time

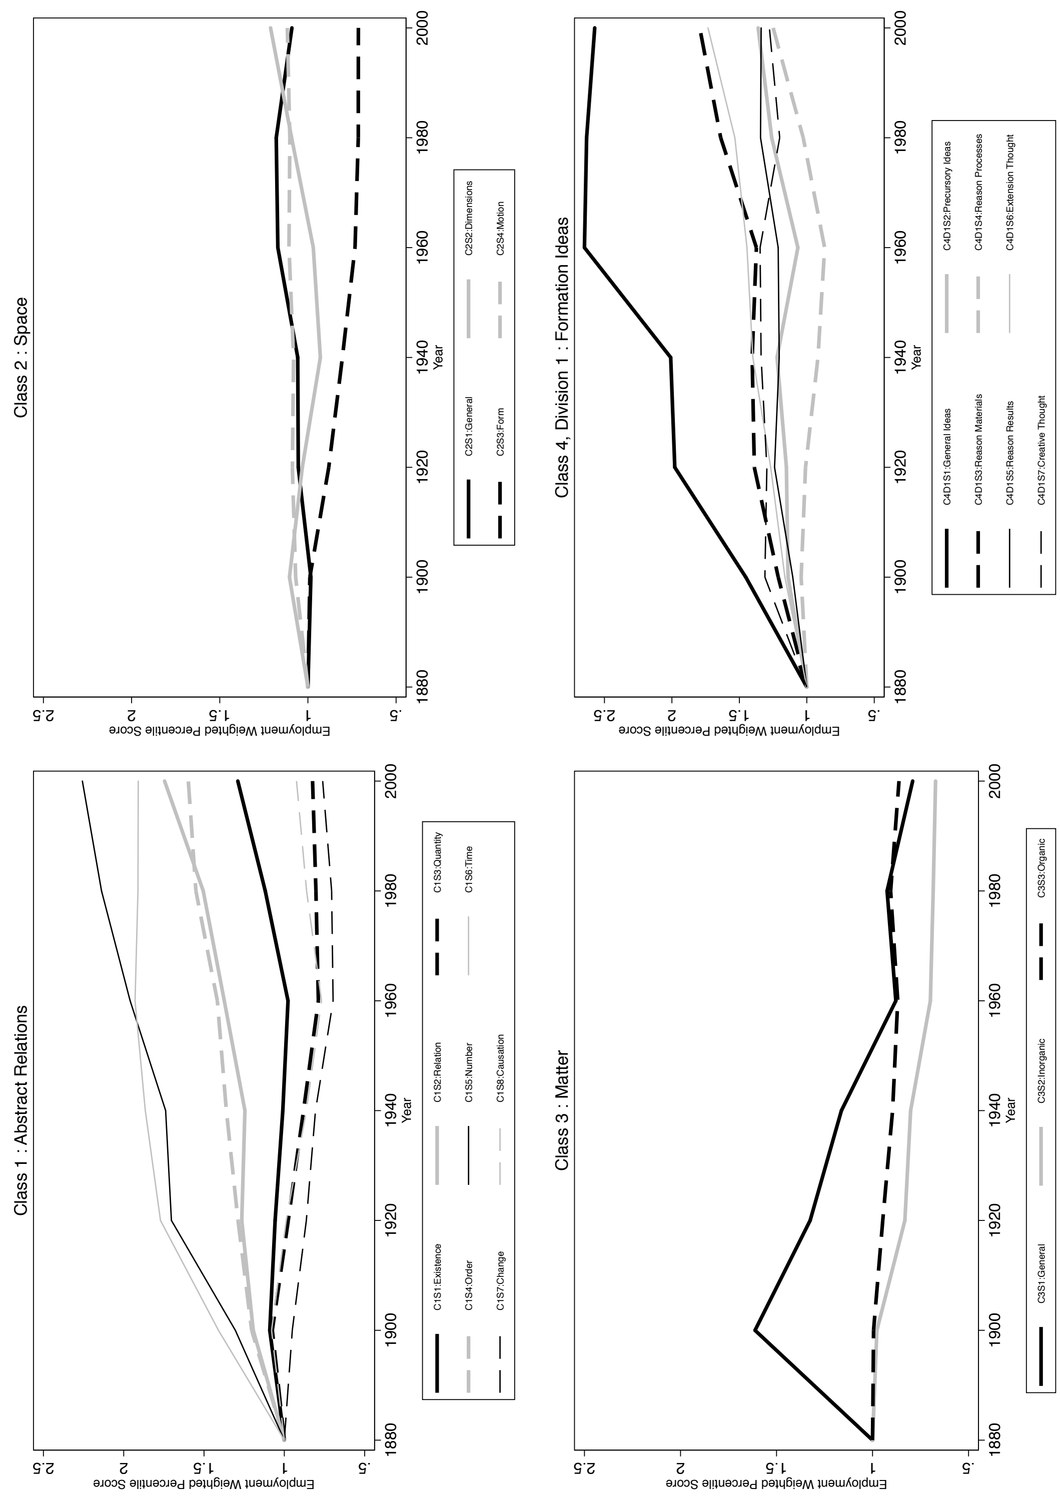

Note: Employment-weighted average for each year of thesaurus task k measure for occupation o (employment-weighted average of ThesFreq$_{ko}$ from equation (24)), expressed as an index that equals one in 1880. Each time-invariant thesaurus task measure is converted into the percentile of its distribution across occupations. Employment in each occupation and year is measured using IPUMS population census data for each twenty-year interval from 1880-2000.

Figure 3: Task Input by Thesaurus Section Over time (Continued)

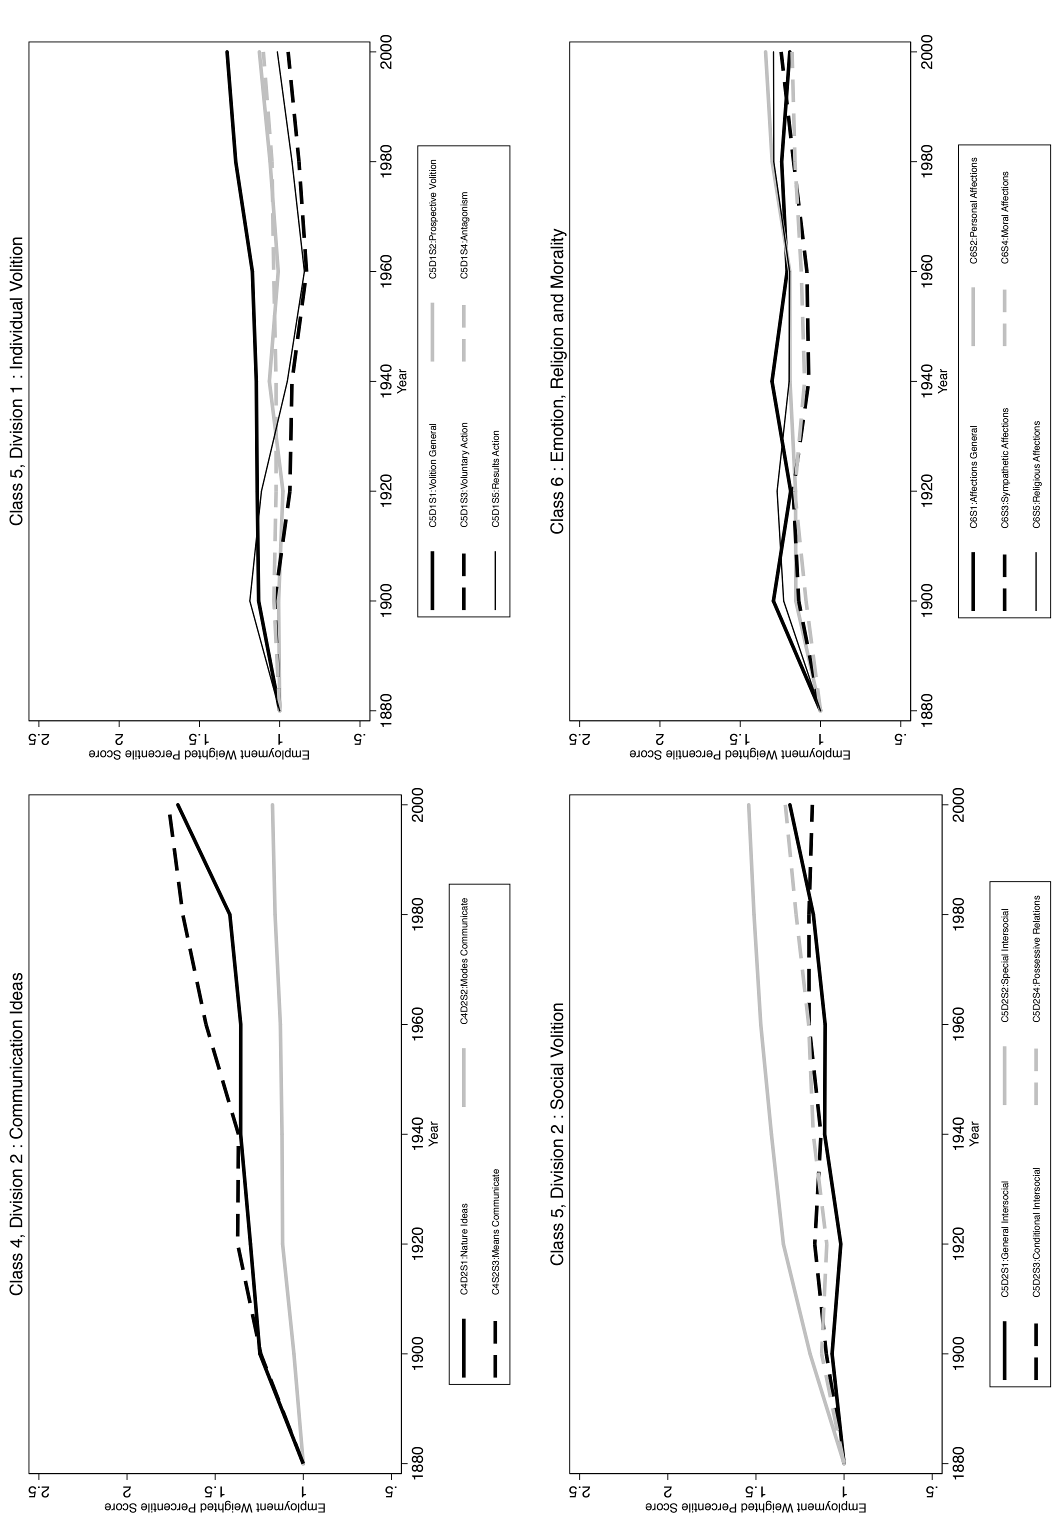

Note: Employment-weighted average for each year of thesaurus task k measure for occupation o (employment-weighted average of ThesFreq_{ko} from equation (24)), expressed as an index that equals one in 1880. Each time-invariant thesaurus task measure is converted into the percentile of its distribution across occupations. Employment in each occupation and year is measured using IPUMS population census data for each twenty-year interval from 1880-2000.

Figure 4: Cumulative Distribution of Occupation Employment Across Percentiles of the Occupation Wage Distribution (1880, 1940 and 2000)

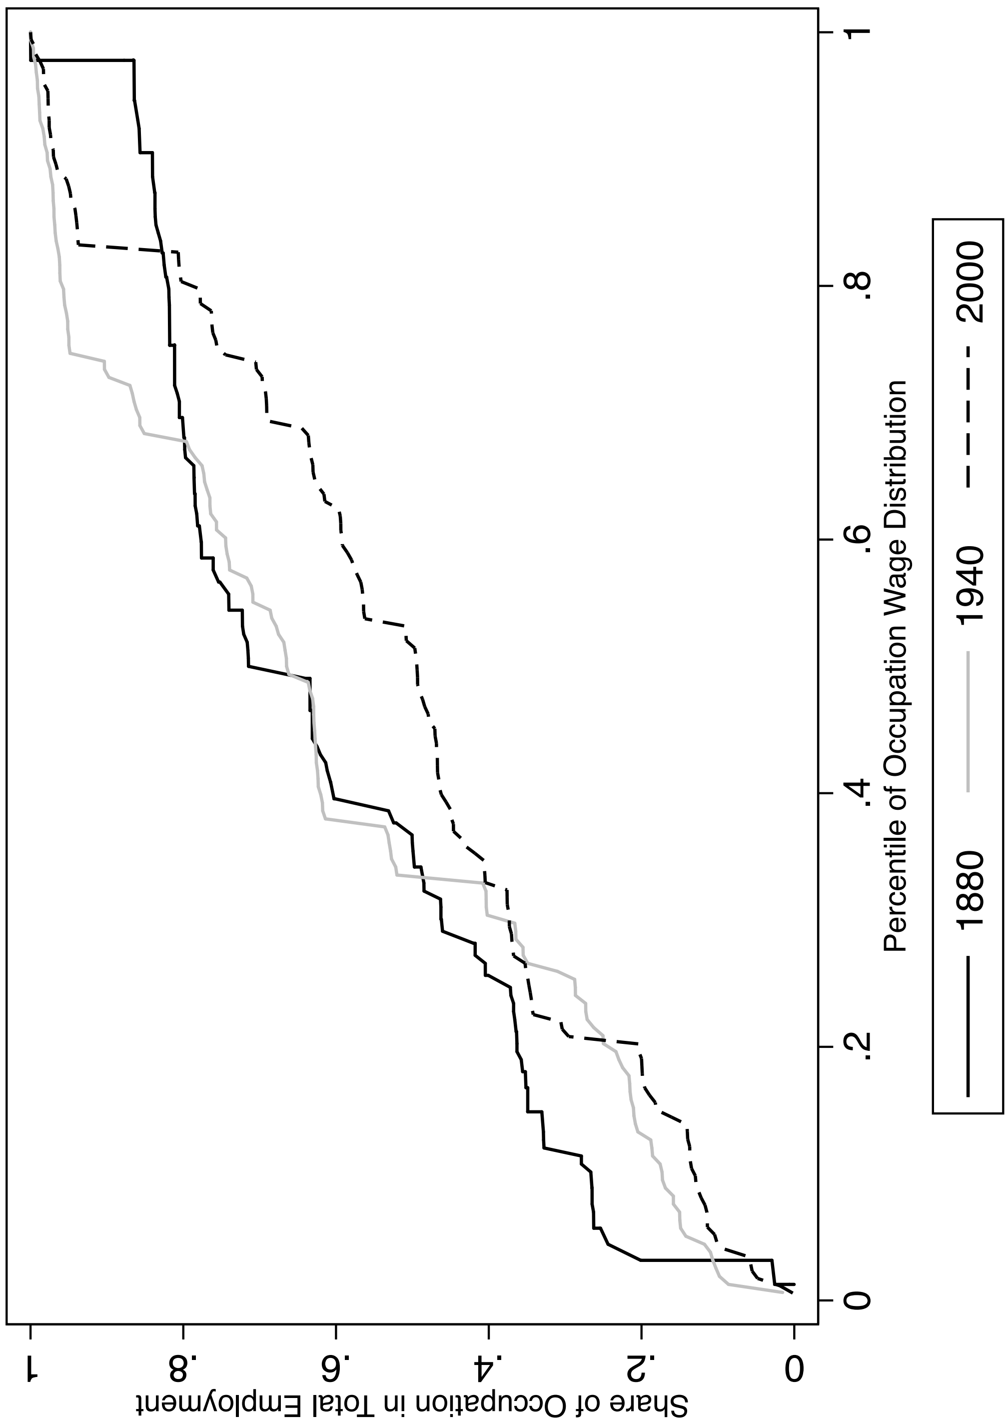

Note: On the horizontal axis, occupations are sorted in each year according to their percentile of the occupation wage distribution in that year. The vertical axis shows the cumulative share of the sorted occupations in total employment in that year. Occupation employment (for all years) and average occupation wages (for 1940 and 2000) are measured using the IPUMs population census data. Average occupation wages for 1880 are from Preston and Haines (1991), as used in Abramitzky, Boustan, and Eriksson (2012, 2014).

Figure 5: Cumulative Distribution of Occupation Employment Across Percentiles of the Task Distribution (1880, 1940 and 2000)

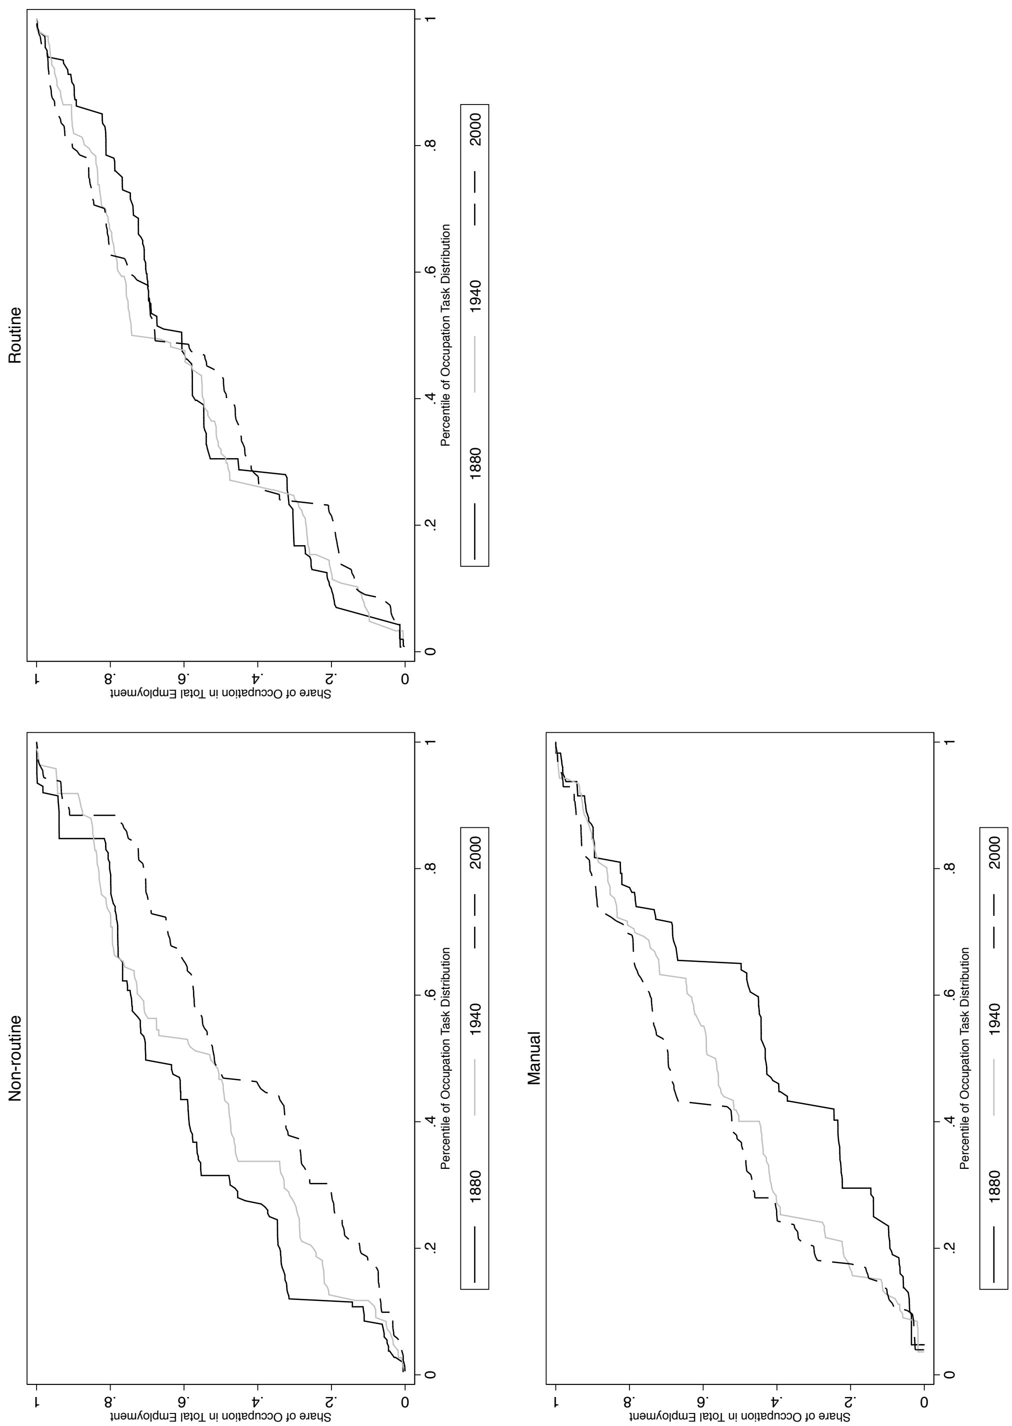

Note: On the horizontal axis, occupations are sorted according to their percentile of the occupation task distribution, as measured using the numerical scores from the Dictionary of Occupational Titles (DOTs) for 1991. “non-routine” is (non-routine analytic+non-routine interactive)/2; “routine” is (routine cognitive+routine manual)/2; and “manual” is non-routine manual. The vertical axis shows the cumulative share of the sorted occupations in total employment in each year. Occupation employment is measured using the IPUMs population census data for each year.

Figure 6: Cumulative Distribution of Tasks Across Percentiles of the Occupation Wage Distribution

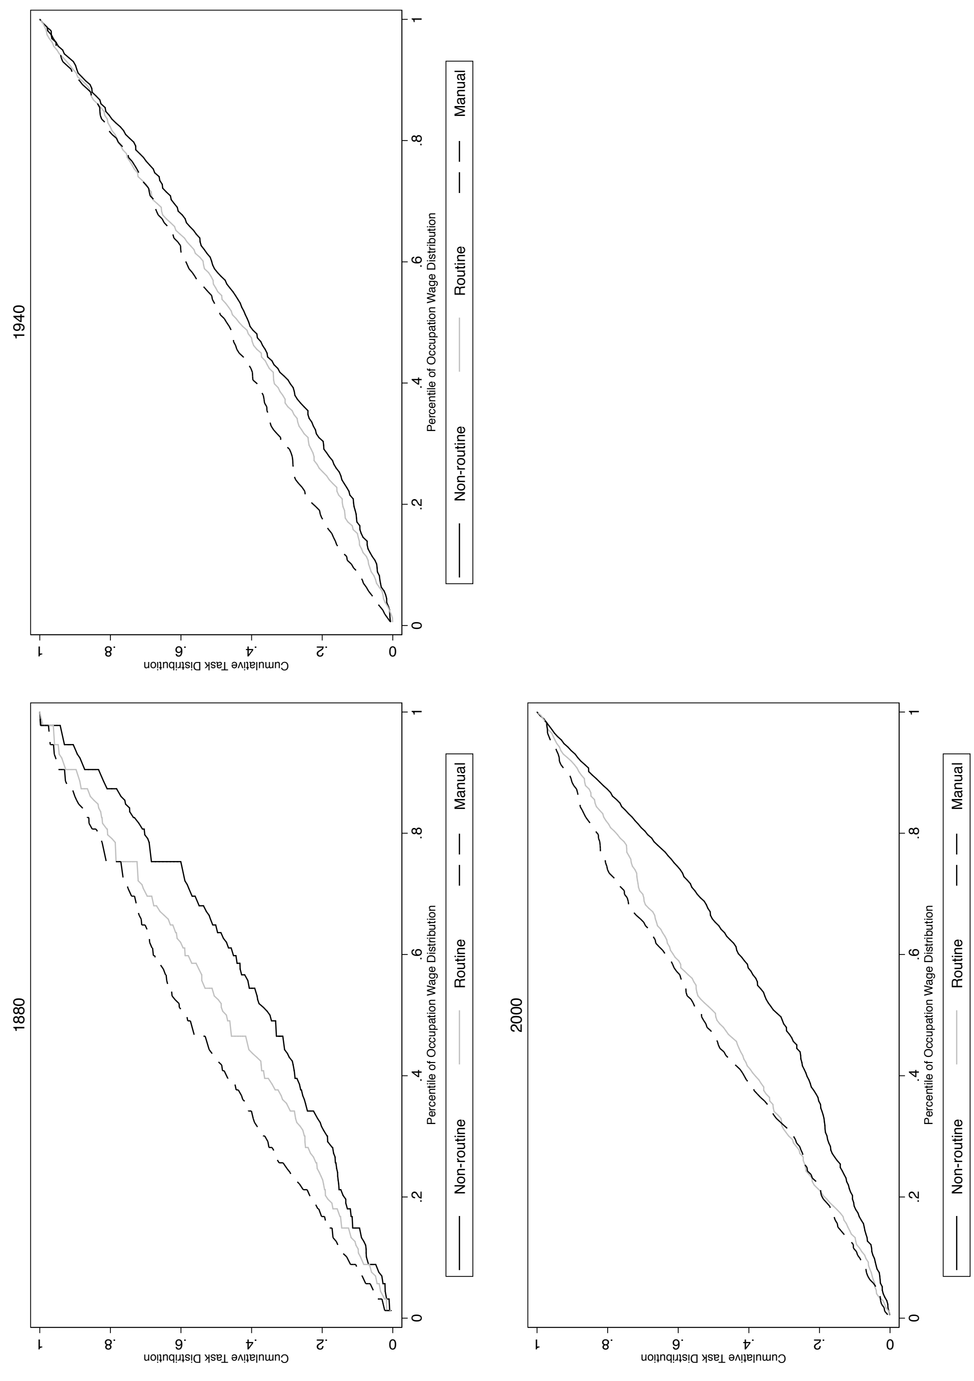

Note: On the horizontal axis, occupations are sorted in each year according to their percentile of the occupation wage distribution in that year. The vertical axis shows the cumulative task distribution of the sorted occupations (the cumulative sum of the percentile numerical scores for the sorted occupations, scaled to add up to one). Numerical scores from the Dictionary of Occupational Titles (DOTs) for 1991. Each time-invariant numerical score is converted into the percentile of its distribution across occupations. "Non-routine" is (non-routine analytic+non-routine interactive)/2; "routine" is (routine cognitive+routine manual)/2; and "manual" is non-routine manual. Average occupation wages for 1940 and 2000 are measured using the IPUMs population census data. Average occupation wages for 1880 are from Preston and Haines (1991), as used in Abramitzky, Boustan, and Eriksson (2012, 2014).

Figure 7: Actual and Counterfactual Distribution of Wages Across Occupations

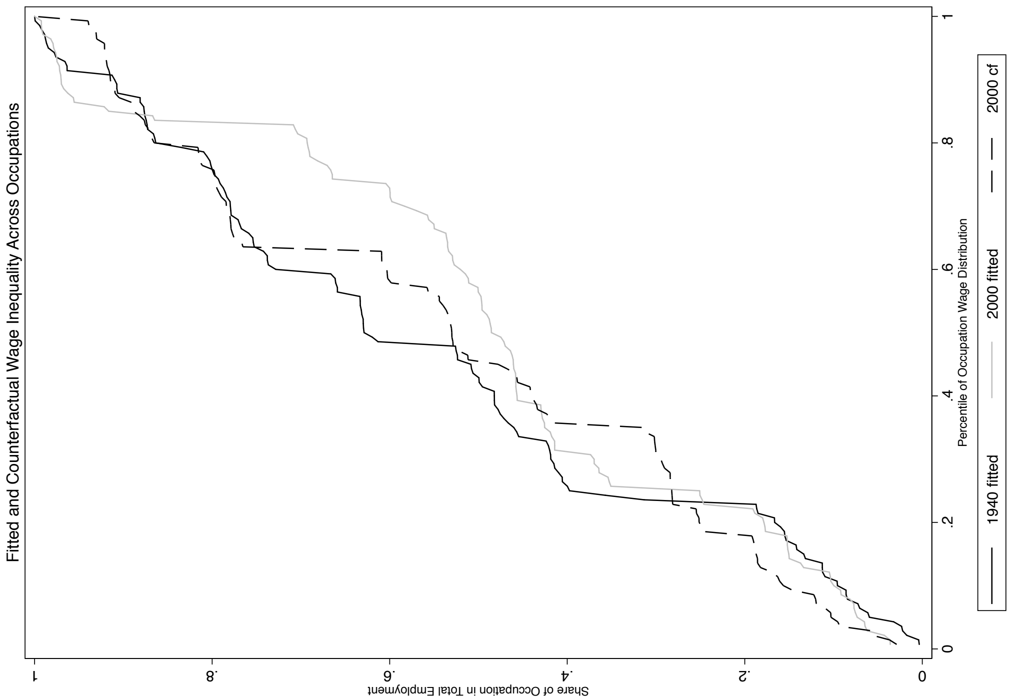

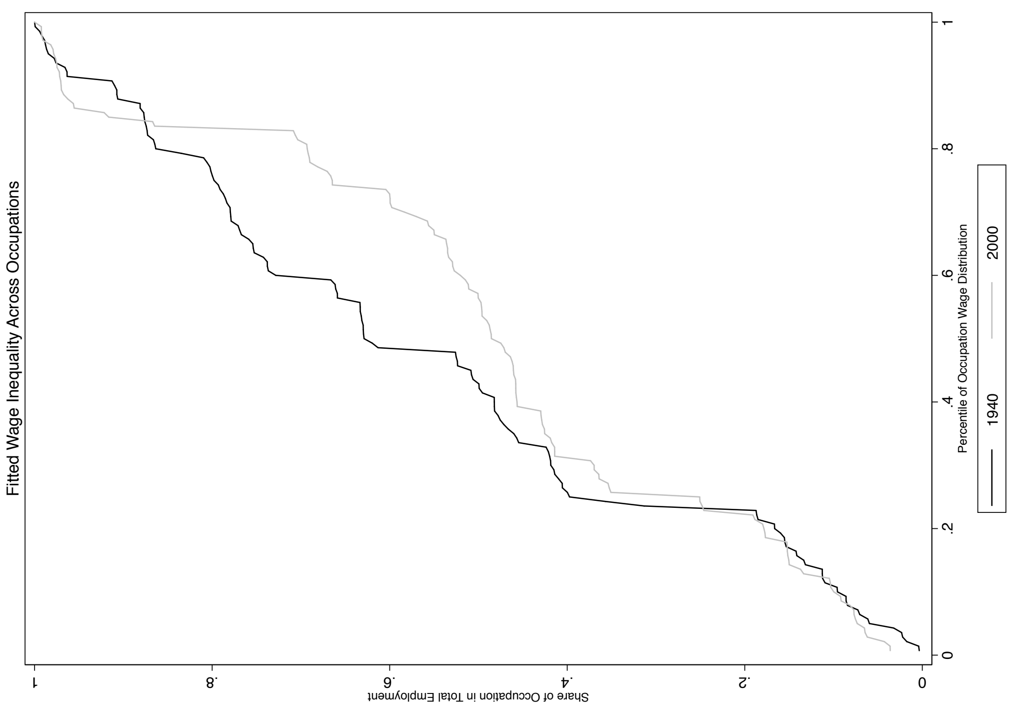

Note: On the horizontal axis, occupations are sorted in each year according to their percentile of the occupation actual, fitted or counterfactual wage distribution. The vertical axis shows the cumulative share of the sorted occupations in total employment in that year. Employment and average actual wages for each occupation are measured using the IPUMs population census data. The fitted [counterfactual] wage for each occupation is the average fitted [counterfactual] wage for each worker within that occupation. The fitted wage for each worker is from the Mincer regression (26). The counterfactual wage for each worker in 2000 equals the fitted wage in 2000, except that it uses the estimated 1940 task wage premia instead of the estimated 2000 task premia.

Figure 8: Task Input and Industry Office and Computing Machinery Use

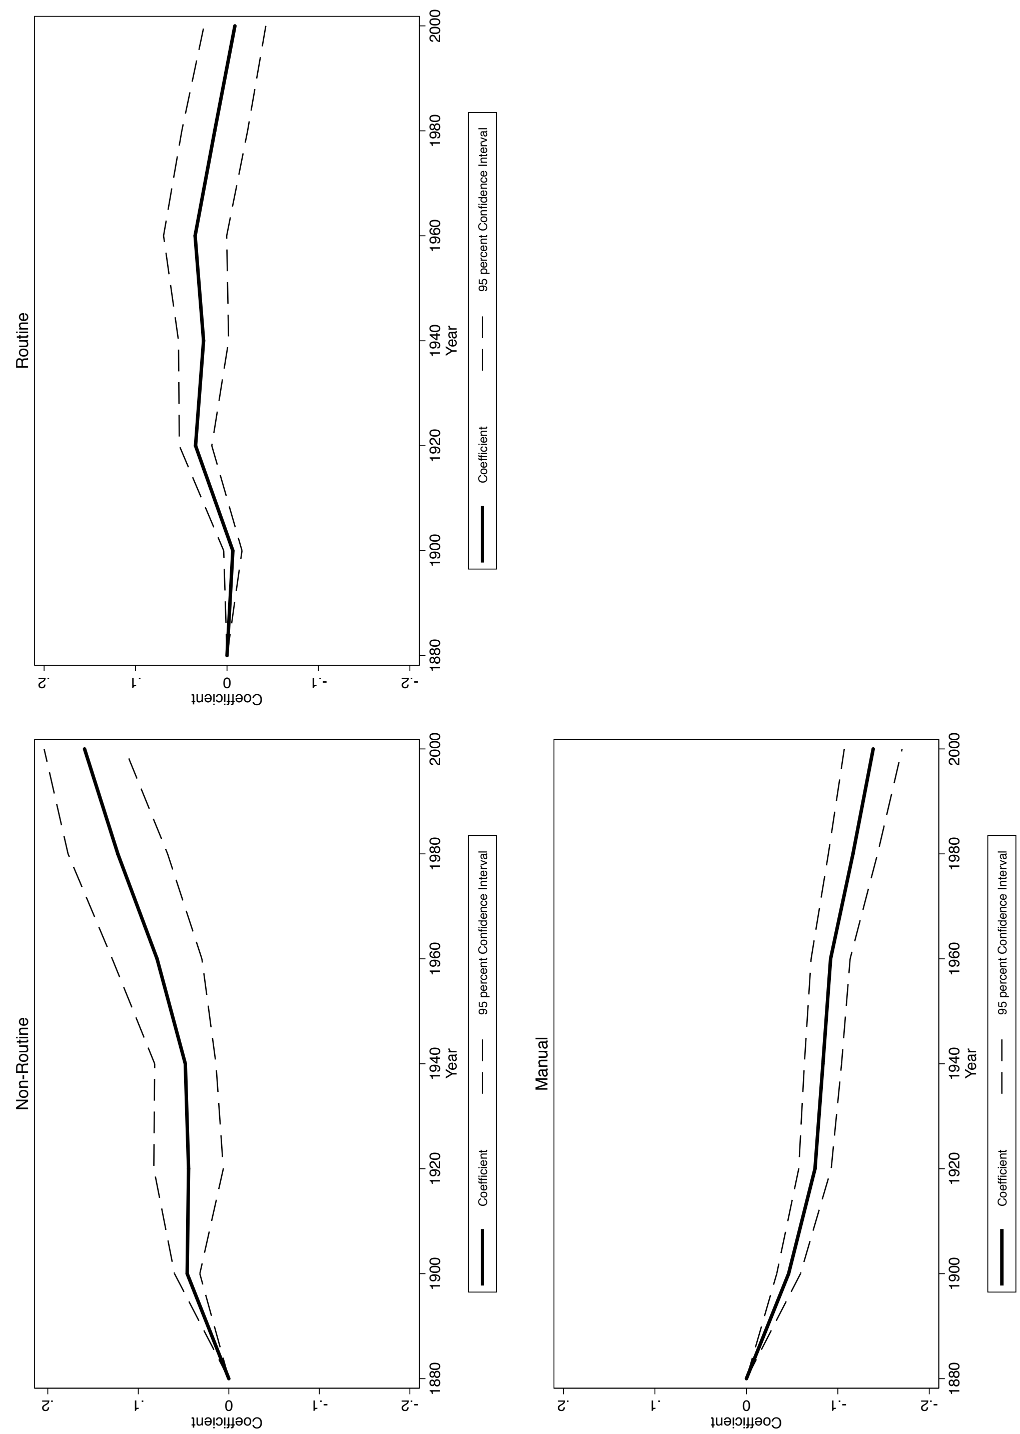

Note: Estimated coefficients (β_{tkm}) and 95 percent confidence intervals from the regression (27) of industry inputs of task k (non-routine, routine and manual) on a time-invariant measure of industry use of office and computing machinery m interacted with dummies for year t. Observations are industries and years. 1880 is the excluded year. Standard errors are heteroskedasticity robust and clustered on industry. Industry task inputs are the employment-weighted average of the non-routine, routine and manual numerical scores for each occupation from the Dictionary of Occupational Titles (DOTs) for 1991. Employment is measured using IPUMS population census data for each twenty-year interval from 1880-2000.

Figure 9: Task Input and All Technologies

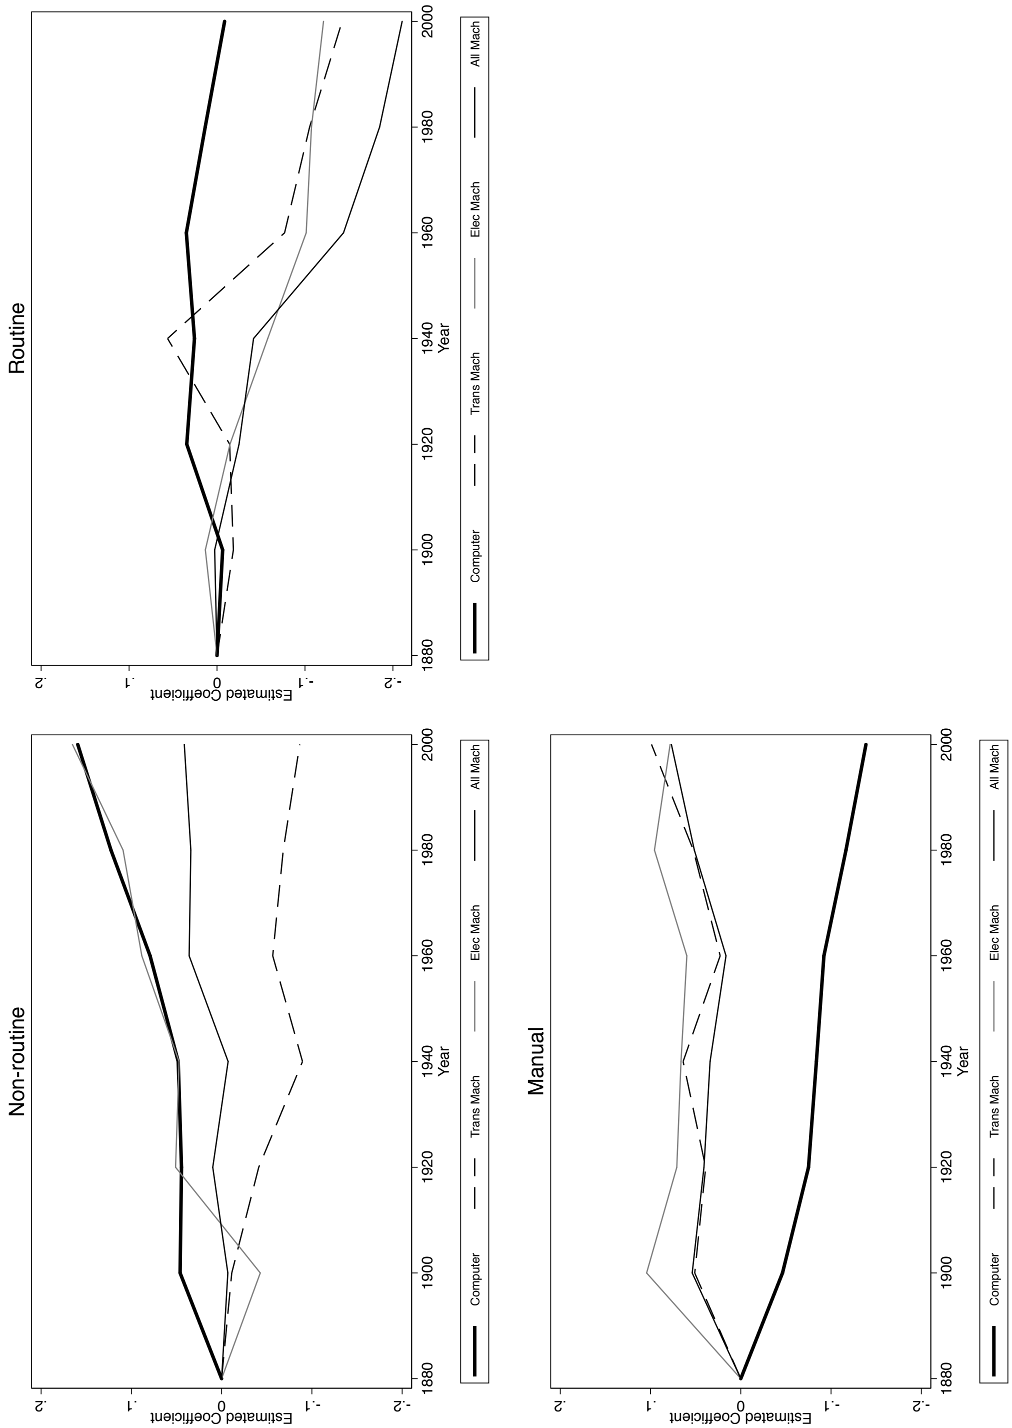

Note: Estimated coefficients (β_{tkm}) from the regression (27) of industry inputs of task k (non-routine, routine and manual) on time-invariant measures of industry use of technology m (office and computing machinery, electrical machinery, transport machinery, and all machinery) interacted with dummies for year t. Observations are industries and years. 1880 is the excluded year. Standard errors are heteroskedasticity robust and clustered on industry. Industry task inputs are the employment-weighted average of the non-routine, routine and manual numerical scores for each occupation from the Dictionary of Occupational Titles (DOTs) for 1991. Employment is measured using IPUMS population census data for each twenty-year interval from 1880-2000.

Figure 10: Metro and Non-metro Task Input by Numerical Score over Time

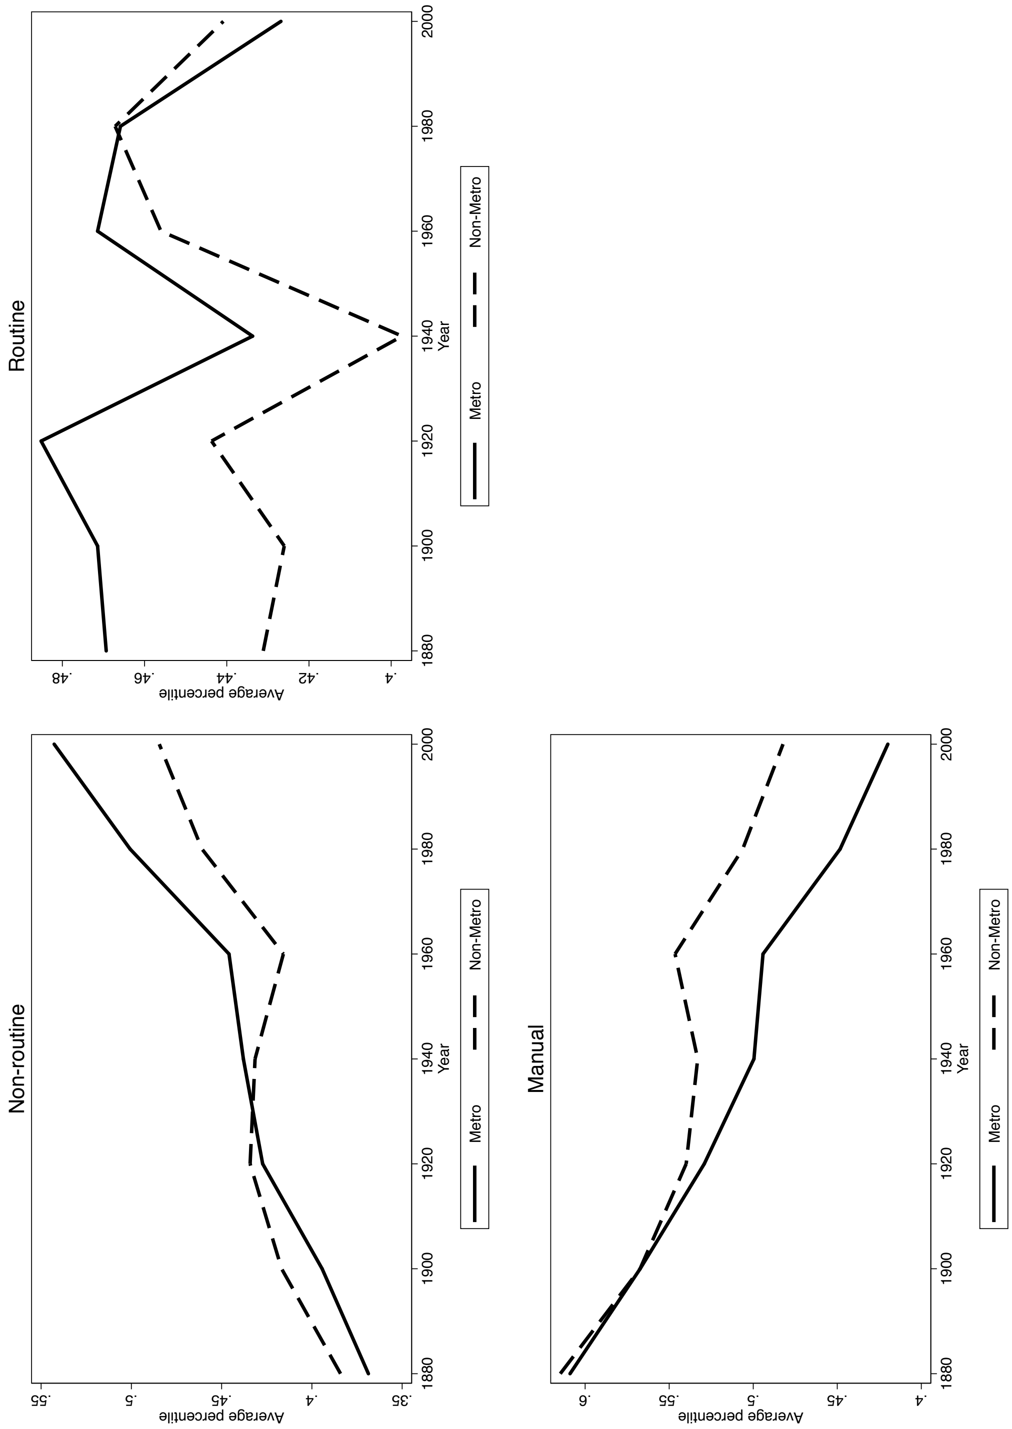

Note: Employment-weighted average for metro and non-metro areas separately of numerical scores summarizing job requirements for each occupation. "Non-routine" is (non-routine analytic+non-routine interactive)/2, "routine"' is (routine cognitive+routine manual)/2, and "manual" is non-routine manual. Numerical scores from the Dictionary of Occupational Titles (DOTs) for 1991. Each time-invariant numerical score is converted into the percentile of its distribution across occupations. Employment for metro and non-metro areas in each occupation and year is measured using IPUMS population census data for each twenty-year interval from 1880-2000.

Figure 11: Cumulative Distribution of Occupation Employment Across Percentiles of the Occupation Wage Distribution in Metro and Non-metro Areas (1880, 1940 and 2000)

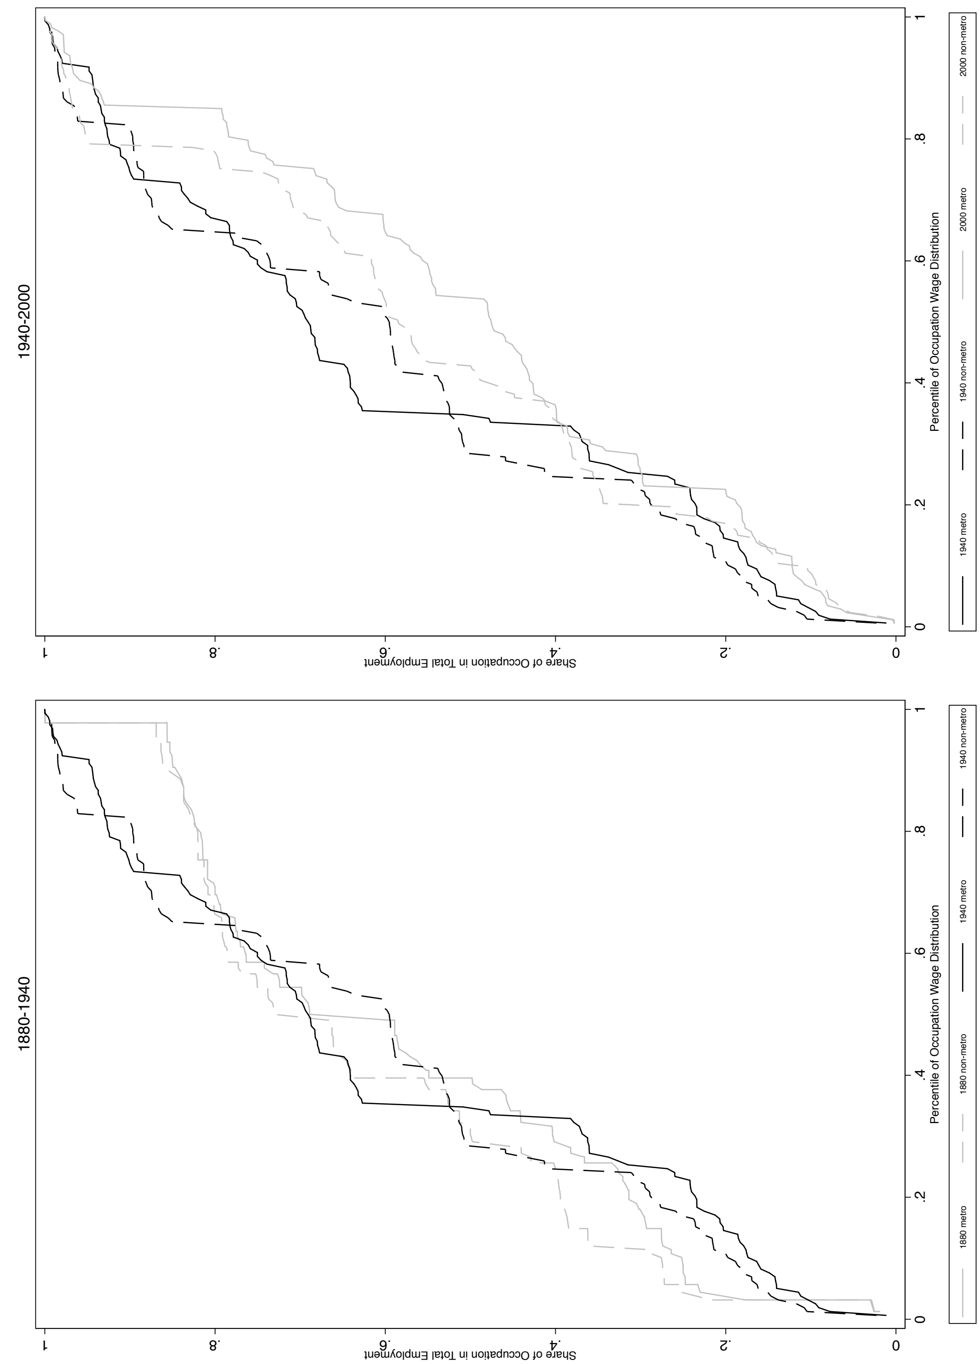

Note: On the horizontal axis, occupations are sorted in each year according to their percentile of the occupation wage distribution for metro or non-metro areas in that year. The vertical axis shows the cumulative share of the sorted occupations in total employment in that year for metro and non-metro areas. Occupation employment (for all years) and average occupation wages (for 1940 and 2000) for metro and non-metro areas are measured using the IPUMs population census data. Average occupation wages for 1880 are the same in metro and non-metro areas, and are from Preston and Haines (1991), as used in Abramitzky, Boustan, and Eriksson (2012, 2014).

Figure 12: Actual and Counterfactual Distribution of Wages Across Occupations for Metro and Non-metro Areas (1940 and 2000)

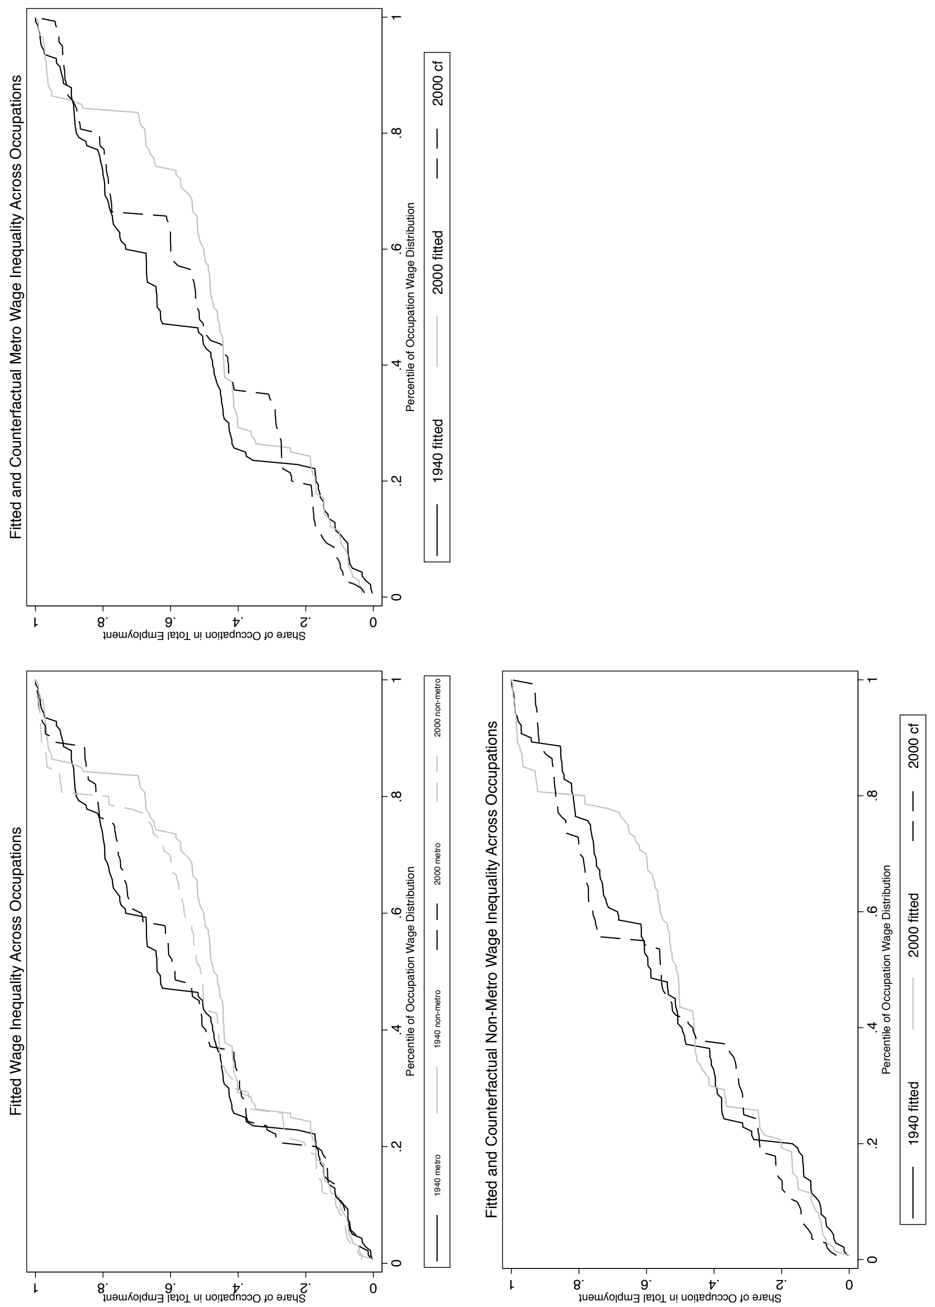

Note: On the horizontal axis, occupations are sorted in each year according to their percentile of the occupation actual, fitted or counterfactual wage distribution for metro and non-metro areas. The vertical axis shows the cumulative share of the sorted occupations in total employment in that year for metro and non-metro areas. Employment and average actual wages for each occupation are measured using the IPUMs population census data. The fitted [counterfactual] wage for each occupation is the average of fitted [counterfactual] wage for each worker within that occupation. The fitted wage for each worker is from the Mincer regression (26). The counterfactual wage for each worker in 2000 equals the fitted wage in 2000, except that it uses the estimated 1940 task wage premia instead of the estimated 2000 task premia.

Table 1: Top Twenty Tasks for each DOTs Numerical Score

	(1)		(2)		(3)		(4)		(5)	
Rank	Non-routine Interactive		Non-routine Analytic		Routine Cognitive		Routine Manual		Non-Routine Manual	
1	Analyze	0.51	Analyze	0.65	Use	0.53	Use	0.48	Prevent	0.42
2	Develop	0.50	Conduct	0.60	Master	0.51	Sketch	0.40	Move	0.34
3	Conduct	0.48	Plan	0.58	Cut	0.42	Master	0.39	Line	0.33
4	Plan	0.47	Develop	0.56	Fit	0.41	Measure	0.38	Climb	0.31
5	Direct	0.47	Research	0.56	Title	0.40	Machine	0.37	Wall	0.31
6	Prepare	0.46	Determine	0.56	Shape	0.40	Color	0.36	Erect	0.30
7	Research	0.44	Test	0.52	Screw	0.40	Lay	0.36	Winch	0.30
8	Confer	0.43	Program	0.51	Machine	0.38	Solder	0.33	Repair	0.30
9	Improve	0.43	Result	0.49	Install	0.38	Weld	0.32	Position	0.29
10	Program	0.42	Improve	0.48	Angle	0.37	Fabricate	0.32	Hole	0.28
11	Determine	0.42	Interpret	0.44	Weld	0.37	Mount	0.32	Power	0.28
12	Evaluate	0.41	Recommend	0.44	Hammer	0.36	Alter	0.31	Point	0.27
13	Project	0.41	Evaluate	0.44	Lay	0.36	Cut	0.31	Drive	0.27
14	Budget	0.41	Formulate	0.44	Clamp	0.35	Fit	0.31	Perform	0.26
15	Concern	0.41	Establish	0.43	Hole	0.35	Title	0.31	Water	0.26
16	Review	0.41	Engineer	0.42	Wire	0.35	Engrave	0.31	Hand	0.26
17	Report	0.41	Specialize	0.42	Produce	0.35	Coat	0.31	Attach	0.26
18	Recommend	0.40	Affect	0.42	Join	0.35	Plate	0.31	Couple	0.25
19	Advise	0.40	Factor	0.40	Replace	0.34	Diagnose	0.30	Maneuver	0.23
20	Approve	0.39	Confer	0.40	Power	0.34	Stage	0.30	Warn	0.23

Note: Top twenty verbs most correlated with each numerical score summarizing job requirements across the cross-section of occupations in 2000. Frequency with which each verb is used for each occupation is measured using VerbFreq_{vo} from equation (22). Numerical scores are non-routine interactive, non-routine analytic, routine cognitive, routine manual and non-routine manual from the 1991 DOTs. Both verb frequencies and numerical scores are converted into percentiles of their distribution across occupations. Correlation coefficients between each verb and each numerical score are reported next to each verb.

Table 2: Top Five Verbs for Each Thesaurus Section

Thesaurus Section	Verb 1	Verb 2	Verb 3	Verb 4	Verb 5
1 Class I. Words Expressing Abstract Relations					
1.1 Section I. Existence	Exist	Zero	Obtain	Posture	Continue
1.2 Section II. Relation	Correspond	Adapt	Reprint	Contrast	Harmonize
1.3 Section III. Quantity	Solder	Cleave	Blend	Clamp	Latch
1.4 Section IV. Order	Sample	Specialize	Include	Disperse	Cluster
1.5 Section V. Number	Recur	Halve	Schedule	Invoice	Compute
1.6 Section VI. Time	Date	Modernize	Synchronize	Dial	Late
1.7 Section VII. Change	Undergo	Transform	Vary	Change	Arrive
1.8 Section VIII. Causation	Generate	Sheathe	Energize	Fertilize	Heir
2 Class II. Words Relating to Space					
2.1 Section I. Space in General	Pouch	Reside	Locate	Camp	Vacate
2.2 Section II. Dimensions	Mesh	Widen	Bunk	Tape	Layer
2.3 Section III. Form	Curl	Spike	Envelope	Scoop	Gouge
2.4 Section IV. Motion	Bob	Shunt	Ramp	Dive	Export
3 Class III. Words Relating to Matter					
3.1 Section I. Matter in General	Weigh	Float	Swim	Balloon	Pound
3.2 Section II. Inorganic Matter	Grease	Irrigate	Soak	Liquefy	Lard
3.3 Section III. Organic Matter	Tint	Glare	Smell	Chime	Bleach
4 Class IV. Words Relating to the Intellectual Faculties					
4.1 Division I. Formation of Ideas					
4.1.1 Section I. Operations of Intellect in General	Occur	Discuss	Weigh	Loop	Digest
4.1.2 Section II. Precursory Conditions and Operations	Assay	Examine	Scrutinize	Trawl	Experiment
4.1.3 Section III. Materials for Reasoning	Ensure	Testify	Attest	Authenticate	Insure
4.1.4 Section IV. Reasoning Processes	Disprove	Guess	Defeat	Demonstrate	Mystify
4.1.5 Section V. Results Of Reasoning	Conform	Minimize	Adjudicate	Detect	Unlock
4.1.6 Section VI. Extension of Thought	Predict	Memorize	Forecast	Announce	Anticipate
4.1.7 Section VII. Creative Thought	Visualize	Guess	Originate	Fabricate	Devise
4.2 Division II. Communication of Ideas					
4.2.1 Section I. Nature of Ideas Communicated	Decipher	Annotate	Interpret	Fudge	Clarify
4.2.2 Section II. Modes of Communication	Disguise	Fake	Learn	Teach	Educate
4.2.3 Section III. Means of Communicating Ideas	Write	Describe	Relate	Narrate	Underlay
5 Class V. Words Relating to the Voluntary Powers					
5.1 Division I. Individual Volition					
5.1.1 Section I. Volition in General	Familiarize	Incline	Volunteer	Deter	Warn
5.1.2 Section II. Prospective Volition	Rot	Drug	Poison	Purify	Misuse
5.1.3 Section III. Voluntary Action	Manage	Consult	Fatigue	Transact	Confer
5.1.4 Section IV. Antagonism	Contest	Bombard	Assist	Avert	Obstruct
5.1.5 Section V. Results of Voluntary Action	Abort	Accomplish	Defeat	Drown	Blossom
5.2 Division II. Social Volition					
5.2.1 Section I. General Intersocial Volition	Restrain	Liberate	Ballot	Delegate	Curb
5.2.2 Section II. Special Intersocial Volition	Petition	Prohibit	Permit	Authorize	Invite
5.2.3 Section III. Conditional Intersocial Volition	Underwrite	Pawn	Endorse	Observe	Sponsor
5.2.4 Section IV. Possessive Relations	Afford	Finance	Liquidate	Grab	Clutch
6 Class VI. Emotion, Religion and Morality					
6.1 Section I. Affections in General	Awaken	Impress	Stimulate	Animate	Excite
6.2 Section II. Personal Affections	Enliven	Fear	Reassure	Beautify	Decorate
6.3 Section III. Sympathetic Affections	Snarl	Welcome	Kiss	Visit	Butcher
6.4 Section IV. Moral Affections	Switch	Thresh	Police	Tipple	Disapprove
6.5 Section V. Religious Affections	Anoint	Induct	Translate	Justify	Cure

Note: Verbs most concentrated in each thesaurus section (verbs with the top five values of ThesMean_{vk} from equation (23) for each thesaurus section, where verb 1 is the highest ranked). Verbs are first sorted by their number of occurrences in a thesaurus section divided by their total number of occurrences in the Dictionary of Occupational Titles (DOTs) for 1991. If two or more verbs have the same value of this fraction, they are next sorted by their number of occurrences in the DOTs.

Table 3: Tasks with Top and Bottom Twenty Increases in Task Input with Computer Use 1880-2000

Ranking	Verb	Coefficient	Standard Error	Ranking	Verb	Coefficient	Standard Error
1	Program	0.5319	0.0172	1830	Truck	-0.0373	0.0192
2	Direct	0.2197	0.0499	1829	Serve	-0.0214	0.0169
3	Test	0.1679	0.0280	1828	Clean	-0.0171	0.0169
4	Use	0.1614	0.0555	1827	Pump	-0.0166	0.0058
5	Engineer	0.1542	0.0094	1826	Cook	-0.0143	0.0075
6	Chart	0.1471	0.0047	1825	Ticket	-0.0135	0.0063
7	Plan	0.1450	0.0159	1824	Rest	-0.0124	0.0051
8	Determine	0.1436	0.0162	1823	Machine	-0.0108	0.1223
9	Prepare	0.1434	0.0170	1822	Deliver	-0.0100	0.0103
10	Analyze	0.1352	0.0070	1821	Grind	-0.0099	0.0115
11	Review	0.1242	0.0132	1820	Weld	-0.0099	0.0127
12	Develop	0.1171	0.0088	1819	Drive	-0.0094	0.0116
13	Work	0.1098	0.0361	1818	Set	-0.0094	0.0274
14	Operate	0.1050	0.0521	1817	Wheel	-0.0090	0.0076
15	Design	0.0979	0.0091	1816	Assemble	-0.0090	0.0234
16	Project	0.0957	0.0135	1815	Race	-0.0089	0.0046
17	Enter	0.0927	0.0057	1814	Nurse	-0.0088	0.0087
18	Report	0.0875	0.0133	1813	Collect	-0.0079	0.0100
19	Run	0.0864	0.0035	1812	Fuel	-0.0077	0.0088
20	Process	0.0826	0.0059	1811	Help	-0.0076	0.0102

Note: Table reports the tasks with the twenty highest and lowest increases in task input in industries intensive in the use of office and computing machinery. Estimated coefficients correspond to β_{tkm} from the regression (27) of industry inputs of task k on a time-invariant measure of industry use of office and computing machinery m interacted with dummies for year t. Each cell corresponds to a separate regression. Observations are industries and years. Coefficients are for 2000 and 1880 is the excluded year. Industry task inputs are the employment-weighted average of each occupation's task input as measured using the verbs from the occupational descriptions in the Dictionary of Occupational Titles (DOTs) for 1991 (VerbFreq_{vo} from equation (22)). Employment is measured using IPUMS population census data for each twenty-year interval from 1880-2000.

Table 4: Estimated Numerical Score Task Premia from Mincer Regression

Panel A: Numerical Score Task Measures				
	1940	1960	1980	2000
Non-routine	0.532***	0.523***	0.610***	0.700***
	(0.078)	(0.082)	(0.095)	(0.120)
Routine	0.096	0.276**	0.178**	0.061
	(0.103)	(0.110)	(0.089)	(0.110)
Manual	−0.207**	0.004	0.115	0.140
	(0.087)	(0.107)	(0.088)	(0.100)
Gender	yes	yes	yes	yes
Age	yes	yes	yes	yes
Education	yes	yes	yes	yes
Ethnicity	yes	yes	yes	yes
Observations	9,897,503	312,765	529,328	905,623
R-squared	0.356	0.377	0.271	0.231
Panel B: Thesaurus Task Measures				
	1940	1960	1980	2000
1. Abstract Relations	0.196***	0.219***	0.124	−0.017
	(0.072)	(0.056)	(0.085)	(0.100)
2. Space	−0.011	−0.095	−0.109	−0.126
	(0.081)	(0.080)	(0.104)	(0.110)
3. Matter	−0.226***	−0.309***	−0.329***	−0.396***
	(0.100)	(0.109)	(0.082)	(0.090)
4.1. Formation of Ideas	0.361***	0.383***	0.414***	0.550***
	(0.081)	(0.070)	(0.095)	(0.103)
4.2. Communication of Ideas	0.179*	0.026	−0.049	−0.073
	(0.098)	(0.101)	(0.097)	(0.124)
5.1. Individual Volition	−0.213**	−0.169*	−0.032	−0.059
	(0.093)	(0.097)	(0.086)	(0.101)
5.2. Social Volition	−0.008	−0.139**	−0.161**	−0.099
	(0.055)	(0.059)	(0.071)	(0.910)
6. Emotion, Religion and Morality	0.037	−0.000	0.082	0.076
	(0.079)	(0.084)	(0.077)	(0.075)
Gender	yes	yes	yes	yes
Age	yes	yes	yes	yes
Education	yes	yes	yes	yes
Ethnicity	yes	yes	yes	yes
Observations	8,616,583	257,101	474,758	825,225
R-squared	0.375	0.409	0.282	0.246

Note: Table reports the estimated task premia (ζ_t) on percentile task scores ($\mathbb{T}_{o(i)}$) from the Mincer regression (26) for workers i in occupations o in year t, including controls for worker education, age, gender and ethnicity. Each column in each panel corresponds to a separate regression for a separate year (column) and specification (panel). Panel A reports results using percentiles of the numerical score measures of tasks (non-routine, routine and manual) across occupations. Panel B reports results using percentiles of the thesaurus task content measure (ThesFreq_{ko} from equation (24)) across occupations. Standard errors in parentheses are heteroscedasticity robust and clustered on occupation. *** denotes significance at the 1 percent level; ** indicates significance at the 5 percent level; and * corresponds to significance at the 10 percent level.

Table 5: Top and Bottom Ten Verbs Most Concentrated in Metro Areas (1991 DOT)

Panel A: Verbs Most Strongly Correlated with Metro Area Employment Shares							
Rank	1880	1900	1920	1940	1960	1980	2000
1	Thread	Thread	File	File	Document	Identify	Develop
2	Stretch	Stretch	Distribute	Compile	Schedule	Document	Determine
3	Sew	Sew	Telephone	Bill	File	Advise	Analyze
4	Stitch	Stitch	Record	Letter	Record	Report	Factor
5	Ravel	Pulse	Notice	Notice	Compile	Concern	Review
6	Flap	Hand	Bill	Record	Identify	Schedule	Confer
7	Shoulder	Shoulder	Document	Identify	Notice	Analyze	Advise
8	Row	Visit	Number	Distribute	Distribute	Develop	Report
9	Sack	Sack	Envelope	Send	Send	Determine	Concern
10	Snap	Telephone	Identify	Learn	Notify	Notify	Plan
Panel B: Verbs Least Strongly Correlated with Metro Area Employment Shares							
Rank	1880	1900	1920	1940	1960	1980	2000
N-9	Channel	Tread	Teach	Enlist	Destroy	Water	Restrain
N-8	Sound	Tunnel	Matter	Pinch	Escape	Pour	Cut
N-7	See	Settle	Tunnel	Labor	Tread	Power	Power
N-6	Rule	Matter	Assign	State	Traverse	Erect	Massage
N-5	Matter	Pinch	Administer	Tread	Loosen	Pump	Remove
N-4	Tunnel	Assign	Consolidate	Tunnel	Range	Cut	Feed
N-3	Tread	Sound	Sound	Malt	Accord	Lever	Clean
N-2	Pinch	See	Discuss	Rock	Activate	Feed	Pump
N-1	Drill	Rule	Rule	Establish	Turn	Turn	Move
N	Sole	Sole	State	Move	Move	Move	Turn

Note: Panel A reports the ten verbs with the highest correlations with metro area employment shares, as measured by the ten verbs v with the largest estimated coefficients in year t (α_{vt}) in equation (29). Panel B reports the ten verbs with the lowest correlations with metro area employment shares, as measured by the ten verbs v with the smallest estimated coefficients in year t (α_{vt}) in equation (29). Observations are sectors and occupations in each year. Metro area employment shares measured as the share of employment within a sector-occupation that is located in a metro area in each year. The reported coefficients are standardized by variable standard deviations ("beta coefficients"). Verb frequencies for each occupation are measured from the occupational descriptions in the Dictionary of Occupational Titles (DOTs) for 1991 (VerbFreq_{vo} from equation (22)).

Table 6: Metro Area Employment Shares and Thesaurus Sections (1991 DOT)

	Class	Section	(1) Rank Section 1880	(2) Rank Section 2000	(3) Rank 1880 - Rank 2000
1.1	Abstract relations	Existence	35	8	27
1.2	Abstract relations	Relation	12	14	-2
1.3	Abstract relations	Quantity	3	36	-33
1.4	Abstract relations	Order	7	9	-2
1.5	Abstract relations	Number	20	10	10
1.6	Abstract relations	Time	15	33	-18
1.7	Abstract relations	Change	34	32	2
1.8	Abstract relations	Causation	38	27	11
2.1	Space	Space in General	2	18	-16
2.2	Space	Dimensions	33	20	13
2.3	Space	Form	16	39	-23
2.4	Space	Motion	27	22	5
3.1	Matter	Matter in General	19	34	-15
3.2	Matter	Inorganic Matter	4	38	-34
3.3	Matter	Organic Matter	5	37	-32
4.1.1	Intellect	Operations of the Intellect in General	1	15	-14
4.1.2	Intellect	Precursory Conditions and Operations	28	7	21
4.1.3	Intellect	Materials for Reasoning	32	2	30
4.1.4	Intellect	Reasoning Processes	39	6	33
4.1.5	Intellect	Results of Reasoning	17	21	-4
4.1.6	Intellect	Extension of Thought	13	29	-16
4.1.7	Intellect	Creative Thought	22	26	-4
4.2.1	Intellect	Nature of Ideas Communicated	30	11	19
4.2.2	Intellect	Modes of Communication	36	25	11
4.2.3.	Intellect	Means of Communicating Ideas	10	3	7
5.1.1	Volition	Volition in General	29	4	25
5.1.2	Volition	Prospective Volition	25	30	-5
5.1.3	Volition	Voluntary Action	37	1	36
5.1.4	Volition	Antagonism	6	23	-17
5.1.5	Volition	Results of Voluntary Action	26	16	10
5.2.1	Volition	General Intersocial Volition	21	17	4
5.2.2	Volition	Special Intersocial Volition	23	13	10
5.2.3	Volition	Conditional Intersocial Volition	9	28	-19
5.2.4	Volition	Posessive Relations	31	5	26
6.1	Emotion	Affections in General	8	35	-27
6.2	Emotion	Personal Affections	14	24	-10
6.3	Emotion	Sympathetic Affections	24	19	5
6.4	Emotion	Moral Affections	18	12	6
6.5	Emotion	Religious Affections	11	31	-20

Note: Columns (1) and (2) report the ranking of thesaurus sections by correlations with metro area employment shares in 1880 and 2000 respectively, as measured by the ranking of thesaurus sections k in year t in terms of their estimated coefficients (β_{kt}) in equation (30). The highest value is assigned a rank of one. Column (3) reports the change in ranking (1880 minus 2000, so that positive numbers correspond to thesaurus sections that have become more concentrated in metro areas). Observations are sectors and occupations in each year. Metro area employment shares measured as the share of employment within a sector-occupation that is located in a metro area in each year. The reported coefficients are standardized by variable standard deviations ("beta coefficients"). Thesaurus task content for each occupation is measured using the verbs from occupational descriptions in the Dictionary of Occupational Titles (DOTs) for 1991 and Roget's thesaurus (ThesFreq_{ko} from equation (24)).