

Task Specialization in U.S. Cities from 1880-2000*

Guy Michaels

London School of Economics and CEP[†]

Ferdinand Rauch

University of Oxford and CEP[‡]

Stephen J. Redding

Princeton University and CEP[§]

January 10, 2013

Abstract

We develop a new methodology for quantifying the tasks undertaken within occupations using 3,000 verbs from around 12,000 occupational descriptions in the Dictionary of Occupational Titles (DOTs). Using micro-data from the United States from 1880-2000, we find an increase in the employment share of interactive occupations within sectors over time that is larger in metro areas than non-metro areas. We provide evidence that this increase in the interactivity of employment is related to the dissemination of improvements in transport and communication technologies. Our findings highlight a change in the nature of agglomeration over time towards an increased emphasis on human interaction.

KEYWORDS: Economic Development, Human Interaction, Urbanization.

JEL CLASSIFICATION: N92, O18, R12

*We are grateful to the Centre for Economic Performance and Princeton University for research support. We are also grateful to Gilles Duranton and Matt Turner for sharing data. We thank William D. Caughlin and George Kupczak for their help with data. We thank Don Davis, Gilles Duranton, Gene Grossman, Gordon Hanson, Marco Manacorda, Alan Manning, Nathan Nunn, Vernon Henderson, Rick Hornbeck, Gianmarco Ottaviano, Steve Pischke, Esteban Rossi-Hansberg, Yona Rubinstein, Will Strange, Nancy Qian, Tony Venables and seminar and conference participants at Copenhagen, the London School of Economics, NARSC, Oxford, Sorbonne and Sussex 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: ferdi-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

Agglomeration forces are widely understood to play a central role in sustaining the dense concentrations of population observed in urban areas. Much less is known about the detailed tasks undertaken in urban areas and how these have changed over time. Yet understanding the task content of employment in urban and rural areas is central to evaluating alternative theories of agglomeration and assessing the likely impact of improvements in transport and communication technologies on spatial concentrations of economic activity. In this paper, we provide new evidence on the detailed tasks undertaken by workers in urban and rural areas over a long historical time period in the United States (U.S.). We develop a new methodology for measuring the tasks undertaken within occupations using the verbs from occupational descriptions in the Dictionary of Occupational Titles (DOTs). We implement this methodology using micro data on employment in disaggregated occupations and sectors in metro and non-metro areas from 1880-2000.

To measure the tasks undertaken by workers within each occupation, we use 3,000 verbs from around 12,000 occupational descriptions in both historical and contemporary editions of the Dictionary of Occupational Titles (DOTs). Using these verbs, we find a systematic change in the composition of employment across tasks in urban versus rural areas over time. We quantify this change in the task composition of employment using the meaning of verbs from the standard classification of English language use in Roget's Thesaurus. In both metro and non-metro areas, we find a systematic reallocation of employment over time towards interactive occupations, which involve tasks described by verbs that appear in thesaurus categories concerned with thought, communication and inter-social activity. At the beginning of our sample period, non-metro areas actually have higher shares of employment in interactive occupations than metro areas. Over time, employment growth in interactive occupations is much higher in metro areas, so that by the end of our sample period the initial pattern of specialization is reversed, and metro areas are more interactive than non-metro areas. This increasing interactiveness of employment at higher population densities is observed not only between metro and non-metro areas but also across metro areas of different population densities. While in 1880 there is little relationship between specialization in interactive occupations and population density, by 2000 this relationship is positive, strong and statistically significant. Taken together, these results suggest that human interaction has become increasingly important in agglomerations of economic activity over time.

We provide evidence that this change in the task composition of employment reflects a secular and systematic process of structural change. We find a similar pattern using variation across sectors, within sectors and within occupations and sectors. The increased interactiveness of employment in metro areas is not driven by any one occupation or sector, and is not explained by an expansion in the geographical boundaries of metro areas over time, since we find similar results when we restrict attention to central cities. We examine and provide evidence against a number of potential explanations, such as the dispersion of

manufacturing from urban areas, structural transformation from manufacturing to services, and an increase in female labor force participation with the attendant need for “power couples” to collocate. For the more limited time period for which we have wage data, we find that the increase in the relative employment share of interactive occupations in metro areas goes hand in hand with an increase in their relative wage bill share. This pattern of results is consistent with the observed change in the task composition of employment being driven by a change in relative demand that increases both the relative employment and relative wagebills of interactive occupations rather than by a change in relative supply that increases the relative employment but decreases the relative wagebills of interactive occupations.

Within our sample period, we find that the increased interactiveness of employment in metro areas is particularly strong from 1880-1930, which is a period of rapid change in communication and transport technology. Following the award of Alexander Graham Bell’s patent for the telephone in 1886, the U.S. telephone network grew rapidly in the opening decades of the twentieth century.¹ After the award of Karl Benz’s patent for the internal combustion engine in 1879 and after the passage of the Federal Aid Road Act of 1916 and the Federal Highway Act of 1921, the U.S. road network and automobile use expanded rapidly over the same period.² To provide evidence on the role played by these improvements in communication and transport technologies in influencing the task composition of employment, we combine county data on employment by occupation and sector for 1880 and 1930 with newly-collected county data on telephone use and the road network in the 1930s. To address the concern that both telephone use and the road network could be endogenously influenced by changes in interactiveness as a result of omitted third factors, we develop instruments for the geographical dissemination of these improvements in technology. For the telephone, we use its network properties to construct an instrument based on proximity to nodes on the American Telegraph and Telephone’s (AT&T) company’s long distance trunk network, whose construction was influenced by the strategic objectives of connecting the nation as a whole. For roads, we use the 1922 “Pershing Map” of highway routes of military importance for coastal and border defense. Consistent with these improvements in communication and transport technologies influencing the task content of employment, we find a strong, positive and statistically significant relationship between increases in the interactiveness of employment and telephone use and road construction predicted by our instruments.

To interpret our empirical results, we develop a model that captures specialization across the many sectors, occupations and locations observed in our data. We use the model to rationalize a secular reallocation of employment towards interactive occupations across all locations and a greater reallocation of employment towards interactive occupations in more densely-populated locations. The distribution of economic activity across locations is determined by productivity differences (including agglomeration forces which

¹See, for example, John (2010). The electric telegraph was patented much earlier in 1837 by Samuel Morse and the U.S. telegraph network was largely complete by 1880 (see Standage 2007).

²See, for example, Swift (2011).

concentrate population in productive locations) and an inelastic supply of land (which favors population dispersion). Each location can produce in a number of sectors (e.g. Manufacturing, Services, which correspond to the two-digit sectors in our data) and each sector includes a number of more disaggregated goods (e.g. Motor Vehicles, Drugs and Medicines, which correspond to the three-digit sectors in our data). Production of each good involves a number of stages of production supplied by occupations (e.g. Managers, Operatives, which correspond to the two-digit occupations in our data) and workers within in each occupation perform a number of tasks (e.g. advising, typing, stamping, stretching, which correspond to the verbs in the descriptions for the disaggregated occupations in our data). Final goods can be traded between locations subject to goods trade costs that can differ across sectors. Tasks also can be traded between locations subject to task trade costs that differ across occupations.

As in the macroeconomics literature on structural transformation, the model can account for a secular reallocation of employment across sectors and occupation in all locations. When demand is inelastic for sectors and occupations, more rapid productivity growth in some sectors and occupations leads to a more than proportionate fall in price, which reallocates employment towards other sectors and occupations. As in the international trade literature on offshoring, the model can account for a greater reallocation of employment towards interactive locations in more densely-populated locations through improvements in communications and transport technologies. Reductions in final goods trade costs induce specialization across sectors according to standard theories of comparative advantage. Reductions in task trade costs induce an analogous process of specialization across occupations within sectors. When task trade costs are prohibitively high, all tasks are performed in the location in which the final good is produced. As task trade costs fall, it becomes feasible to unbundle production across locations and trade tasks between these locations. To the extent that densely populated locations have a comparative advantage in interactive tasks (e.g. because agglomeration forces are stronger for interactive tasks), reductions in task trade costs induce densely-populated locations to specialize in more-interactive occupations, while more sparsely-populated locations specialize in less-interactive occupations.

Our paper is related to a number of literatures. We build on the wider literature on agglomeration economies, as surveyed in Duranton and Puga (2004) and Rosenthal and Strange (2004). One strand of this literature emphasizes differences in the composition of economic activity between urban and rural areas. Studies emphasizing the role of human capital and skills in promoting agglomeration include Glaeser and Saiz (2003), Glaeser and Resseger (2009), Bacolod, Blum and Strange (2009a), Glaeser, Ponzetto and Toblo (2011) and Moretti (2004). Particular types of skills are highlighted in Bacalod, Blum and Strange (2009b), which introduces the concept of soft skills that enable agents to interact in cities and industry clusters. More generally, the role of idea generation and exchange is emphasized in Davis and Dingel (2012), which develops a system of cities model in which costly idea exchange is the agglomeration force. Our contribution relative to this literature is to provide detailed microeconomic evidence on the tasks

undertaken in urban and rural areas over time and the role of improvements in communication and transport technologies in influencing these tasks.

Another line of research has distinguished different dimensions along which cities specialize. Duranton and Puga (2005) provides theory and evidence that in recent decades cities have shifted from specializing by sector – with integrated headquarters and plants – to specializing mainly by function - with headquarters and business services clustered in larger cities and plants clustered in smaller areas.³ Rossi-Hansberg, Sarte and Owens (2009) develops a model in which firms choose locations of their headquarters and production facilities, and argues that the increased separation of these locations accounts for observed changes in patterns of residential and business activity. Ota and Fujita (1993) models the distinction between the front-unit (e.g. business office) and back-unit (e.g. plant or back-office) of firms and explores its implications for city structure. Helsley and Strange (2007) explicitly analyzes the vertical integration decision of the firm in conjunction with its location decision. Fujita and Tabuchi (1997) provides evidence that the increased separation of headquarters and production has contributed to observed changes in the distribution of economic activity across Japanese regions. Our contribution is again to provide microeconomic evidence on task specialization. We show that the changes in interactiveness within sectors observed during our sample period are not driven by a single sector or occupation, such as Manufacturing or Managers, but are rather pervasive across sectors and occupations.

Existing research has explored a variety of other dimensions of specialization in urban areas. Following Jacobs (1969), a long line of research distinguishes between localization externalities within industries and urbanization externalities across industries, as for example in Henderson (2003). Duranton and Jayet (2011) examine Adam Smith’s idea that the division of labor is determined by the extent of the market and find that rare occupations are over-represented in large cities. Duranton and Puga (2001) contrast innovation and production and introduce a distinction between diversified and specialized cities.⁴ Finally, Combes et al. (2012) provide evidence on the extent to which the productivity advantages of large cities are accounted for by a rightwards shift in the firm productivity distribution (agglomeration) or a greater left-truncation of the firm productivity distribution (selection). None of these studies provide evidence on task specialization or interactiveness.

Our analysis is also related to the labor economics literature on the task content of employment, including Autor, Levy and Murnane (2003), Autor and Handel (2009), and Gray (2010). Whereas this prior research has focused on the numerical scores from the DOTs such as “Direction, Control and Planning (DCP),” we make use of the detailed occupational descriptions. As a result, we are able to characterize changes in task specialization in rich detail using over 3,000 verbs and 12,000 occupational descriptions.

³See Henderson (1974) for the classic analysis of industry specialization and the size distribution of cities.

⁴Lin (2011) develops a measure of new work based on revisions to occupational classifications and finds that new work is more likely to be observed in locations initially dense in college graduates and industry variety.

This approach enables us to develop our new measure of interactiveness, which we compare to existing occupational characteristics from the DOTs. We provide evidence on the evolving task content of employment over a much longer time period than usually considered in the labor economics literature. Our focus is on differences in task specialization between urban and rural areas over time, which has received little attention in the existing labor literature, and yet is central to debates about the nature of agglomeration.

An advantage of our empirical setting using U.S. micro data over a long historical time period is that we can provide direct evidence on the role of improvements in transport technology (roads) and communication technology (telephones) in explaining the observed increase in interactiveness. Our work is therefore related to the literature on roads and urban growth (e.g. Baum-Snow 2007 and Duranton and Turner 2012) and the literature on innovations in communication technology and urban growth (e.g. Pool 1977, Fischer 1992, Gaspar and Glaeser 1998, Leamer and Storper 2001). Since these improvements in transport and communication technologies facilitate the remote sourcing of tasks, our work is also related to the international trade literature on offshoring, including Becker, Ekholm and Muendler (2009), Blinder (2009), Blinder and Krueger (2012), Grossman and Rossi-Hansberg (2008), and Ottaviano, Peri and Wright (2010). Our micro data enable us to explore the effects of these innovations on task specialization at a fine spatial scale, to show that remote sourcing can occur within as well as across countries, and to develop instruments for the dissemination of these improvements in transport and communication technologies.

The remainder of the paper is structured as follows. Section 2 introduces the theoretical framework that we use to interpret our empirical results. Section 3 discusses the data. Section 4 provides some motivating evidence on specialization across occupations and sectors in urban and rural areas that motivates our analysis of task specialization. Section 5 introduces our methodology for measuring the tasks undertaken within each occupation using the verbs from occupational descriptions. Section 6 develops quantitative measures of task specialization and presents our baseline evidence on the increase in the relative interactiveness of urban areas over time. Section 7 reports a number of robustness tests and compares our measure of interactiveness with other measures of occupational characteristics. Section 8 provides evidence on the determinants of the increase in interactiveness over time. Section 9 concludes.

2 Theoretical Model

In this section, we outline a theoretical model that we use to understand the distribution of employment across occupations, sectors and locations in our data.⁵ As in standard models of agglomeration, the distribution of population across locations is driven by a tension between agglomeration forces (productivity differences that depend on production externalities) and congestion forces (an inelastic supply of land). In

⁵A more detailed exposition of the model including the technical derivations of relationships is contained in a web-based technical appendix.

contrast to much of the theoretical literature on agglomeration, our framework allows for a large number of locations and incorporates multiple layers of specialization across both sectors and occupations within sectors. Nonetheless, the analysis remains tractable because of the stochastic formulation of productivity differences across sectors, occupations and locations.

2.1 Preferences and Endowments

The economy consists of many locations indexed by $n \in N$. Each location n is endowed with an exogenous supply of land \bar{H}_n . The economy as a whole is endowed with a measure of workers \bar{L} , who are perfectly mobile across locations. Workers' preferences are defined over a goods consumption index (C_n) and residential land use (H_n) and are assumed to take the Cobb-Douglas form:⁶

$$U_n = C_n^\alpha H_n^{1-\alpha}, \quad 0 < \alpha < 1. \quad (1)$$

The goods consumption index (C_n) is assumed to be a constant elasticity of substitution (CES) function of consumption indices for a number of sectors (e.g. Manufacturing, Services) indexed by $s \in S$:

$$C_n = \left[\sum_{s \in S} C_{ns}^{\frac{\beta-1}{\beta}} \right]^{\frac{\beta}{\beta-1}}, \quad (2)$$

where β is the elasticity of substitution between sectors. Sectors can be either substitutes ($\beta > 1$) or complements in goods consumption ($0 < \beta < 1$), where the standard assumption in the literature on structural transformation in macroeconomics is complements (e.g. Ngai and Pissarides 2007, Yi and Zhang 2010).

The consumption index for each sector is in turn a CES function of consumption of a continuum of goods (e.g. Motor Vehicles, Drugs and Medicines) indexed by $j \in [0, 1]$:

$$C_{ns} = \left[\int_0^1 c_{ns}(j)^{\frac{\sigma_s-1}{\sigma_s}} dj \right]^{\frac{\sigma_s}{\sigma_s-1}}, \quad (3)$$

where the elasticity of substitution between goods σ_s varies across sectors. While in the data we observe a finite number of goods within sectors, we adopt the theoretical assumption of a continuum of goods for reasons of tractability, because it enables us to make use of law of large numbers results in determining specialization at the sectoral level. Goods can be either substitutes ($\sigma_s > 1$) or complements ($0 < \sigma_s < 1$) and we can allow any ranking of the elasticities of substitution between goods and sectors, although the conventional assumption in such a nested CES structure is a higher elasticity of substitution at the more disaggregated level ($\sigma_s > \beta$).

⁶For empirical evidence using U.S. data in support of the constant expenditure share implied by the Cobb-Douglas functional form, see Davis and Ortalo-Magne (2011).

Expenditure on residential land in each location is assumed to be redistributed lump-sum to residents of that location, as in Helpman (1998). Therefore total income in each location equals payments to labor used in production plus expenditure on residential land:

$$v_n L_n = w_n L_n + (1 - \alpha) v_n L_n = \frac{w_n L_n}{\alpha}, \quad (4)$$

where w_n is the wage; L_n is the population of location n ; and equilibrium land rents in each location (r_n) are determined by land market clearing.

2.2 Production

Goods are homogeneous in the sense that one unit of a given good is the same as any other unit of that good. Production occurs under conditions of perfect competition and constant returns to scale. The cost to a consumer in location n of purchasing one unit of good j within sector s from location i is therefore:

$$p_{nis}(j) = \frac{d_{nis} G_{is}(j)}{z_{is}(j)}, \quad (5)$$

where d_{nis} are iceberg goods trade costs, such that $d_{nis} > 1$ must be shipped from location i to location n within sector s in order for one unit to arrive; $z_{is}(j)$ is productivity for good j within sector s in location i ; and $G_{is}(j)$ is the unit cost of the composite factor of production used for good j within sector s in location i , as determined below.

Final goods productivity is stochastic and modeled as in Eaton and Kortum (2002) and Costinot, Donaldson and Komunjer (2012). Final goods productivity for each good, sector and location is assumed to be drawn independently from a Fréchet distribution:⁷

$$F_{is}(z) = e^{-T_{is} L_{is}^{\eta_s} z^{\theta_s}}, \quad (6)$$

where the shape parameter ($\theta_s > 1$) controls the dispersion of productivity across goods within each sector, which determines comparative advantage across goods. In contrast, the scale parameter ($T_{is} L_{is}^{\eta_s}$, where $\eta_s > 0$) determines average productivity within each sector for each location, which determines comparative advantage across sectors. We allow average productivity in a sector and location to be increasing in employment in that sector and location to capture agglomeration forces in the form of external economies of scale in final goods production (e.g. Ethier 1982).

As in the Ricardian model of trade, our framework features comparative advantage across goods and sectors, which explains the observed variation in the shares of sectors and goods in employment across locations in the data. But we also observe variation in the shares of occupations and tasks in employment across locations in the data, which is not captured in the standard Ricardian framework. To account for

⁷To simplify the exposition, we use i to denote locations of production and n to denote locations of consumption, except where otherwise indicated.

this additional layer of specialization, we assume that each good is produced using a number of stages of production, where each stage of production within a sector is supplied by a separate occupation indexed by $o \in O_s$ (e.g. Managers, Operatives). Output of good j within sector s in location i ($y_{is}(j)$) is a CES function of the inputs of each occupation ($X_{iso}(j)$):

$$y_{is}(j) = \left[\sum_{o \in O_s} X_{iso}(j)^{\frac{\mu_s - 1}{\mu_s}} \right]^{\frac{\mu_s}{\mu_s - 1}}, \quad (7)$$

where μ_s is the elasticity of substitution between occupations and again we can allow occupations to be either substitutes ($\mu_s > 1$) or complements ($0 < \mu_s < 1$). We allow sectors to differ in terms of the set of occupations O_s , and firms within each sector adjust the proportions with which workers in different occupations are employed depending their cost.

Workers within each occupation perform a continuum of tasks $t \in [0, 1]$ as in Grossman and Rossi-Hansberg (2008) (e.g. Advising, Typing, Stretching, Stamping). The input for occupation o and good j within sector s and location i ($X_{iso}(j)$) is a CES function of the inputs for these tasks ($x_{iso}(j, t)$):

$$X_{iso}(j) = \left[\int_0^1 x_{iso}(j, t)^{\frac{\nu_{so} - 1}{\nu_{so}}} dt \right]^{\frac{\nu_{so}}{\nu_{so} - 1}} \quad (8)$$

where the elasticity of substitution between tasks ν_{so} varies across sectors and occupations. While in the data we observe a finite number of tasks within occupations, we again adopt the theoretical assumption of a continuum of tasks for reasons of tractability, because it enables us to make use of law of large numbers results in determining specialization at the occupational level.⁸ We allow tasks within occupations to be either substitutes ($\nu_{so} > 1$) or complements ($0 < \nu_{so} < 1$), and we can consider any ranking of the elasticities of substitution between tasks and occupations, although the conventional assumption in such a nested CES structure is again a higher elasticity of substitution at the more disaggregated level ($\nu_{so} > \mu_s$).⁹

Tasks are performed by labor using a constant returns to scale technology and can be traded between locations. For example, product design can be undertaken in one location, while production and assembly occur in another location. The cost to a firm in location n of sourcing a task t from location i within occupation o and sector s is:

$$g_{niso}(j, t) = \frac{\tau_{niso} w_i}{a_{iso}(j, t)}, \quad (9)$$

where w_i is the wage; τ_{niso} are iceberg task trade costs, such that $\tau_{niso} > 1$ units of the task must be performed in location i in order for one unit to be completed in location n for occupation o and sector s ; $a_{iso}(j, t)$ is productivity for task t and good j within occupation o and sector s in location i .

⁸To reduce the notational burden, we assume the same $[0, 1]$ interval of tasks for all occupations, but it is straightforward to allow this interval to vary across occupations.

⁹While we interpret production as being undertaken by workers in occupations that perform many tasks, an equivalent interpretation is that each occupation corresponds to a stage of production and each task corresponds to an intermediate input within that stage of production.

Input productivity for each task, occupation, sector and location is also stochastic and is assumed to be drawn independently from a Fréchet distribution:

$$\mathcal{F}_{iso} = e^{-U_{iso}L_{iso}^{\chi_{so}}a^{-\epsilon_{so}}}, \quad (10)$$

where the shape parameter ($\epsilon_{so} > 1$) controls the dispersion of productivity across tasks within occupations, which determines comparative advantage across tasks. In contrast, the scale parameter ($U_{iso}L_{iso}^{\chi_{so}} > 0$, where $\chi_{so} > 0$) controls average productivity within each occupation, which determines comparative advantage across occupations. We allow average productivity in an occupation, sector and location to be increasing in employment in that occupation, sector and location ($\chi_{so} > 0$) to capture external economies of scale in task production (e.g. Grossman and Rossi-Hansberg 2012).

2.3 Trade in Tasks and Input Costs

Firms within a given location n source each task t within an occupation o , good j and sector s from the lowest cost source of supply for that task:

$$g_{nso}(j, t) = \min \{g_{niso}(j, t); i \in N\}.$$

Given finite task trade costs, locations supply tasks for which they have high productivity draws themselves, and source other tasks for which they have low productivity draws from other locations. Under our assumption of a Fréchet distribution of input productivity, the share of firm costs in location n accounted for by tasks sourced from location i within occupation o and sector s (λ_{niso}) is equal to the fraction of tasks sourced from that location:¹⁰

$$\lambda_{niso} = \frac{U_{iso}L_{iso}^{\chi_{so}}(\tau_{niso}w_i)^{-\epsilon_{so}}}{\sum_{k \in N} U_{kso}L_{kso}^{\chi_{so}}(\tau_{nkso}w_k)^{-\epsilon_{so}}}, \quad (11)$$

and the unit cost for occupation o and sector s in location n can be written as:

$$G_{nso} = \gamma_{so} \left(\frac{U_{nso}L_{nso}^{\chi_{so}}}{\lambda_{nnso}} \right)^{-\frac{1}{\epsilon_{so}}} w_n. \quad (12)$$

where $\gamma_{so} = \left[\Gamma \left(\frac{\epsilon_{so} + 1 - \nu_{so}}{\epsilon_{so}} \right) \right]^{\frac{1}{1 - \nu_{so}}}$ and $\Gamma(\cdot)$ is the gamma function.

Intuitively, costs for occupation o in sector s and location n are low when the average input productivity for performing tasks with that occupation, sector and location ($U_{nso}L_{nso}^{\chi_{so}}$) is high, when the wage (w_n) is low, and when that occupation, sector and location spends a low share of its costs on itself (λ_{nnso}). The

¹⁰Since the Fréchet distribution is unbounded from above, each location draws an arbitrarily high input productivity for a positive measure of tasks. To allow for the possibility that a location may not have positive employment in an occupation o and sector s , we take $\lim U_{iso} \rightarrow 0$, in which case the location's employment in that occupation and sector converges to zero. Similarly, to allow for the possibility that an occupation o may not be traded, we take $\lim d_{nis} \rightarrow \infty$, in which case trade in that occupation converges to zero.

presence of the own trade share reflects the fact that remote sourcing tasks from other locations lowers unit costs, and hence acts like an increase in productivity that reduces unit costs.

Using unit costs for each occupation (11) and the final goods production technology (7), the share of occupation o in sector costs can be written as:

$$e_{nso} = \frac{\gamma_{so}^{1-\mu_s} \left(\frac{U_{nso} L_{nso}^{\chi_{so}}}{\lambda_{nnsso}} \right)^{-\frac{1-\mu_s}{\epsilon_{so}}}}{\sum_{m \in O_s} \gamma_{sm}^{1-\mu_s} \left(\frac{U_{nsm} L_{nsm}^{\chi_{sm}}}{\lambda_{nnsms}} \right)^{-\frac{1-\mu_s}{\epsilon_{sm}}}}. \quad (13)$$

Intuitively, high-unit-cost occupations (low $U_{nso} L_{nso}^{\chi_{so}} / \lambda_{nnsso}$) account for high shares of firm costs and employment if occupations are complements ($0 < \mu_s < 1$) and low shares of firm costs and employment if occupations are substitutes ($\mu_s > 1$).

2.4 Trade in Final Goods and Price Indices

Consumers within a given location n source each final good j within a sector s from the lowest cost source of supply for that final good:

$$p_{ns}(j) = \min \{p_{nis}(j); i \in N\}.$$

For finite final goods trade costs, locations supply final goods for which they have low unit costs themselves, and source other final goods for which they have high units costs from other locations. These unit costs for final goods depend on input productivities and trade in tasks, as characterized in the previous subsection, as well on final goods productivities. Under our assumption of a Fréchet distribution of final goods productivity, the share of location n 's expenditure on final goods produced in location i within sector s (π_{nis}) is equal to the fraction of final goods sourced from that location:¹¹

$$\pi_{nis} = \frac{T_{is} L_{is}^{\eta_s} (d_{nis} \Phi_{is} w_i)^{-\theta_s}}{\sum_{k \in N} T_{ks} L_{ks}^{\eta_s} (d_{nks} \Phi_{ks} w_k)^{-\theta_s}}, \quad (14)$$

and the price index for sector s in location n can be written as:

$$P_{ns} = \kappa_s \left(\frac{T_{ns} L_{ns}^{\eta_s}}{\pi_{nns}} \right)^{-\frac{1}{\theta}} \Phi_{ns} w_n. \quad (15)$$

where Φ_{is} is a summary statistic for occupation unit costs in sector s in location i :

$$\Phi_{is} = \left[\sum_{o \in O_s} \gamma_{so}^{1-\mu_s} \left(\frac{U_{iso} L_{iso}^{\chi_{so}}}{\lambda_{iiso}} \right)^{-\frac{1-\mu_s}{\epsilon_{so}}} \right]^{\frac{1}{1-\mu_s}}. \quad (16)$$

¹¹Since the Fréchet distribution is unbounded from above, each location draws an arbitrarily high final goods productivity for a positive measure of final goods. To allow for the possibility that a location may not have positive employment in a sector s , we take $\lim T_{is} \rightarrow 0$, in which case the location's employment in that sector converges to zero. Similarly, to allow for the possibility that a sector s may not be traded, we take $\lim d_{nis} \rightarrow \infty$, in which case trade in that sector converges to zero.

and $\kappa_s = \left[\Gamma \left(\frac{\theta_s + 1 - \sigma_s}{\theta_s} \right) \right]^{\frac{1}{1 - \sigma_s}}$ where $\Gamma(\cdot)$ is the Gamma function.

Intuitively, the price index for sector s in location n is low when average productivity within that sector and location ($T_{ns}L_{ns}^{\eta_s}$) is high, when the own trade share for final goods within that sector and location (π_{nns}) is low, and when unit costs for that sector and location ($\Phi_{ns}w_n$) are low. Unit costs in turn can be low because of a low wage (w_n), high average input productivities for occupations within that sector and location ($U_{nso}L_{nso}^{\chi_{so}}$), and low own trade shares for occupations within that sector and location (λ_{nns}).

Using the sectoral price index (15), the share of sector s in aggregate goods expenditure is:

$$E_{ns} = \frac{\kappa_s^{1-\beta} \left(\frac{T_{ns}L_{ns}^{\eta_s}}{\pi_{nns}} \right)^{-\frac{1-\beta}{\theta_s}} \Phi_{ns}^{1-\beta}}{\sum_{r \in S} \kappa_r^{1-\beta} \left(\frac{T_{nr}L_{nr}^{\eta_r}}{\pi_{nnr}} \right)^{-\frac{1-\beta}{\theta_r}} \Phi_{nr}^{1-\beta}}. \quad (17)$$

Intuitively, high-price sectors (low $T_{ns}L_{ns}^{\eta_s}$, high π_{nns} and high Φ_{ns}) account for high shares of aggregate goods expenditure and employment if sectors are complements ($0 < \beta < 1$) and low shares of aggregate goods expenditure and employment if sectors are substitutes ($\beta > 1$).

2.5 Population Mobility

Population mobility implies that workers must receive the same indirect utility in all populated locations:

$$V_n = \frac{v_n}{P_n^\alpha r_n^{1-\alpha}} = \bar{V}, \quad (18)$$

where labor market clearing requires:

$$\sum_{n \in N} L_n = \bar{L}. \quad (19)$$

In the web appendix, we determine the general equilibrium of the model, including the distribution of population across locations (L_n) and each location's own trade share for final goods (π_{nns}) and tasks (λ_{nns}). In the next two sections, we use the model's predictions to interpret the data.

2.6 Structural Transformation across Sectors and Occupations

One key feature of our data is a reallocation of employment towards more interactive occupations over time, which occurs across all locations and both between and within sectors. As in the macroeconomics literature on structural transformation, between-sector reallocations of employment for all locations can be explained by differences in rates of productivity growth across sectors and inelastic demand between sectors (e.g. Ngai and Pissarides 2007, Yi and Zhang 2010). To show this formally, partition final goods productivity in a sector-location into a sector component (\tilde{T}_s), a location component (\tilde{T}_n) and a residual (\tilde{T}_{ns}): $T_{ns} = \tilde{T}_s \tilde{T}_n \tilde{T}_{ns}$. Since the sector component \tilde{T}_s is common to all locations, it cancels from the numerator and denominator of the location trade share (π_{nis}) and hence does not directly effect π_{nis} for

given wages. But this sector component \tilde{T}_s directly affects the share of sectors in aggregate expenditure (E_{ns}) for all locations. Taking the partial derivative of this expenditure share (17) with respect to \tilde{T}_s at the initial equilibrium vectors of wages (\mathbf{w}) and employment (\mathbf{L}_{so}), faster productivity growth in sector s reduces the share of that sector in expenditure and increases the share of all other sectors in expenditure when sectors are complements and has the reverse effect when sectors are substitutes:

$$\begin{aligned} \left. \frac{\partial E_{ns}}{\partial \tilde{T}_s} \frac{\tilde{T}_s}{E_{ns}} \right|_{\mathbf{w}, \mathbf{L}_{so}} &= - \left(\frac{1-\beta}{\theta_s} \right) (1 - E_{ns}) < 0, & 0 < \beta < 1, \\ \left. \frac{\partial E_{nr}}{\partial \tilde{T}_s} \frac{\tilde{T}_s}{E_{nr}} \right|_{\mathbf{w}, \mathbf{L}_{so}} &= \left(\frac{1-\beta}{\theta_s} \right) E_{nr} > 0, & r \neq s, 0 < \beta < 1, \end{aligned}$$

where this productivity growth and the changes in expenditure shares it induces in turn have general equilibrium effects for wages and employment allocations.

While the macroeconomics literature on structural transformation typically focuses on sectors, secular changes in the shares of occupations in employment within sectors also can be explained by differences in productivity growth across occupations. To show this formally, partition average input productivity in an occupation-sector-location into an occupation component (\tilde{U}_o), sector component (\tilde{U}_s), location component (\tilde{U}_n) and a residual (\tilde{U}_{nso}): $U_{nso} = \tilde{U}_o \tilde{U}_s \tilde{U}_n \tilde{U}_{nso}$. Since the occupation component \tilde{U}_o is common to all locations, it cancels from the numerator and denominator of the location trade share (λ_{niso}) and hence does not directly effect λ_{niso} for given wages. But this occupation component \tilde{U}_o directly affects the share of occupations in sectoral expenditure (e_{nso}) for all locations. Taking the partial derivative of this expenditure share (13) with respect to \tilde{U}_o at the initial equilibrium vectors of wages (\mathbf{w}) and employment (\mathbf{L}_{so}), faster productivity growth in occupation o reduces the share of that occupation in costs and increases the share of all other occupations in costs when occupations are complements and has the reverse effect when occupations are substitutes:

$$\begin{aligned} \left. \frac{\partial e_{nso}}{\partial \tilde{U}_o} \frac{\tilde{U}_o}{e_{nso}} \right|_{\mathbf{w}, \mathbf{L}_{so}} &= - \left(\frac{1-\mu_s}{\epsilon_{so}} \right) (1 - e_{nso}) < 0, & 0 < \mu_s < 1, \\ \left. \frac{\partial e_{nsm}}{\partial \tilde{U}_o} \frac{\tilde{U}_o}{e_{nsm}} \right|_{\mathbf{w}, \mathbf{L}_{so}} &= \left(\frac{1-\mu_s}{\epsilon_{so}} \right) e_{nsm} > 0, & m \neq o, 0 < \mu_s < 1, \end{aligned}$$

where this productivity growth and the changes in expenditure shares it induces in turn have general equilibrium effects for wages and employment allocations.

Thus differences in productivity growth between interactive and non-interactive occupations can generate a reallocation of employment towards more-interactive occupations for all locations.

2.7 Trade in Tasks and Final Goods

Another key feature of our data is that the reallocation of employment towards more interactive occupations over time is more pronounced in more densely-populated locations. The model provides a natural explanation for these differences in the evolution of employment patterns across locations in the form of reductions in trade costs for final goods and tasks and specialization according to comparative advantage.

The pattern of trade across occupations can be characterized by a double difference across exporting locations and occupations within an importing location. The first difference computes the ratio of exports of tasks from two locations i and k in a third market n in a single occupation; the second difference compares this ratio of exports of tasks for two separate occupations o and m :

$$\frac{\lambda_{niso}/\lambda_{nkso}}{\lambda_{nism}/\lambda_{nksm}} = \frac{[U_{iso}L_{iso}^{\chi_{so}}(\tau_{niso}w_i)^{-\epsilon_{so}}] / [U_{kso}L_{kso}^{\chi_{so}}(\tau_{nkso}w_k)^{-\epsilon_{so}}]}{[U_{ism}L_{ism}^{\chi_{so}}(\tau_{nism}w_i)^{-\epsilon_{sm}}] / [U_{ksm}L_{ksm}^{\chi_{so}}(\tau_{nksm}w_k)^{-\epsilon_{sm}}]}.$$

From the above double difference, locations export relatively more tasks in occupations in which they have relatively lower costs of supply, where these costs of supply depend on relative productivities (which in turn depend on relative employments through the external economies of scale), relative wages and task trade costs. Depending on the pattern of relative costs of supply, each location is a net exporter of tasks in some occupations and a net importer of tasks in other occupations (inter-occupation trade in tasks). Even if a location is a net importer of tasks in an occupation, it still exports some tasks within that occupation for which it has high productivity draws. Similarly, even if a location is a net exporter of tasks in another occupation, it still imports some tasks within that occupation for which it has low productivity draws. Therefore the model also features two-way exporting and importing of tasks within occupations (intra-occupation trade in tasks).

The pattern of trade across sectors can be characterized by an analogous double difference across exporting locations and sectors within an importing location. The first difference computes the ratio of exports of final goods from two locations i and k in a third market n in a single sector; the second difference compares this ratio of exports of final goods for two separate sectors s and r :

$$\frac{\pi_{niss}/\pi_{nkss}}{\pi_{nirr}/\pi_{nkrr}} = \frac{[T_{is}L_{is}^{\eta_s}(d_{niss}\Phi_{is}w_i)^{-\theta_s}] / [T_{ks}L_{ks}^{\eta_s}(d_{nkss}\Phi_{ks}w_k)^{-\theta_s}]}{[T_{ir}L_{ir}^{\eta_r}(d_{nirr}\Phi_{ir}w_i)^{-\theta_r}] / [T_{kr}L_{kr}^{\eta_r}(d_{nkrr}\Phi_{kr}w_k)^{-\theta_r}]}.$$

From the above double difference, locations export relatively more final goods in sectors in which they have relatively lower costs of supply, where these costs of supply depend on relative productivities (which in turn depend on relative employments through the external economies of scale), relative unit costs (which depend on wages and trade in tasks), and final goods trade costs. Depending on the pattern of relative costs of supply, each location is a net exporter of final goods in some sectors and a net importer of final goods in other sectors (inter-industry trade in goods). Even if a location is a net importer of final goods in a sector, it still exports some final goods within that sector for which it has high productivity draws. Similarly, even if a location is a net exporter of final goods in another sector, it still imports some final goods within that sector for which it has low productivity draws. Therefore the model also features two-way exporting and importing of final goods within sectors (intra-industry trade in goods).

Our long historical time period encompasses the development of new communication technologies (e.g. telephones) and transport technologies (e.g. roads and the automobile) that are likely to have reduced both

final goods and task trade costs. Reductions in final goods trade costs (d_{nis}) induce specialization across sectors according to standard theories of comparative advantage. Reductions in task trade costs (τ_{niso}) induce an analogous process of specialization across occupations within sectors. When task trade costs are prohibitively high, all tasks are performed in the location in which the final good is produced. As task trade costs fall, it becomes feasible to unbundle production across locations and trade tasks between these locations.¹² To the extent that densely-populated locations are relatively more productive in interactive tasks (e.g. because agglomeration forces χ_{so} are stronger for interactive tasks), reductions in task trade costs induce densely-populated locations to specialize in more-interactive occupations, while more sparsely-populated locations specialize in less-interactive occupations.¹³ According to this explanation, densely-populated locations are always relatively more productive in interactive occupations, but it is only as task trade costs fall that it becomes feasible for them to specialize and reallocate employment within sectors towards these more interactive occupations.

3 Data Description

Our empirical analysis uses two main sources of data. The first is individual-level records from the U.S. Population Census for twenty-year intervals from 1880-2000 from Integrated Public Use Microdata Series (IPUMS): see Ruggles et al. (2010). These census micro data report individuals' location, occupation and sector, as well as other demographic information. We use these data to determine whether an individual is located in a metro area as well as the occupation and sector in which an individual is employed.¹⁴ We weight individuals by their person weights to ensure the representativeness of the sample. Our main dataset is a panel from 1880-2000 that uses information on the share of employment within an occupation and sector in metro areas, for which the 1 percent IPUMS samples are representative. To provide evidence on improvements in communication and transportation technologies, we also use long-differenced data from 1880-1930 that uses information on employment by occupation, sector and county, for which we use samples of 10 percent for 1880 and 5 percent for 1930 to again ensure representativeness.

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

¹²For further discussion of the increased unbundling of production, see for example Baldwin (2012).

¹³While our model focuses on trade in tasks across locations within countries, a similar process of trade in tasks could also occur between countries. To the extent that there is greater offshoring of tasks in less interactive occupations from metro areas than from non-metro areas, this provides a related explanation in terms of the same mechanism for the increased concentration of employment in interactive occupations in metro areas relative to non-metro areas over time.

¹⁴Metro areas are defined in IPUMS based on Census Bureau Metropolitan Statistical Areas (MSAs).

(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 and 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 task specialization in urban and rural areas over time, and agriculture is unsurprisingly overwhelmingly located in rural areas.¹⁶ While our baseline sample uses time-varying boundaries of metro areas to ensure that these correspond to meaningful economic areas, we also report robustness checks using administrative cities whose boundaries are more stable over time.

Our second main data source is the Dictionary of Occupational Titles (U.S. Department of Labor 1991), which contains detailed descriptions of more than 12,000 occupations. Following Autor et al. (2003), previous research using DOTs typically uses the numerical scores that were constructed for each occupation by the Department of Labor (e.g. a Non-routine Interactive measure based on the Direction, Control and Planning of Activities (DCP) numerical score). In contrast, we use verbs from the detailed occupational descriptions in DOTs to directly measure the tasks performed by workers in each occupation. We use a list of over 3,000 English verbs from “Writing English,” a company that offers English language consulting.¹⁷ This approach enables us to provide a rich analysis of the tasks undertaken in urban and rural areas using the 3,000 verbs and 12,000 occupational descriptions without being restricted to the numerical scores. Nonetheless, we also compare our measures of occupational characteristics to those from the numerical scores. We match the DOTs occupations to the three-digit occupations in our census data using the cross-walk developed by Autor, Levy and Murnane (2003). In our baseline specification, we use a time-invariant measure of tasks based on the occupational descriptions from the digital edition of the 1991 DOTs, which ensures that our results are not driven by changes in language use over time. In sensitivity checks, we also report results using digitized occupational descriptions from the first edition of the DOTs in 1939 (U.S. Department of Labor 1939).

We complement these two main data sources with information from a variety of other sources. We use the standard reference for word usage in English (Roget’s Thesaurus) to quantify the meanings of verbs from the occupational descriptions.¹⁸ We use ArcGIS shapefiles from the National Historical Geographical Information System (NHGIS) to track the evolution of county boundaries over time. We also use measures of improvements in transport and communication technologies. We measure the length of roads in each county using a georeferenced 1931 road map (Gallup 1931).¹⁹ At the beginning of our sample in 1880,

¹⁵See 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. All our results are robust to restricting attention to occupations and sectors that are present in all years.

¹⁶Our key findings, however, are robust to the inclusion of these agricultural workers. For further analysis of the relationship between urbanization and structural transformation away from agriculture, see Michaels et al. (2012).

¹⁷See <http://www.writingenglish.com/englishverbs.htm>.

¹⁸We use the online computer-searchable edition of Roget (1911): <http://machaut.uchicago.edu/rogets>.

¹⁹Recent economics research on the U.S. road network has largely concentrated on the later development of the interstate highway system, as in Baum-Snow (2007), Michaels (2008) and Duranton and Turner (2012).

most U.S. roads were little more than dirt tracks (see, for example, Swift 2011) and widespread paved road construction only occurred following the Federal Aid Road Act of 1916 and the Federal Highway Act of 1921. Therefore we use the 1931 map to construct a measure of the growth of the paved road network from 1880-1930. We measure the number of residence telephones in each county in 1935 using newly-digitized data from American Telephone and Telegraph Company (AT&T 1935). The telephone was not patented until 1876 just before the beginning of our sample period and the telephone network developed rapidly from 1890 onwards (see, for example, Fischer 1992). Therefore we use the data on telephones to construct a measure of the growth of telephones from 1880-1930. To address the concern that the road network could be influenced by changes in the interactiveness of economic activity, we use an instrument based on the “Pershing” map of highway routes of military importance for coastal and border defense. To address similar concerns for the telephone, we use an instrument based on proximity to primary and secondary outlets on AT&T’s long distance trunk network, whose construction was influenced by the strategic objective of connecting the nation as a whole.

4 Specialization Across Occupations and Sectors

We begin by providing some motivating evidence of changes in specialization across occupations and sectors in metro areas relative to non-metro areas. To do so, we estimate the following regression for each year t separately using data across occupations o and sectors s :

$$\text{MetroShare}_{ost} = \mu_{ot} + \eta_{st} + \varepsilon_{ost}, \quad (20)$$

where MetroShare_{ost} is the share of employment in metro areas in occupation o , sector s and year t ; observations are weighted by person weights; μ_{ot} are occupation-year fixed effects; η_{st} are sector-year fixed effects; and ε_{ost} is a stochastic error. We normalize the sector-year and occupation-year fixed effects so that they each sum to zero in each year, and hence they capture deviations from the overall mean in each year. While we estimate the above regression using a share as the left-hand side variable so that the estimated coefficients have a natural interpretation as frequencies, we find a very similar pattern of results in a robustness test in which we use a logistic transformation of the left-hand side variable: $\text{MetroShare}_{ost}/(1 - \text{MetroShare}_{ost})$.

The occupation-year fixed effects (μ_{ot}) capture the average probability of being in a metro area for workers in each occupation in each year, after controlling for differences across sectors in metro probabilities. Similarly, the sector-year fixed effects (η_{st}) capture the average probability of being located in a metro area for workers in each sector in each year, after controlling for differences across occupations in metro probabilities. The sector and occupation fixed effects are separately identified because there is substantial overlap in occupations and sectors, such that each sector contains multiple occupations and each occupa-

tion is employed in several sectors.²⁰ We estimate this regression using both the aggregate (two-digit) and disaggregate (three-digit) definitions of occupations and sectors discussed above.

As reported in Table 1 for two-digit occupations and sectors, we find substantial changes in specialization across occupations and sectors in metro areas relative to non-metro areas over time. From Panel A, in 1880, “Clerical and Kindred” workers were the most likely to be located in metro areas. In contrast, by 2000, “Clerical and Kindred” workers were ranked only fourth, and “Professional and Technical” workers were the most likely to be located in metro areas. From 1880-2000, declines in ranks were observed for “Craftsmen” (from 2 to 6) and “Operatives” (from 3 to 7), while increases in ranks were observed for “Professional and Technical” workers (from 7 to 1) and “Managers, Officials and Proprietors” (from 6 to 3). As apparent from the first and fourth columns of the table, these changes in ranks reflect substantial changes in the probabilities of workers in individual occupations being located in metro areas over time.

Since the regression (20) includes sector-year fixed effects, these changes in the metro probabilities for each occupation are not driven by changes in sector composition, but rather reflect changes in the organization of economic activity within sectors. Nonetheless, we also observe substantial changes in sector structure in metro areas relative to non-metro areas over time. From Panel B, declines in ranks from 1880-2000 were observed for “Wholesale and Retail Trade” (from 2 to 6) and “Manufacturing” (from 4 to 10). In contrast, increases in ranks from 1880-2000 were observed for “Transportation, Communication and Other Utilities” (from 6 to 3) and “Business and Repair Services” (from 9 to 1).

In Figures A1 and A2 of the web appendix, we show the evolution of the occupation and sector coefficients across each of the twenty-year intervals in our data. While “Professional and Technical” workers display an increased propensity to locate in metro areas from 1880-1960, the probability that “Managers, Officials and Proprietors” are located in urban areas increases particularly sharply from 1940-2000. In contrast, the likelihood that “Craftsmen” are found in metro areas declines throughout our sample period, while the probability for “Clerical and Kindred” workers declines from 1900 onwards, and the probability for “Service” workers initially rises until 1920 and later declines until around 1960.

Such changes in specialization are not limited to the aggregate categories considered so far, but are also found using more disaggregated measures of occupations and sectors. In Table A1 of the web appendix, we report the results of estimating the regression (20) including three-digit-occupation-year and three-digit-sector-year fixed effects. Panels A and B report the twenty occupations within the largest increases and decreases respectively in the within-sector probability of being located in a metro area from 1880-2000. Both the top agglomerating occupations in Panel A and the top dispersing occupations in Panel B are diverse and span multiple sectors. For example, the top agglomerating occupations include “Editors and Reporters”, “Judges and Lawyers” and “Pattern and Model Makers,” while the top dispersing occupations contain “Of-

²⁰The average three-digit sector employs workers from 111 three-digit occupations, while the average three-digit occupation contains workers employed in 81 sectors.

“Office Machine Operators” and “Upholsterers.” In our empirical analysis below, we provide evidence on the systematic characteristics shared by occupations that agglomerate versus disperse over time.

5 Measuring the Tasks Undertaken by Occupations

To measure the tasks undertaken by each occupation, we use the detailed descriptions of more than 12,000 disaggregated occupations included in the DOTs. We use the verbs from each occupation’s description to measure the tasks performed by workers within that occupation, because verbs capture an action (bring, read, walk, run, learn), an occurrence (happen, become), or a state of being (be, exist, stand), and hence capture the task being performed. To focus on persistent characteristics of occupations and abstract from changes in word use over time, our baseline analysis uses time-invariant occupational descriptions from the 1991 digital edition of the DOTs. While the tasks undertaken within each occupation can change over time, the relative task content of occupations is likely to be more stable. To provide evidence on the extent to which this is the case, we have also digitized the occupational descriptions from the first edition of the DOTs in 1939. Although the descriptions of occupations are less detailed and the boundaries between occupations are less clear in the historical DOTs, we find a similar pattern of results using both sets of occupational descriptions and provide evidence below on the correlation of the relative task content of occupations over time.

The first step of our procedure uses a list of over 3,000 English verbs from “Writing English,” a company that offers English language consulting. Using this list of verbs, we search each occupational description in the 1991 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).²¹ For example, 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.²²

²¹An emerging literature in economics and the social sciences uses textual search as the basis for quantitative analysis: see for example Gentzkow et al. (2012) on political influence and Michel et al. (2011) on culture.

²²As an indication of the wide coverage of our list of over 3,000 verbs, only 1,830 appear in the 1991 DOTs occupational

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.

From this first step, we obtain the number of occurrences of each verb for each DOTs occupation. We next match the more than 12,000 DOTs occupations to IPUMS standardized 1950 occupations using the crosswalk developed by Autor, Levy and Murnane (2003). Finally, we calculate the frequency with which each verb v is used for each IPUMS occupation o :

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

where we focus on the frequency rather than the number of verb uses to capture the relative importance of tasks for an occupation and to control for potential variation in the length of the occupational descriptions matched to each IPUMS occupation.

We provide evidence on changes in task specialization in metro areas relative to non-metro areas over time by estimating the following regression for each verb v and year t separately using data across occupations o and sectors s :

$$\text{MetroShare}_{ost} = \alpha_{vt} \text{VerbFreq}_{vo} + \eta_{vst} + \varepsilon_{ost}, \quad (21)$$

where MetroShare_{ost} is again the share of employment in metro areas in occupation o , sector s and year t ; VerbFreq_{vo} is defined above for verb v and occupation o ; η_{vst} are verb-sector-year fixed effects; 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 verb-sector-year fixed effects (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 Panels A and B of Table 2, we report for each year the ten verbs with the highest and lowest standardized coefficient α_{vt} (the estimated coefficient multiplied by the standard deviation of VerbFreq_{vo}).²³ As apparent from Panel A, we find substantial changes in the tasks most concentrated in metro areas within sectors over time. In 1880, the verbs with the highest metro employment shares typically involve physical tasks such as “Braid,” “Sew,” “Stretch” and “Thread.” By 1920, the top ten verbs include an increased number of clerical tasks, such as “Bill,” “File,” “Notice,” and “Record.” By 1980 and 2000, the leading metro

descriptions.

²³We find a similar pattern of results just using the estimated coefficients instead of the estimated coefficients times the standard deviation of VerbFreq_{vo} .

verbs include a proliferation of interactive tasks, such as “Analyze,” “Advise,” “Confer” and “Report.” As shown in Panel B, we also find some changes in the tasks least concentrated in metro areas, although here the pattern is less clear cut (e.g. “Tread” appears from 1880-1960 and “Turn” appears from 1960-2000).

6 Quantifying Task Specialization

The approach developed in the previous section allows us to provide a detailed characterization of the tasks performed in urban and rural areas using the 3,000 verbs and 12,000 occupational descriptions. In this section, we develop a quantitative measure of task specialization based on the meanings of these verbs. To do so, we use the online computer-searchable version of Roget’s Thesaurus (1911), which has been the standard reference for English language use for more than a century, and 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 this classification is that it explicitly takes into account that words can have different meanings depending on context by including extensive cross-references to link related groups of words.²⁴

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 I (Abstract Relations) deals with ideas such as number, order and time; Class II (Space) is concerned with movement, shapes and sizes; and Class III (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 IV, Intellect), the human will (Class V, Volition) and the human heart and soul (Class VI, Emotion, Religion and Morality).

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

$$\text{ThesFreq}_{vk} = \frac{\text{Appearances of verb } v \text{ in subdivision } k \text{ of thesaurus}}{\text{Total appearances of verb } v \text{ in thesaurus}}, \quad (22)$$

where our use of a frequency 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.²⁵

²⁴For further discussion of the genesis of Roget’s Thesaurus, see for example Hüllen (2003).

²⁵For example, the verb “Consult” appears in six thesaurus Categories. The entry followed by a comma is 695 Advice, which

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

$$\text{TaskContent}_{ko} = \sum_{v \in V} \text{VerbFreq}_{vo} \times \text{ThesFreq}_{vk}.$$

We use this measure to examine changes in task specialization in metro areas relative to non-metro areas over time by estimating an analogous regression for each thesaurus subdivision k and year t as for each verb and year in the previous section:

$$\text{MetroShare}_{ost} = \beta_{kt} \text{TaskContent}_{ko} + \eta_{kst} + \varepsilon_{ost}, \quad (23)$$

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

The coefficient of interest β_{kt} again captures a conditional correlation: the correlation between occupations’ shares of employment in metro areas and their frequency of use of verbs in thesaurus subdivision k . The thesaurus-subdivision-sector-year fixed effects (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 TaskContent_{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 3, we report the estimation results for the thirty-eight Sections of the thesaurus. We calculate the standardized coefficient for each Section of the thesaurus (the estimated coefficient β_{kt} multiplied 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).²⁶ 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. The table highlights the top-five thesaurus Sections in 1880 in bold-italics and the top-five thesaurus Sections in 2000 in bold.

The results in Table 3 reveal a sharp change the relative ranking of thesaurus Sections involving the external world (Classes I-III) and those involving the internal world of human beings (Classes IV-VI). In

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]”).

²⁶Again we find a similar pattern of results using just the estimated coefficient instead of the estimated coefficient times the standard deviation of TaskContent_{ko} .

1880, the top-five thesaurus Sections most concentrated in metro areas were: Quantity (Class I), Time (Class I), Matter in General (Class III), Dimensions (Class II), and Inorganic Matter (Class III). In contrast, in 2000, the top-five thesaurus Sections were: Results of Reasoning (Class IV), Means of Communicating Ideas (Class V), Moral Affections (Class VI), Voluntary Action (Class IV) and Precursory Conditions and Operations (Class IV). The correlation between the rankings of the thesaurus sections in 1880 and 2000 is negative and statistically significant (-0.43).

Positive changes in ranks in Table 3 are typically concentrated in thesaurus Classes IV and V, which correspond to the human mind and the human will respectively. These Classes include Class IV, Division 1 (Formation of Ideas), Class IV, Division 2 (Communication of Ideas) and Class V, Division 2 (Intersocial Volition). We summarize this combination of tasks – thought, communication and intersocial activity – as “interactiveness.” We use as our baseline measure of the interactiveness of a verb whether it appears in Classes IV and V of the thesaurus and measure the interactiveness of an occupation using the verbs in its occupational description. Specifically, we measure the interactiveness of an occupation using the frequency with which verbs appear in that occupation’s description and the frequency with which those verbs appear in thesaurus Classes IV and V:

$$\text{Interactive}_o = \sum_{v \in V} \text{FreqVerb}_{vo} \times \text{FreqInteractive}_v, \quad (24)$$

where FreqVerb_{vo} is the frequency with which verb v is used for occupation o from above; FreqInteractive_v is the frequency with which verb v appears in thesaurus Classes IV and V (computed as in (22)). We also report results below breaking out interactiveness into thought (Class IV, Division 1), communication (Class IV, Division 2) and intersocial activity (Class V, Division 2).

In Panels A and B of Table 4, we report the top ten and bottom ten interactive occupations according to this measure. While any single quantitative measure of interactiveness is unlikely to capture the full meaning of this concept, the occupations identified by our procedure as having high and low levels of interactiveness appear intuitive. “Buyers and Department Heads”, “Clergymen” and “Pharmacists” arguably perform more interactive tasks than “Blasters and Powdermen”, “Roofers and Slaters” and “Welders and Flame Cutters.”

In Figure 1, we measure the interactiveness of metro areas, non-metro areas and the economy as a whole using the employment-weighted average of interactiveness for each occupation. In this measure, interactiveness only differs between metro and non-metro areas to the extent that they have different distributions of employment across occupations:

$$\text{Interactive}_{jt} = \sum_{o=1}^O \frac{E_{ojt}}{E_{jt}} \text{Interactive}_o, \quad j \in \{M, N\}, \quad (25)$$

where j indexes a type of location and we again denote metro areas by M and non-metro areas by N ; E_{ojt} corresponds to employment in occupation o in location type $j \in \{M, N\}$ in year t .

In 1880, metro and non-metro areas have similar levels of interactiveness, with if anything non-metro areas having higher interactiveness than metro areas. Over time, interactiveness increases in both sets of locations, but this increase is greater in metro areas than in non-metro areas. This increase in the relative interactiveness of metro areas is particularly sharp from 1900-1920, which coincides with the dissemination of improvements in communication and transport technologies in the form of the telephone and roads and the automobile. The model can account for this increase in the interactiveness of employment in both sets of locations in terms of more rapid productivity growth in non-interactive occupations and inelastic demand between occupations. To the extent that this differential productivity growth is stronger in metro than in non-metro areas, this mechanism could also account for the relative increase in the interactiveness of metro areas. Alternatively, if densely-populated locations are relatively more productive in interactive occupations and improvements in communication and transport technologies make it feasible to specialize and reallocate employment towards these occupations, this provides another potential explanation for the relative increase in the interactiveness of employment in metro areas. In our empirical analysis below, we provide direct evidence on the role played by improvements in communication and transport technologies in the period for which we observe the largest increase in the relative interactiveness of metro areas.

7 Robustness

Having presented our baseline evidence of an increase in the interactiveness of employment in metro areas relative to non-metro areas over time, we now document the robustness of this finding across a large number of different samples and specifications.

7.1 1939 DOTs

Our baseline specification measures the task content of employment using time-invariant occupational descriptions from the 1991 DOTs. While this approach ensures that our findings are not driven by changes in language use over time, it assumes that the relative task content of occupations is persistent over time. One concern is that the interactiveness of occupations could have changed over time and these changes in interactiveness could be correlated with occupations' shares of employment in metro areas.

To address this concern, we replicated our analysis using the first edition of the DOTs from 1939. We digitized the more than 12,000 occupational descriptions in the 1939 DOTs and implemented our procedure of searching for verbs in each occupational description. The boundaries between occupations are less well defined and the occupational descriptions are less detailed in the 1939 DOTs, which implies that the resulting measures of the task content of employment are likely to be less precise than those using the 1991 DOTs. Nonetheless, as reported in Table A2 of the web appendix, we find similar changes in task specialization in this robustness test. The verbs most correlated with metro employment shares in 1880 include physical

tasks such as “Slot,” “Thread,” “Straighten” and “Stitch.” In contrast, the verbs most correlated with metro employment shares in 2000 include interactive tasks such as “Advise,” “Present,” “Question” and “Report.”

Using the verbs from the 1939 occupational descriptions and the frequency with which these verbs appear in Classes IV and V of the thesaurus, we again find increase in the interactiveness of employment over time that is more rapid in metro areas than in non-metro areas, as shown in Figure A3 in the web appendix. This similarity of the results using both the 1939 and 1991 occupational descriptions suggests that our findings are unlikely to be driven by changes in the relative interactiveness of occupations over time. Indeed, although the layout of the occupational descriptions implies that our measure of interactiveness using the 1939 DOTs is less precise than our baseline measure using the 1991 DOTs (which by itself would induce an imperfect correlation), we find that they are positively and statistically significantly correlated. As reported in Table A3 of the web appendix, the unweighted correlation coefficient between the 1939 and 1991 measures across the sample of occupations in 2000 is 0.622.

7.2 Metro Areas and Administrative Cities

Our analysis has so far used variation between metro and non-metro areas. To provide further evidence of a relative increase in the interactiveness of employment in densely-populated locations, we now present evidence using a different source of variation across metro areas of differing population densities.

In Panels A and B of Figure 2, we display mean interactiveness for each metro area (as calculated using (25)) against log population density for 1880 and 2000 respectively, as well as the fitted values and confidence intervals from locally-weighted linear least squares regressions. To ensure that metro areas correspond to meaningful economic units, we use time-varying definitions of metro areas from IPUMS, and hence the number of observations changes over time as new metro areas enter the sample. In 1880, we find little relationship between interactiveness and log population density across metro areas, which is reflected in a negative but statistically insignificant OLS coefficient (standard error) of -0.0002 (0.0013). In contrast, in 2000, we find a positive and statistically significant relationship between interactiveness and log population density, which is reflected in an OLS coefficient (standard error) of 0.0018 (0.0002). Even when we restrict the 2000 sample to metro areas that exist in 1880, we continue to find a positive relationship that is statistically significant at the 10 percent level, confirming that these findings are not driven by a change in the composition of metro areas. Therefore the increase in the relative interactiveness of densely-populated locations over time is observed not only comparing metro and non-metro areas but also comparing metro areas of differing population densities. Metro areas with relatively high levels of interactiveness conditional on population density in 2000 include Boston (BOS, MA) and New York (NYC, CT/NY/NJ), while those with low levels of interactiveness conditional on population density include Anniston (ANN, AL) and Mansfield (MAN, OH).

In our baseline specification, we use time-varying definitions of the boundaries of metro areas, which ensures that they correspond to meaningful economic units. One concern is that the change in the relative interactiveness of metro areas could be driven by a change in the geographical boundaries of metro areas as they have expanded to include surrounding suburbs. To address this concern, we replicated our analysis using an alternative definition of urban areas as administrative cities, which have much more stable geographical boundaries over time. Again we find an increase in the relative interactiveness of urban areas over time, whether we compare administrative cities to all other locations (Figure A4 in the web appendix) or only to non-metro areas (Figure A5 in the web appendix). Therefore the increase in the relative interactiveness of urban areas reflects a change in the organization of economic activity within existing geographical boundaries.

7.3 Other Occupational Characteristics

Our approach of using verbs from the occupational descriptions enables us to provide a detailed characterization of task specialization in urban and rural areas over a long historical time period. Based on this detailed characterization, we have developed a new measure of the interactiveness of occupations, which we now compare with existing measures of occupational characteristics, including the numerical scores from the DOTs used by Autor, Levy and Murnane (2003). Since these numerical scores are not available in the first edition of the DOTs in 1939, we use their values from the 1991 digital edition of the DOTs.

In Table A3 of the appendix, we report the correlation coefficients between interactiveness and other measures of occupation characteristics across the sample of occupations in 2000. We report both unweighted correlations and correlated weighted by occupation employment. The highest correlation coefficients are for the Non-routine Interactive (Direction, Control and Planning (DCP)) and Non-routine Analytic (GED-MATH) used by Autor, Levy and Murnane (2003). While both of these measures are related to the concepts of thought, communication and intersocial activity captured by our interactiveness measure, the correlations are around 0.5. Therefore our interactiveness measure captures distinctive information about the tasks performed by workers within occupations. While DCP is orientated towards top-down interactions between workers (e.g. between a manager and her subordinates), our measure captures all interactions between workers (e.g. between members of a product design team). While GED-MATH is orientated towards thought, our measure of interactiveness also captures communication and intersocial activity.

8 Explaining Increased Interactiveness

Having demonstrated the robustness of the increase in the relative interactiveness of metro areas across a number of different samples and specifications, we now provide further evidence on explanations for the observed change in interactiveness. First, we decompose the overall change in interactiveness into

the contributions of individual occupations and sectors, which enables us to explore explanations that emphasize particular occupations and sectors. Second, we report regression specifications using variation in interactivens between sectors, within sectors, and within sectors and occupations over time. Using these regressions, we explore the importance of the constituent components of interactivens (thought, communication and intersocial) and present evidence on a number of potential explanations. Third, we provide direct evidence on the role played by improvements in communication and transport technologies in the period for which we observe the largest increase in the relative interactivens of metro areas.

8.1 Decomposing Interactivens

We begin by decomposing the change in the overall interactivens of metro and non-metro areas into the contributions of each two-digit occupation and sector. Overall interactivens for metro and non-metro areas is the employment-weighted average of interactivens for each two-digit-sector-occupation cell:

$$I_{jt} = \sum_{z \in \Omega} \sum_{o \in \Omega_z} \frac{E_{ojt}}{E_{jt}} I_o, \quad j \in \{M, N\}, \quad (26)$$

where z indexes two-digit-sector-occupation cells; o indexes disaggregated three-digit occupations within these cells; and t indexes time; Ω is the set of two-digit-sector-occupation cells; Ω_z is the set of three-digit occupations within each cell z ; the interactivens of each three-digit occupation is measured using (24) based on the time-invariant occupational descriptions from the 1991 DOTs.

Taking differences between times T and $t > T$, the change in the overall interactivens of metro and non-metro areas can be decomposed as follows:

$$\Delta I_{jt} = \sum_{z \in \Omega} \sum_{o \in \Omega_z} \left[\Delta \left(\frac{E_{ojt}}{E_{jt}} \right) \right] I_o, \quad j \in \{M, N\}, \quad (27)$$

where $\Delta I_{jt} = I_{jt} - I_{jT}$; $\Delta (E_{ojt}/E_{jt})$ is the change in the employment share of occupation o in location $j \in \{M, N\}$; and we have used the fact that occupation interactivens is constant over time. Taking differences again between metro and non-metro areas, we obtain an analogous decomposition of the change in the relative interactivens of metro and non-metro areas:

$$\Delta I_{Mt} - \Delta I_{Nt} = \sum_{z \in \Omega} \sum_{o \in \Omega_z} \left[\Delta \frac{E_{oMt}}{E_{Mt}} - \Delta \frac{E_{oNt}}{E_{Nt}} \right] I_o, \quad (28)$$

where the right-hand sides of the decompositions (27) and (28) are summations over the contributions from each two-digit-sector-occupation-cell. These contributions correspond to a matrix with two-digit sectors for rows and two-digit occupations for columns, where the right-hand side is a summation across both rows and columns. Metro areas display a larger increase in interactivens than non-metro areas to the extent that they experience a greater reallocation of employment shares towards high-interactivens occupations.

Figure 3 summarizes the results from the decompositions of the change in the relative interactiveness of metro and non-metro areas (28). Panels A and B show the contributions for each two-digit occupation (summing across sectors in the rows of the matrix of contributions) for each twenty-year interval in our sample, while Panels C and D show the corresponding contributions for each two-digit sector (summing across occupations in the columns of the matrix of contributions).²⁷ Figures A6 and A7 in the web appendix report analogous results from the decompositions of the change in interactiveness for metro and non-metro areas separately (27).

Panels A and B of Figure 3 show that the sharp increase in the relative interactiveness of metro areas from 1880-1920 is largely driven by positive contributions from Clerks (and to a lesser extent Professionals), with Operatives, Sales Workers and Managers all making negative contributions. From 1920-1960, Professionals (and to a lesser but growing extent Managers) make the largest positive contributions, while Craftsmen and Service Workers make negative contributions. From 1960-2000, Professionals and Managers have the largest positive contributions, while Clerks have the largest negative contribution.

Panels C and D of Figure 3 show that Professional and Business services are the two sectors that make the largest contributions to the increase in the relative interactiveness of metro areas over the sample as a whole. Professional Services are more important earlier in the sample period, while Business Services become more important later on. The sector that makes the largest negative contribution over the sample period as a whole is Wholesale and Retail trade, with the absolute magnitude of its contribution diminishing over time. While the contribution from Manufacturing is initially positive (up to 1920), it becomes negative from 1940 onwards.

Taking these decomposition results together, the increase in the relative interactiveness of metro areas is not driven by any one occupation or sector. Our results are not solely explained by Managers (whose contribution only becomes positive towards the end of our sample period). Clerks and Professionals make notable positive contributions towards the beginning and end of our sample period respectively. Our results are also not simply driven by a decline of Manufacturing in urban areas (indeed Manufacturing was expanding in the early decades of our sample when some of the largest changes in interactiveness were observed). Similarly, our findings are not simply attributable to an expansion of Services in urban areas (indeed Services was a relatively small share of employment in the early decades of our sample when some of the largest changes in interactiveness were observed). Instead we find evidence of a pervasive reallocation of employment towards more interactive occupations within sectors.

²⁷Since the change in overall interactiveness is the sum across all elements in the matrix, adding the sums for occupations and the sums for sectors would result in double-counting (since each element would be counted twice).

8.2 Variation Within and Between Sectors

To further explore the determinants of the increase in the relative interactiveness of metro areas, we begin by examining between-sector variation. Sector interactiveness is measured as the employment-weighted mean of the interactiveness of each occupation:

$$\text{Interactive}_{st} = \sum_o \frac{E_{ost}}{E_{st}} \text{Interactive}_o,$$

We run a regression across sectors of the share of a sector's employment in metro areas (MetroShare_{st}) on its interactiveness (Interactive_{st}) for each year separately:

$$\text{MetroShare}_{st} = \alpha_t \text{Interactive}_{st} + \varepsilon_{st},$$

where ε_{st} is a stochastic error; α_t captures the correlation between sectors' shares of employment in metro areas and their interactiveness in each year. While we estimate the above regression and the remaining regressions in this section using a share as the left-hand side variable so that the estimated coefficients have a natural interpretation as frequencies, we again find a very similar pattern of results in a robustness test in which we use a logistic transformation of the left-hand side variable: $\text{MetroShare}_{st}/(1 - \text{MetroShare}_{st})$.

Panel A of Table 5 reports the results, where each cell in the table corresponds to a separate regression. In 1880, there is a negative but statistically insignificant correlation between a sector's metro employment share and its interactiveness. Starting in 1900, there is an increase in the correlation between a sector's metro employment share and its interactiveness, which is particularly sharp from 1900-1940, and becomes positive and statistically significant at conventional critical values in 1960. Therefore more interactive sectors become increasingly concentrated in metro areas over time.

Panel A of Table 5 also breaks out overall interactiveness into thought (Class IV, Division 1), communication (Class IV, Division 2) and intersocial (Class V, Division 2). As shown in the table, we find that the sectors increasingly concentrating in metro areas over time involve each of these components of interactiveness: thought, communication and intersocial activity.

These changes in patterns of specialization in metro versus non-metro areas are explained in the model by changes in the relative demand for occupations as a result of either differential changes in the relative productivity of occupations or improvements in communication and transport technology. An increase in the relative demand for an occupation raises both its employment and its wage (and hence raises its wagebill). In contrast, an increase in the relative supply of an occupation raises its employment but reduces its wage (and hence reduces its wagebill if the demand for occupations is inelastic). To assess the relative importance of these two explanations, Panel A of Table 5 also reports the results of regressions in which we use the share of the sector wagebill in metro areas as the left-hand side variable. Although the wage data are available for a much shorter time period than the employment data, we find a similar pattern of results using

this alternative left-hand side variable, which is consistent with relative demand moving relative wagebills and employment in the same direction.

Having established these relationships between sectors, we next examine within-sector variation. We run a regression across sectors and occupations of the share of a sector-occupation's employment in metro areas (MetroShare_{ost}) on occupation interactiveness (Interactive_o) for each year separately:

$$\text{Metro}_{ost} = \alpha_t \text{Interactive}_o + \eta_{st} + \varepsilon_{ost},$$

where η_{st} are sector-year fixed effects and ε_{ost} is a stochastic error. The sector-year fixed effects (η_{st}) control for changes in sector composition over time, so that the coefficient α_t is identified solely from variation within sectors. The coefficient α_t captures the within-sector correlation between the share of employment in metro areas and the interactiveness of occupations.

Panel B of Table 5 reports the results, where each cell in the table again corresponds to a separate regression. In line with our previous results, the correlation between metro employment shares and interactiveness is negative in 1880. Over time, this correlation becomes more positive and becomes statistically significant by 1960. Therefore, within sectors, more interactive occupations become increasingly concentrated in metro areas over time. This finding of the same pattern of reallocation across occupations both between and within sectors is consistent with a wide-ranging secular process favoring specialization in interactive occupations in metro areas.

Panel B of Table 5 also breaks out overall interactiveness into thought (Class IV, Division 1), communication (Class IV, Division 2) and intersocial (Class V, Division 2). Again we find that the increased concentration of employment in interactive occupations in metro areas is largely explained by an increased specialization of metro areas in occupations involving thought and communication. Panel B of Table 5 also reports the results of regressions in which we use the share of a sector-occupation's wagebill in metro areas (rather than its share of employment in metro areas) as the left-hand side variable. For the shorter period over which we have the wage data, we again find a similar pattern of results using this alternative left-hand side variable, which is consistent with relative demand moving relative wagebills and employment in the same direction.

Finally, to use variation within sectors and occupations, we pool our sector-occupation data over time and estimate a panel data regression that facilitates the inclusion of sector, occupation and year fixed effects. We regress the share of a sector-occupation's employment in metro areas on these fixed effects and interaction terms between time dummies and our measure of occupation interactiveness:

$$\text{MetroShare}_{ost} = \alpha_t [\text{Interactive}_o \times \text{Year}_t] + \mu_o + \eta_s + \delta_t + \varepsilon_{ost},$$

where ε_{ost} is a stochastic error; we choose 1880 as the excluded year from the interaction terms. The occupation fixed effects (μ_o) control for time-invariant differences between metro and non-metro areas in

the share of an occupation in employment and capture the main effect of occupation interactiveness. The sector fixed effects (η_s) control for time-invariant differences between metro and non-metro areas in the share of a sector in employment. The year fixed effects (δ_t) control for changes in the shares of metro areas in employment across all occupations and sectors. The coefficients α_t capture the change in the correlation between metro employment shares and interactiveness relative to 1880.

Table 6 reports the estimation results. Column (1) confirms our findings above of an increasing correlation between metro employment shares and occupation interactiveness over time, which becomes positive and statistically significant by 1960. As shown in Column (2), this increasing correlation between metro employment shares and occupation interactiveness is robust to replacing the sector and year fixed effects with sector-year fixed effects to control for changes in sector composition over time.

While the model's explanation for the increase in the relative interactiveness of metro areas emphasizes changes in the relative demand for occupations as a result of either differential changes in the relative productivity of occupations or improvements in communication and transport technology, Columns (3)-(4) consider an alternative explanation based on changes in female labor force participation. Over our long historical time period, female labor force participation increased substantially, which implies that more and more couples face a colocation problem where both partners are looking for work in a common location (e.g. Costa and Kahn 2000). Since solving such a colocation problem is likely to be easier in more densely-populated locations, one concern is that the movement of such "power couples" into densely-populated locations could be driving the increase in the relative concentration of employment in interactive occupations in metro areas. Although it is not necessarily the case that power couples work in interactive occupations, Columns (3) and (4) provide evidence against this concern by estimating the specification in Column (2) separately for single and married people. Comparing the two columns, we find a similar pattern of results irrespective of marital status, which suggests that our findings are not being driven by the location decisions of power couples.

In Columns (5)-(6), we provide further evidence against explanations based on individual sectors. In Column (5), we include only workers in the manufacturing sector and demonstrate a similar pattern of results, which corroborates that our findings are not simply being driven by the rise of the services sector in urban areas. In Column (6), we include only workers in the services sector, which confirms that our findings are not simply being driven by a decline in manufacturing in urban areas. In Columns (7)-(8), we examine the role of differences in human capital across cities. Glaeser and Resseger (2009) find that the positive average relationship between productivity and metro area population is driven by a strong positive relationship for more-skilled metro areas, whereas this relationship is almost non-existent for less-skilled metro areas. Using Glaeser and Resseger (2009)'s classification of metro areas by skill, Columns (7) and (8) re-estimate the specification in Column (2) excluding more and less-skilled metro areas respectively.²⁸

²⁸In Glaeser and Resseger (2009)'s classification, more-skilled Metropolitan Statistical Areas (MSAs) have a share of adults

In both samples, we find a positive and statistically significant increase in the relative concentration of employment in interactive occupations in metro areas over time. Therefore, although the size of this increase is larger in the sample excluding less-skilled metro areas, even in the sample excluding more-skilled metro areas we find the same reallocation of employment towards interactive occupations in metro areas.

8.3 Improvements in Transport and Communication Technologies

To provide direct evidence on the role played by improvements in transport and communication technologies, we combine data on employment by occupation, sector and county for 1880 and 1930 with information on the spatial diffusion of the telephone and road network in the opening decades of the twentieth century. We focus on this period because both the telephone and paved highways were virtually non-existent in 1880 and diffused rapidly from 1880-1930. This is also the period for which we observe the largest increase in the relative interactiveness of metro areas, and 1930 is the last year for which county identifiers are available in IPUMS, and hence the last year for which we can measure changes in interactiveness by county.²⁹

Our baseline specification regresses the change in interactiveness in each county from 1880-1930 ($\Delta\text{Interactive}_c$) on log telephones per capita (Phonepc_c) and highways per kilometer (Highwaypa_c) in the 1930s:

$$\Delta\text{Interactive}_c = \alpha_P \ln(\text{Phonepc}_c) + \alpha_H \text{Highwaypa}_c + X_c \alpha_X + u_c, \quad (29)$$

where Phonepc_c is residence telephones in 1935 divided by population in 1930; Highwaypa_c is the length of highways from the Gallup (1931) map in each county divided by county area; X_c are controls for other county characteristics; u_c is a stochastic error; since telephones and paved highways were both essentially non-existent in 1880, the values of these variables in the 1930s capture their growth from 1880-1930.

Telephones and highways are unlikely to be randomly assigned to counties. Therefore a concern is that changes in interactiveness and the diffusion of these technologies both could be influenced by omitted third factors that enter the error term u_c and hence induce a correlation between the diffusion of these technologies and the error term. In particular, we have already shown that more densely-populated locations experienced an increase in their relative interactiveness over time, and telephones and highways may have also diffused more rapidly to more densely-populated locations. For this reason, we include among our controls X_c each county's initial log population in 1880 and its log area.

To further address the concern that telephones and roads are non-randomly assigned, we develop instruments based on institutional features of the development of the telephone and highway network. We include these instruments alongside our controls in the following first-stage regressions:

$$\ln(\text{Phonepc}_c) = \beta_P Z_{Pc} + \beta_H Z_{Hc} + X_c \beta_X + \varepsilon_c, \quad (30)$$

with college degrees of greater than 25.025 percent in 2006. The year 1960 is omitted in Columns (7) and (8) because the IPUMS 1960 data do not contain the identifiers for individual MSAs.

²⁹While identifiers are available for some counties in the IPUMS data for 1940, these counties are a selected subsample of all counties.

$$\text{Highwaypa}_c = \gamma_P Z_{Pc} + \gamma_H Z_{Hc} + X_c \gamma_X + \omega_c, \quad (31)$$

where Z_{Pc} is our instrument for telephones (P is mnemonic for phones) and Z_{Hc} is our instrument for highways (H is mnemonic for highways); ε_c and ω_c are stochastic errors.

To develop an instrument for log telephones per capita, we exploit the network structure of telephone communication. Following Alexander Graham Bell’s successful filing for a patent in 1876, the Bell Telephone Company was incorporated in 1877, and the first telephone exchange was opened under license from Bell Telephone in New Haven, CT in 1878. As local telephone exchanges began to emerge in major U.S. cities, the American Telephone and Telegraph Company (AT&T) was formed in 1885 as a subsidiary of American Bell Telephone to build and operate a long distance telephone network. In these early years, there was considerable debate within American Bell Telephone about the strategic rationale for developing a long distance network and whether such a network would be profitable given that much of the initial demand for telephones appeared to be local (see for example John 2010).

By the end of 1885, the first long distance line was completed between New York and Philadelphia with an initial capacity of one telephone call, and it was not until 1892 that a long distance line to Chicago was finished again with an initial capacity of one call. Following Theodore Vail’s accession to the Presidency of AT&T in 1907, the company aggressively pursued the development of its long distance network, with the strategic goals of connecting the nation as a whole (e.g. Osbourne 1930) and pressing for nationwide monopoly powers under Vail’s slogan of “One System, One Policy, Universal Service.” Ultimately this goal was achieved in 1913 with the issuance of the Kingsbury Commitment, which established AT&T as a government-sponsored monopoly, in return for it divesting its interests in the manufacture of telephone and telegraph equipment and allowing independent telephone companies to connect with its long distance network. By 1915, the first transcontinental long distance line to San Francisco was completed.

As our instrument for county log telephones per capita, we use county proximity to AT&T’s long distance network (see Map A1 in the web appendix). We measure proximity using the log of the sum of the distances from each county’s centroid to the nearest primary and secondary outlets on this network, which captures the centrality of each county relative to the network. This instrument uses the fact that AT&T’s long distance network was developed with the strategic objective of connecting the nation rather than based on interactiveness in individual counties. Our identifying assumption is that conditional on our controls for initial population and area there is no direct effect of proximity to long distance outlets on county interactiveness other than through log telephones per capita. The locations of these long distance outlets have predictive power for log telephones per capita, because they facilitated the connection of local telephone companies to the long distance network, which increased the value of a telephone connection to local subscribers, and hence increased telephone diffusion. In this way, we exploit the network properties of the telephone, which it shares with for example distribution networks, as in Holmes (2011).

Our instrument for highways per kilometer uses the institutional development of the U.S. highways network. In 1880, paved roads were the exception and were concentrated in the immediate vicinity of central business districts.³⁰ Demand for road improvements grew following the production of the first American gasoline-powered automobile in Chicopee, Massachusetts in 1893 and the rapid growth in car registrations, which reached 8,000 in 1900, nearly 33,000 in 1903 and over 10 million by 1921 (U.S. Department of Transport 1976, Lewis 1997 and Swift 2011). The federal government's involvement in the road network dates back to the formation of the Office of Road Inquiry in 1893, which became the Office of Public Roads in 1905 and the Bureau of Public Roads in 1915. Federal government participation was stimulated in part by its responsibility for the postal service, which was a department of the federal government from 1792-1971. Thus the Federal Aid Road Act of 1916 provided federal funding for rural post roads on the condition that these roads were open to public at no charge and that states submitted plans, surveys and estimates for the approval of the Secretary of Agriculture.

The scale of federal government participation grew with the Federal Aid Highway Act of 1921, which provided 50-50 matching funds for state highway building. Each state was required to propose a system of roads for federal aid that did not exceed 7 percent of its highway mileage, and the Department of Agriculture was authorized to publish a map of the network on which federal aid would be spent by November 1923. As part of the planning process for this network, the Bureau of Public Roads commissioned General John J. Pershing to draw up a map of roads of military importance in the event of war. This "Pershing Map" identified 75,000 miles of road as strategically important for reasons of coastal and border defense (see Map A2 in the web appendix).³¹ More than 10,000 miles of Federal Aid Highways were laid down in 1922 and by 1929 more than 90 percent of the Federal Aid Highways (around 170,000 miles) had been improved.

We instrument the length of highways per kilometer from the Gallup (1931) map using the length of Pershing highways per kilometer within each county from the Pershing Map. Our identifying assumption is that conditional on our controls for initial population and area there is no direct effect of Pershing highways per kilometer on county interactiveness other than through actual highways per kilometer. Pershing highways per kilometer have predictive power for actual highways per kilometer, because these highways of military importance were incorporated into the final network of Federal Aid Highways in the Department of Agriculture's 1923 map.

In Column (1) of Table 7, we begin by running an OLS regression of the change in county interactiveness from 1880-1930 on county log telephones per capita in 1935, highways per kilometer in 1931 and our controls (equation (29)). We find a positive and statistically significant coefficient for telephones and a positive but statistically insignificant coefficient for highways.³² In Column (2), we report our instrumental

³⁰At the end of 1909, concrete accounted for only nine miles of state and county roads (Macdonald 1928).

³¹Consistent with these objectives, the Pershing Map excluded parts of the Deep South and Florida that were considered to be sufficiently swampy as to render foreign invasion impractical.

³²As discussed in the data section above, our telephones data for 1935 are for residence telephones. Separate data for business

variables estimates of equations (29)-(31). We find positive and statistically significant coefficients for both telephones and highways. Therefore both the diffusion of telephones induced by AT&T's long distance network and the development of highways for military reasons raise county interactiveness. The increase in the estimated coefficients on highways between the OLS and IV specifications is consistent with the view that conditional on our controls for population density highways are disproportionately assigned to locations with lower growth in interactiveness. This finding is in line with Duranton and Turner (2012)'s results for the later interstate highway system, in which conditional on their controls highways also appear to be disproportionately assigned to relatively less-developed locations. While these specifications control for population density through the inclusion of log population and log area, we also find a very similar pattern of results if we also include a (0,1) dummy for whether a county is located within a metro area.

In Columns (3)-(4) of Table 7, we report the first-stage regressions for phones and highways respectively, while Column (5) reports the reduced-form regression. We find that proximity to the AT&T long distance network and Pershing highways have predictive power for the endogenous variables, with first-stage F-statistics on the excluded exogenous variables of 38.40 and 26.35 in Columns (3) and (4) respectively. Consistent with this, we reject the null hypotheses of underidentification and weak identification in the Kleibergen-Paap test statistics reported in Column (2).

Taken together, these results provide evidence that the increase in interactiveness from 1880-1930 is indeed related to the diffusion of improvements in communication and transport technologies.

9 Conclusions

While there is a large literature on agglomeration, there is relatively little evidence on the tasks undertaken within agglomerations and how these have changed over time. We develop a new methodology for quantifying the tasks undertaken in urban and rural areas that uses 3,000 verbs from over 12,000 occupational descriptions over a period of more than a century. We use this methodology to construct a quantitative measure of the interactiveness of occupations that uses the frequency with which verbs from the occupational descriptions appear in thesaurus categories involving thought, communication and intersocial activity. We find an increase in the employment share of interactive occupations within sectors over time that is larger in metro areas than non-metro areas.

These findings are consistent with a simple model of trade in final goods and tasks between locations. The model emphasizes changes in the relative productivity of occupations and changes in the relative cost of trading goods and tasks that affect the relative demand for occupations. We provide evidence in support of this explanation and against alternative explanations. We show that the increase in the interactiveness of

and residence telephones are available for 1945 and we find a strong correlation between them. Regressing log residence telephones on log business telephones across counties in 1945, we find an estimated coefficient (standard error) of 0.8950 (0.0090) and a regression R-squared of 0.87.

employment in metro areas is pervasive, is not driven by any one occupation or any one sector, and occurs both between and within sectors. We find the same pattern when we exclude manufacturing or services, whether we focus on married or single workers, and whether we consider more or less-skilled metro areas. Consistent with an explanation based on relative demand, we find similar results using relative employment and relative wagebill shares.

We find that the increase in the relative interactiveness of employment in metro areas is particularly rapid in the early decades of the twentieth century and provide evidence for this period that changes in interactiveness are related to the diffusion of improvements in communication and transport technologies in the form of telephones and highways. Our findings highlight the increasing role that human interaction plays in agglomeration and the role of improvements in communication and transport technologies in shaping the task content of employment.

References

- [1] American Telephone and Telegraph Company (1935) *Residence Telephones by Counties*, New York: American Telephone and Telegraph Company
- [2] Autor, David H. and Michael J. Handel (2009) "Putting Tasks to the Test: Human Capital, Job Tasks and Wages," *NBER Working Paper*, 15116.
- [3] Autor, David H., Frank Levy and Richard J. Murnane (2003) "The Skill Content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics*, November, 1279-1333.
- [4] Bacolod, Marigee, Bernardo S. Blum and William C. Strange (2009a) "Skills in the City," *Journal of Urban Economics*, 65, 136-153.
- [5] Bacolod, Marigee, Bernardo S. Blum and William C. Strange (2009b) "Urban Interactions: Soft Skills Versus Specialization," *Journal of Economic Geography*, 9, 227-262.
- [6] Baldwin, Richard E. (2012) "Trade and Industrialisation after Globalisation's 2nd Unbundling: How Building and Joining a Supply Chain are Different and Why it Matters," *CEPR Discussion Paper*, 8768.
- [7] Baum-Snow, Nate (2007) "Did Highways Cause Suburbanization?" *Quarterly Journal of Economics*, 122(2), 775-805.
- [8] Becker, Sascha O., Karolina Ekholm and Marc-Andreas Muendler (2009) "Offshoring and the On-shore Composition of Tasks and Skills," University of California, San Diego.
- [9] Blinder, Alan (2009) "How Many U.S. Jobs Might Be Offshorable," *World Economics*, 10(2), 41-78.
- [10] Blinder, Alan and Alan Krueger (2012) "Alternative Measures of Offshorability: A Survey Approach," *Journal of Labor Economics*, forthcoming.
- [11] Combes, Pierre-Philippe Combes, Gilles Duranton, Laurent Gobillon, Diego Puga and Sbastien Roux (2012) "The Productivity Advantages of Large Cities: Distinguishing Agglomeration from Firm Selection," *Econometrica*, 80(6), 2543-2594.

- [12] Costa, Dora L. and Matthew E. Kahn (2000) "Power Couples: Changes In The Locational Choice Of The College Educated, 1940-1990," *Quarterly Journal of Economics*, 115(4), 1287-1315.
- [13] Costinot, Arnaud, Dave Donaldson and Ivana Komunjer (2012) "What Goods Do Countries Trade? A Quantitative Exploration of Ricardo's Ideas," *Review of Economic Studies*, 79(2), 581-608.
- [14] Davis, Donald R. and Jonathan I. Dingel (2012) "A Spatial Knowledge Economy," *NBER Working Paper*, 18188.
- [15] Davis, Morris A. and François Ortalo-Magne (2011) "Household Expenditures, Wages, Rents," *Review of Economic Dynamics*, 14(2), 248-261.
- [16] Duranton, Gilles and Diego Puga (2001) "Nursery Cities: Urban diversity, process innovation, and the life cycle of products", *American Economic Review*, 91(5), 1454-1477.
- [17] Duranton, Gilles and Diego Puga (2004) "Micro-foundations of Urban Agglomeration Economies," Chapter 48 in (eds.) J. V. Henderson and J.F. Thisse, *Handbook of Regional and Urban Economics*, Volume 4, Amsterdam: Elsevier North Holland.
- [18] Duranton, Gilles and Diego Puga (2005) "From Sectoral to Functional Urban Specialization," *Journal of Urban Economics*, 343-370.
- [19] Duranton, Gilles and Hubert 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.
- [20] Duranton, Gilles and Matthew Turner (2012) "Urban Growth and Transportation," *Review of Economic Studies*, forthcoming.
- [21] Eaton, Jonathan and Samuel S. Kortum (2002) "Technology, Geography, and Trade," *Econometrica*, 70(5) (2002), 1741-1779.
- [22] Ethier, Wilfred J. (1982) "Decreasing Costs in International Trade and Frank Grahams Argument for Protection," *Econometrica*, 50, 1243-1268.
- [23] Fischer, Claude S. (1992) *America Calling: A Social History of the Telephone to 1940*, Los Angeles: University of California Press.
- [24] Fujita, Masahisa and Takatoshi Tabuchi (1997) "Regional Growth in Postwar Japan," *Regional Science and Urban Economics*, 27, 643-670.
- [25] Gallup (1931) *Motor Trails Map of the United States*, Kansas City: Gallup Map Company.
- [26] Gaspar, Jess and Edward L. Glaeser (1998) "Information Technology and the Future of Cities," *Journal of Urban Economics*, 43(1), 136-156.
- [27] Gentzkow, Matthew, Jesse M. Shapiro and Michael Sinkinson (2012) "Competition and Ideological Diversity: Historical Evidence from U.S. Newspapers," Chicago Booth School of Business, mimeograph.
- [28] Glaeser, Edward L. and Albert Saiz (2003) "The Rise of the Skilled City," *NBER Working Papers*, 10191.
- [29] Glaeser, Edward L., A. M. Ponzetto and Kristina Toblo (2011) "Cities, Skills and Regional Change," *NBER Working Paper*, 16934.

- [30] Glaeser, Edward L. and Matthew G. Resseger (2009) "The Complementarity Between Cities and Skills," *NBER Working Paper*, 15103.
- [31] Gray, Rowena (2010) "Taking Technology to Task: The Skill Content of Technological Change in Early Twentieth Century U.S." University of Essex, mimeograph.
- [32] Grossman, Gene and Esteban Rossi-Hansberg (2008) "Trading Tasks: A Simple Theory of Offshoring," *American Economic Review*, 98(5), 1978-1997.
- [33] Grossman, Gene and Esteban Rossi-Hansberg (2012) "Task Trade between Similar Countries," *Econometrica*, 80(2), 593-629.
- [34] Helpman, Elhanan (1998) "The Size of Regions," In *Topics in Public Economics: Theoretical and Applied Analysis*, ed. David Pines, Efraim Sadka, and Itzhak Zilcha, Cambridge: Cambridge University Press.
- [35] Helsley, Robert W. and William C. Strange (2007) "Agglomeration, Opportunism and the Organization of Production," *Journal of Urban Economics*, 62, 55-75.
- [36] Henderson, J. Vernon (1974) "The Sizes and Types of Cities," *American Economic Review*, 64(4), 640-56.
- [37] Henderson, J. Vernon (2003) "Marshall's Scale Economies," *Journal of Urban Economics*, 53, 1-28.
- [38] Holmes, Thomas J. (2011) "The Diffusion of Wal-Mart and Economies of Density," *Econometrica*, 79(1), 253-302.
- [39] Hüllen, Werner (2003) *A History of Roget's Thesaurus: Origins, Development, and Design*, Oxford: Oxford University Press.
- [40] Jacobs, J. (1969) *The Economy of Cities*, Random House.
- [41] John, Richard R. (2010) *Network Nation: Inventing American Telecommunications*, Cambridge MA: Bellknap Press.
- [42] Leamer, Edward and Michael Storper (2001) "The Economic Geography of the Internet Age," *Journal of International Business Studies*, 32(4), 641-665.
- [43] Lewis, Tom (1997) *Divided Highways: Building the Interstate Highways, Transforming American Life*, New York, N.Y.: Viking.
- [44] Lin, Jeffrey (2011) "Technological Adaptation, Cities and New Work," *Review of Economics and Statistics*, 93(2), 554-574.
- [45] Macdonald, Thomas H. (1928) "The History and Development of Road Building in the United States," *Transactions of the American Society of Civil Engineers*, 92, 1181-1206.
- [46] Michaels, Guy (2008) "The Effect of Trade on the Demand for Skill - Evidence from the Interstate Highway System," *Review of Economics and Statistics*, 90(4), 683-701.
- [47] Michaels, Guy, Ferdinand Rauch and Stephen Redding (2012) "Urbanization and Structural Transformation," *Quarterly Journal of Economics*, 127, 535-586.

- [48] Michel, Jean-Baptiste, Yuan Kui Shen, Aviva Presser Aiden, Adrian Veres, Matthew K. Gray, The Google Books Team, Joseph P. Pickett, Dale Hoiberg, Dan Clancy, Peter Norvig, Jon Orwant, Steven Pinker, Martin A. Nowak, and Erez Lieberman Aiden (2011) “Quantitative Analysis of Culture Using Millions of Digitized Books,” *Science*, 331, 176-182.
- [49] Moretti, Enrico (2004) “Human Capital Externalities in Cities,” Chapter 51 in (eds) V. Henderson and J.F. Thisse, *Handbook of Urban and Regional Economics*, Amsterdam: North Holland-Elsevier.
- [50] Ngai, Rachel, and Chris Pissarides (2007) “Structural Change in a Multi-sector Model of Growth,” *American Economic Review*, 97(1), 429-443.
- [51] Osbourne, H. S. (1930) “A General Switching Plan for Telephone Toll Service,” *Bell System Technical Journal*, 9, 429-447.
- [52] Ota, Mitsuru and Masahisa Fujita (1993) “Communication Technologies and Spatial Organization of Multi-unit Firms in Metropolitan Areas,” *Regional Science and Urban Economics*, 23, 695-729.
- [53] Ottaviano, Gianmarco I. P., Giovanni Peri and Greg C. Wright (2010) “Immigration, Offshoring and American Jobs,” *NBER Working Paper*, 16439.
- [54] Pool, Ithiel de Sola (1977) *The Social Impact of the Telephone*, Cambridge MA: MIT Press.
- [55] Roget, Peter Mark (1911) *Thesaurus of English Words and Phrases*, Computer Readable Version of 1911 Edition, Evergreen Review, Inc.
- [56] Rosenthal, Stuart S. and William C. Strange (2004) “Evidence on the Nature and Sources of Agglomeration Economics,” Chapter 49 in (eds.) J. V. Henderson and J.F. Thisse, *Handbook of Regional and Urban Economics*, Volume 4, Amsterdam: Elsevier North Holland.
- [57] Rossi-Hansberg, Esteban, Pierre-Daniel Sarte and Raymond Owens III (2009) “Firm Fragmentation and Urban Patterns,” *International Economic Review*, 50(1), 143-186.
- [58] Ruggles, Steven, J. Trent Alexander, Katie Genadek, Ronald Goeken, Matthew B. Schroeder, and Matthew Sobek (2010) *Integrated Public Use Microdata Series*, Version 5.0 [Machine-readable database]. Minneapolis: University of Minnesota.
- [59] Standage, Tom (2007) *The Victorian Internet: The Remarkable Story of the Telegraph and the Nineteenth Century’s On-line Pioneers*, New York: Walker & Company.
- [60] Swift, Earl (2011) *The Big Roads: The Untold Story of the Engineers, Visionaries, and Trailblazers Who Created the American Superhighways*, Mariner Books.
- [61] U. S. Department of Labor (1939) *Dictionary of Occupational Titles*, Washington DC: U.S. Department of Labor.
- [62] U. S. Department of Labor (1991) *Dictionary of Occupational Titles*, Washington DC: U.S. Department of Labor.
- [63] U.S. Department of Transportation (1976) *America’s Highways 1776-1976: A History of the Federal Aid Program*, Washington DC: U.S. Government Printing Office.
- [64] Yi, Kei-Mu and Jing Zhang (2010) “Structural Change in an Open Economy,” University of Minnesota, mimeograph.

Table 1: Metro Area Specialization for Aggregate Occupations and Sectors

Panel A						
Two-digit occupation	Coefficient	Standard		Coefficient	Standard	
	1880	Error 1880	Rank 1880	2000	Error 2000	Rank 2000
Clerical and Kindred	0.15	0.08	1	0.04	0.01	4
Craftsmen	0.09	0.06	2	-0.01	0.01	6
Operatives	0.06	0.07	3	-0.05	0.01	7
Sales workers	0.01	0.07	4	0.05	0.01	2
Service Workers	0.00	0.08	5	0.00	0.01	5
Managers, Officials, and Proprietors	-0.03	0.08	6	0.05	0.01	3
Professional, Technical	-0.07	0.08	7	0.07	0.01	1
Laborers	-0.20	0.18	8	-0.15	0.07	8

Panel B						
Two-digit sector	Coefficient	Standard		Coefficient	Standard	
	1880	Error 1880	Rank 1880	2000	Error 2000	Rank 2000
Entertainment and Recreation Services	0.29	0.08	1	0.04	0.01	4
Wholesale and Retail Trade	0.13	0.05	2	0.01	0.01	6
Finance, Insurance, and Real Estate	0.13	0.06	3	0.06	0.01	2
Manufacturing	0.06	0.05	4	-0.01	0.01	10
Personal Services	0.01	0.06	5	0.03	0.01	5
Transportation, Communication, Other Utilities	0.01	0.04	6	0.05	0.01	3
Public Administration	-0.03	0.07	7	0.01	0.01	7
Professional and Related Services	-0.03	0.06	8	0.00	0.01	9
Business and Repair Services	-0.12	0.08	9	0.08	0.01	1
Construction	-0.14	0.08	10	0.00	0.01	8
Mining	-0.31	0.05	11	-0.27	0.03	11

Notes: Coefficients estimated from a regression of the share of employment in metro areas in an occupation-sector-year on occupation-year and sector-year fixed effects (regression (20) in the paper). Occupation-year and sector-year fixed effects are each normalized to sum to zero. A separate regression is estimated for each year. Standard errors are clustered by occupation. Occupations and sectors sorted by the rank of their estimated coefficients for 1880.

Table 2: Verbs Most and Least Strongly Correlated with Metro Area Employment Shares

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	Stitch	Distribute	Bill	Schedule	Document	Determine
3	Interfere	Telephone	Record	Take	File	Advise	Analyze
4	Hand	Sew	Notice	Compile	Record	Concern	Factor
5	Ravel	Hand	Telephone	Distribute	Distribute	Report	Review
6	Sew	Assist	Bill	Pay	Compile	Schedule	Confer
7	Braid	Visit	Envelope	Letter	Notice	Develop	Advise
8	Visit	Describe	Document	Notice	Identify	Analyze	Report
9	Receive	Number	Learn	Record	Send	Determine	Concern
10	Sack	Stamp	Number	Send	Notify	Notify	Plan
Panel B: Verbs Least Strongly Correlated with Metro Area Employment Shares							
Rank	1880	1900	1920	1940	1960	1980	2000
1821	Conduct	Abstract	Counsel	Delegate	Accord	Power	Restrain
1822	Teach	Tread	Discuss	Enlist	Feed	Pour	Cut
1823	Channel	Pinch	Hear	Labor	Escape	Erect	Power
1824	Sound	Assign	Assign	Tread	Hook	Clean	Massage
1825	Rule	Settle	Teach	Assign	Traverse	Massage	Remove
1826	Matter	Matter	Matter	Approve	Tread	Pump	Feed
1827	Drill	Tunnel	Consolidate	Extract	Loosen	Cut	Clean
1828	Tread	Sound	Rule	Tunnel	Range	Feed	Pump
1829	Tunnel	Rule	Tunnel	Malt	Activate	Move	Move
1830	Pinch	Sole	Sound	Establish	Turn	Turn	Turn

Notes: Coefficients estimated from a regression of the share of occupation-sector employment in metro areas on the frequency with which a verb is used for an occupation and verb-sector-year fixed effects (regression (21) in the paper). A separate regression is estimated for each verb and verbs are sorted by their estimated coefficients normalized by the standard deviation for the verb frequency. Verbs are from the time-invariant occupational descriptions from the 1991 Dictionary of Occupations (DOTs).

Table 3: Correlation of Thesaurus Sections with Metro Area Employment Shares

Thesaurus Class	Thesaurus Section	Rank	Rank	Rank
		Section 1880	Section 2000	1880 - 2000
C1, Abstract relations	SECTION I. EXISTENCE	9	10	-1
C1, Abstract relations	SECTION II. RELATION	17	6	11
<i>C1, Abstract relations</i>	<i>SECTION III. QUANTITY</i>	<i>1</i>	<i>32</i>	<i>-31</i>
C1, Abstract relations	SECTION IV. ORDER	30	11	19
C1, Abstract relations	SECTION V. NUMBER	23	14	9
<i>C1, Abstract relations</i>	<i>SECTION VI. TIME</i>	<i>2</i>	<i>20</i>	<i>-18</i>
C1, Abstract relations	SECTION VII. CHANGE	36	7	29
C1, Abstract relations	SECTION VIII. CAUSATION	28	21	7
C2, Space	SECTION I. SPACE I N GENERAL	8	33	-25
<i>C2, Space</i>	<i>SECTION II. DIMENSIONS</i>	<i>4</i>	<i>36</i>	<i>-32</i>
C2, Space	SECTION IV. MOTION	25	24	1
<i>C3, Matter</i>	<i>SECTION I. MATTER IN GENERAL</i>	<i>3</i>	<i>31</i>	<i>-28</i>
<i>C3, Matter</i>	<i>SECTION II. INORGANIC MATTER</i>	<i>5</i>	<i>35</i>	<i>-30</i>
C3, Matter	SECTION III. ORGANIC MATTER	14	37	-23
C4, Intellect	SECTION I. NATURE OF IDEAS COMMUNICATED	19	15	4
C4, Intellect	SECTION I. OPERATIONS OF INTELLECT IN GENERAL	15	23	-8
C4, Intellect	SECTION II. MODES OF COMMUNICATION	18	9	9
<i>C4, Intellect</i>	<i>SECTION II. PRECURSORY CONDITIONS & OPERATIONS</i>	<i>34</i>	<i>5</i>	<i>29</i>
C4, Intellect	SECTION III. MATERIALS FOR REASONING	33	8	25
<i>C4, Intellect</i>	<i>SECTION III. MEANS OF COMMUNICATING IDEAS</i>	<i>11</i>	<i>2</i>	<i>9</i>
C4, Intellect	SECTION IV. REASONING PROCESSES	38	22	16
<i>C4, Intellect</i>	<i>SECTION V. RESULTS OF REASONING</i>	<i>7</i>	<i>1</i>	<i>6</i>
C4, Intellect	SECTION VI. EXTENSION OF THOUGHT	24	18	6
C4, Intellect	SECTION VII. CREATIVE THOUGHT	35	17	18
C5, Will	SECTION I. GENERAL INTERSOCIAL VOLITION	20	25	-5
C5, Will	SECTION I. VOLITION IN GENERAL	29	28	1
C5, Will	SECTION II. Prospective Volition 1	21	38	-17
C5, Will	SECTION II. SPECIAL INTERSOCIAL VOLITION	26	13	13
C5, Will	SECTION III. CONDITIONAL INTERSOCIAL VOLITION	27	12	15
<i>C5, Will</i>	<i>SECTION III. VOLUNTARY ACTION</i>	<i>32</i>	<i>4</i>	<i>28</i>
C5, Will	SECTION IV. ANTAGONISM	10	27	-17
C5, Will	SECTION IV. POSSESSIVE RELATIONS	16	16	0
C5, Will	SECTION V. RESULTS OF VOLUNTARY ACTION	31	26	5
C6, Emotion, Religion, Morality	SECTION I. AFFECTIONS IN GENERAL	6	34	-28
C6, Emotion, Religion, Morality	SECTION II. PERSONAL AFFECTIONS	22	30	-8
C6, Emotion, Religion, Morality	SECTION III. SYMPATHETIC AFFECTIONS	12	29	-17
<i>C6, Emotion, Religion, Morality</i>	<i>SECTION IV. MORAL AFFECTIONS</i>	<i>37</i>	<i>3</i>	<i>34</i>
C6, Emotion, Religion, Morality	SECTION V. RELIGIOUS AFFECTIONS	13	19	-6

Notes: Coefficients estimated from a regression of the share of occupation-sector employment in metro areas on the frequency with which the verbs used for an occupation are classified within thesaurus sections and thesaurus-section-sector-year fixed effects (regression (23) in the paper). A separate regression is estimated for each thesaurus section. Verbs are from the time-invariant occupational descriptions from the 1991 Dictionary of Occupations (DOTs). Thesaurus sections ranked in terms of their estimated coefficient normalized by the standard deviation for the thesaurus section frequency, where the largest value is assigned a rank of one. Top-five thesaurus sections in 1880 highlighted in bold and italics. Top-five thesaurus sections in 2000 highlighted in bold.

Table 4: Most and Least Interactive Occupations

Panel A: Top Ten Interactive Occupations

Economists
Nurses, professional
Pharmacists
Clergymen
Religious workers
Accountants and auditors
Postmasters
Buyers and dept heads, store
Aeronautical-Engineers
Statisticians and actuaries

Panel B: Bottom Ten Interactive Occupations

Brickmasons,stonemasons, tile setters
Attendants, auto service and parking
Painters, except construction
Plumbers and pipe fitters
Upholsterers
Asbestos and insulation workers
Welders and flame cutters
Blasters and powdermen
Dressmakers and seamstresses
Roofers and slaters

Notes: The table reports the ten occupations with the lowest and highest interactiveness, as measured by the frequency of verb use in Classes IV and V of Roget's Thesaurus. Verbs are from the time-invariant occupational descriptions from the 1991 Dictionary of Occupations (DOTs).

Table 5: Metro Employment and Wagebill Shares and Interactiveness

Panel A: Between sectors							
Measure	1880	1900	1920	1940	1960	1980	2000
Interactiveness	-0.130 (0.267)	-0.132 (0.239)	0.258 (0.419)	0.556 (0.405)	0.728*** (0.267)	0.901*** (0.200)	0.814*** (0.182)
Thought	-0.649** (0.268)	-1.304*** (0.261)	-1.805*** (0.363)	-0.608 (0.493)	0.179 (0.313)	0.780*** (0.280)	1.202*** (0.237)
Communication	-0.412*** (0.153)	-0.568*** (0.153)	-0.641*** (0.188)	-0.212 (0.272)	0.219 (0.199)	0.359* (0.210)	0.530** (0.233)
Intersocial	-0.292** (0.144)	-0.473*** (0.136)	-0.548*** (0.169)	-0.0624 (0.203)	0.126 (0.133)	0.280** (0.124)	0.342*** (0.109)
Interactiveness				0.557 (0.366)	0.557* (0.283)	0.814*** (0.215)	0.733*** (0.201)
Panel B: Within sectors							
Measure	1880	1900	1920	1940	1960	1980	2000
Interactiveness	-0.410*** (0.120)	-0.261** (0.119)	-0.104 (0.119)	-0.0360 (0.119)	0.190*** (0.0644)	0.274*** (0.0514)	0.317*** (0.0402)
Thought	-0.340** (0.134)	-0.411*** (0.132)	-0.299*** (0.0933)	-0.145 (0.0948)	0.153*** (0.0489)	0.227*** (0.0374)	0.246*** (0.0394)
Communication	-0.0408 (0.144)	-0.0423 (0.118)	0.0249 (0.0977)	0.118 (0.0789)	0.183*** (0.0360)	0.168*** (0.0323)	0.140*** (0.0384)
Intersocial	-0.0300 (0.130)	-0.0809 (0.0780)	-0.0172 (0.0582)	0.0197 (0.0492)	0.105*** (0.0320)	0.0652* (0.0342)	0.0460 (0.0476)
Inteactiveness				0.043 (0.087)	0.207*** (0.053)	0.281*** (0.043)	0.311*** (0.037)
Sector-year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Each cell of each panel of the table corresponds to a separate regression. The left-hand side in the first four rows of each panel is the share of employment in metro areas; the left-hand side in the fifth row of each panel is the share of the wagebill in metro areas; the wagebill data are only available from 1940 onwards; Interactiveness is our baseline measure using the 1991 DOTs and Classes IV-V of the thesaurus; Thought uses Class IV (Division 1) of the thesaurus; Communication uses Class IV (Division 2) of the thesaurus; Intersocial uses Class V (Division 2) of the thesaurus. In Panel A, the sample is a cross-section of three-digit sectors for each year, and the standard errors are heteroscedasticity robust. In Panel B, the sample is a panel of sectors and occupations for each year; sector-year fixed effects are included; and the standard errors are clustered on occupation. See Section 8.2 for further details on the estimated equation. * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 6: Metro Area Employment Shares and Interactiveness, Within-sector and Within-Occupation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Metro	Metro	Metro	Metro	Metro	Metro	Metro	Metro
	Employment	Employment	Employment	Employment	Employment	Employment	Employment	Employment
	Share	Share	Share	Share	Share	Share	Share	Share
Interactiveness x 1900	0.00207 (0.141)	0.104 (0.162)	0.0493 (0.119)	0.292 (0.195)	0.108 (0.178)	0.124 (0.177)	-0.0307 (0.0763)	0.124 (0.146)
Interactiveness x 1920	0.186 (0.202)	0.187 (0.218)	0.132 (0.176)	0.525** (0.223)	0.254 (0.250)	0.272 (0.203)	-0.00455 (0.129)	0.321 (0.206)
Interactiveness x 1940	0.399 (0.243)	0.321 (0.235)	0.287 (0.205)	0.455* (0.254)	0.334 (0.236)	0.324 (0.233)	0.0379 (0.117)	0.369* (0.221)
Interactiveness x 1960	0.573** (0.231)	0.485*** (0.185)	0.316** (0.158)	0.548** (0.261)	0.284 (0.228)	0.410* (0.227)		
Interactiveness x 1980	0.677*** (0.244)	0.560*** (0.174)	0.489*** (0.152)	0.627** (0.258)	0.424* (0.250)	0.515** (0.240)	0.233*** (0.0672)	0.595*** (0.204)
Interactiveness x 2000	0.672*** (0.253)	0.596*** (0.174)	0.609*** (0.141)	0.788*** (0.233)	0.552** (0.276)	0.681*** (0.221)	0.261*** (0.0670)	0.823*** (0.167)
Observations	56,760	56,760	49,108	41,442	25,105	30,593	35,662	44,128
Year fixed effects	yes							
Sector fixed effects	yes							
Occupation fixed effects	yes	yes	yes	yes	yes	yes	yes	yes
Sector-Year fixed effects		yes	yes	yes	yes	yes	yes	yes
Married only sample			yes					
Single only sample				yes				
Manufacturing only sample					yes			
Services only sample						yes		
Excluding more skilled metro areas							yes	
Excluding less skilled metro areas								yes

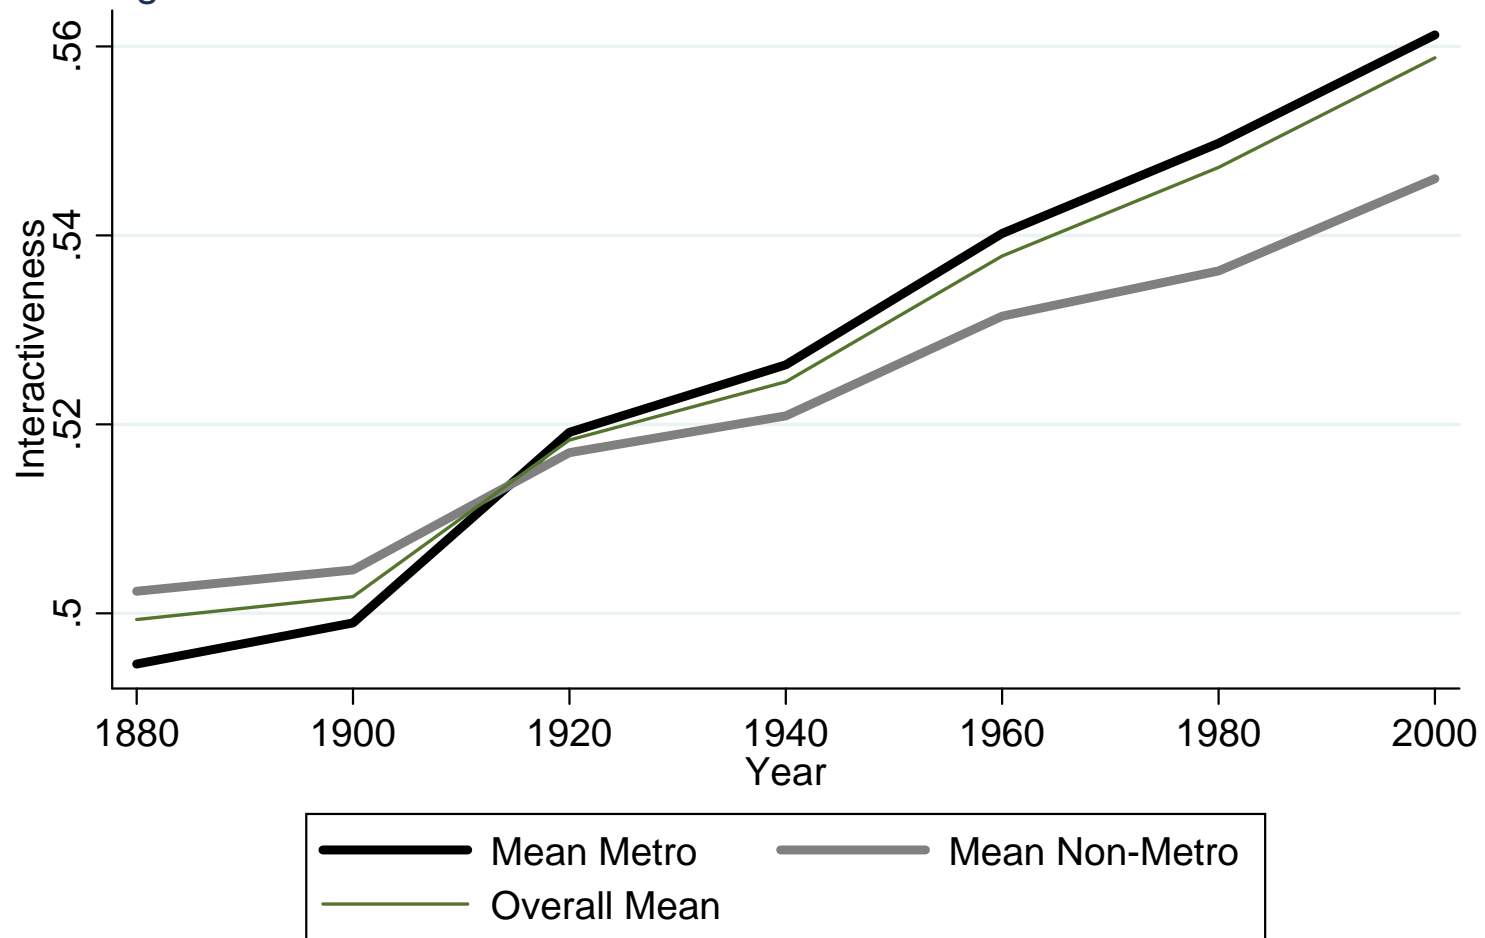
Notes: Sample is a panel of occupation-sector-year observations for twenty-year intervals from 1880-2000; 1880 is the excluded year from the interactions; interactiveness is our baseline measure using the 1991 DOTs and Classes IV-V of the thesaurus. Married only sample includes married workers only. Single only sample excludes married workers. Manufacturing only sample includes workers in manufacturing only. Services only sample includes workers in services only. More and less-skilled metro areas are defined as in Glaeser and Resseger (2009) based on whether the share of adults with a college degree in a Metropolitan Statistical Area (MSA) is greater than or less than 25.025 percent in 2006. The year 1960 is omitted in Columns (7) and (8) because the IPUMS 1960 data do not contain the identifiers for individual MSAs. See Section 8.2 for further details on the estimated equation. Standard errors are clustered on occupation; * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 7: Interactiveness and Improvements in Communication and Transport Technologies

	(1)	(2)	(3)	(4)	(5)
	Change in Interactiveness 1880-1930	Change in Interactiveness 1880-1930	Log phones per capita 1935	Highways per km 1931	Change in Interactiveness 1880-1930
Highways per km	0.007 (0.004)	0.086*** (0.028)			
Log phones per capita	0.022*** (0.002)	0.083*** (0.019)			
Log area	0.007*** (0.001)	0.010*** (0.001)	-0.013** (0.005)	-0.030*** (0.003)	0.007*** (0.001)
Log population 1880	0.004*** (0.001)	0.002* (0.001)	0.006* (0.003)	0.016*** (0.002)	0.004*** (0.001)
Pershing highways per km			-0.113** (0.055)	0.274*** (0.032)	0.015*** (0.005)
Log remoteness from long distance outlet			-0.063*** (0.009)	0.008** (0.004)	-0.005*** (0.001)
Observations	2467	2467	2467	2509	2509
R-squared	0.12	0.12	0.02	0.19	0.09
Estimation	OLS	2SLS	OLS	OLS	OLS
Specification	Second-stage	Second-stage	First-stage	First-stage	Reduced-form
F-statistic instruments			26.35	38.40	14.05
Underidentification test (Kleibergen-Paap LM statistic)		35.63			
Weak identification test (Kleibergen-Papp F-statistic)		18.61			

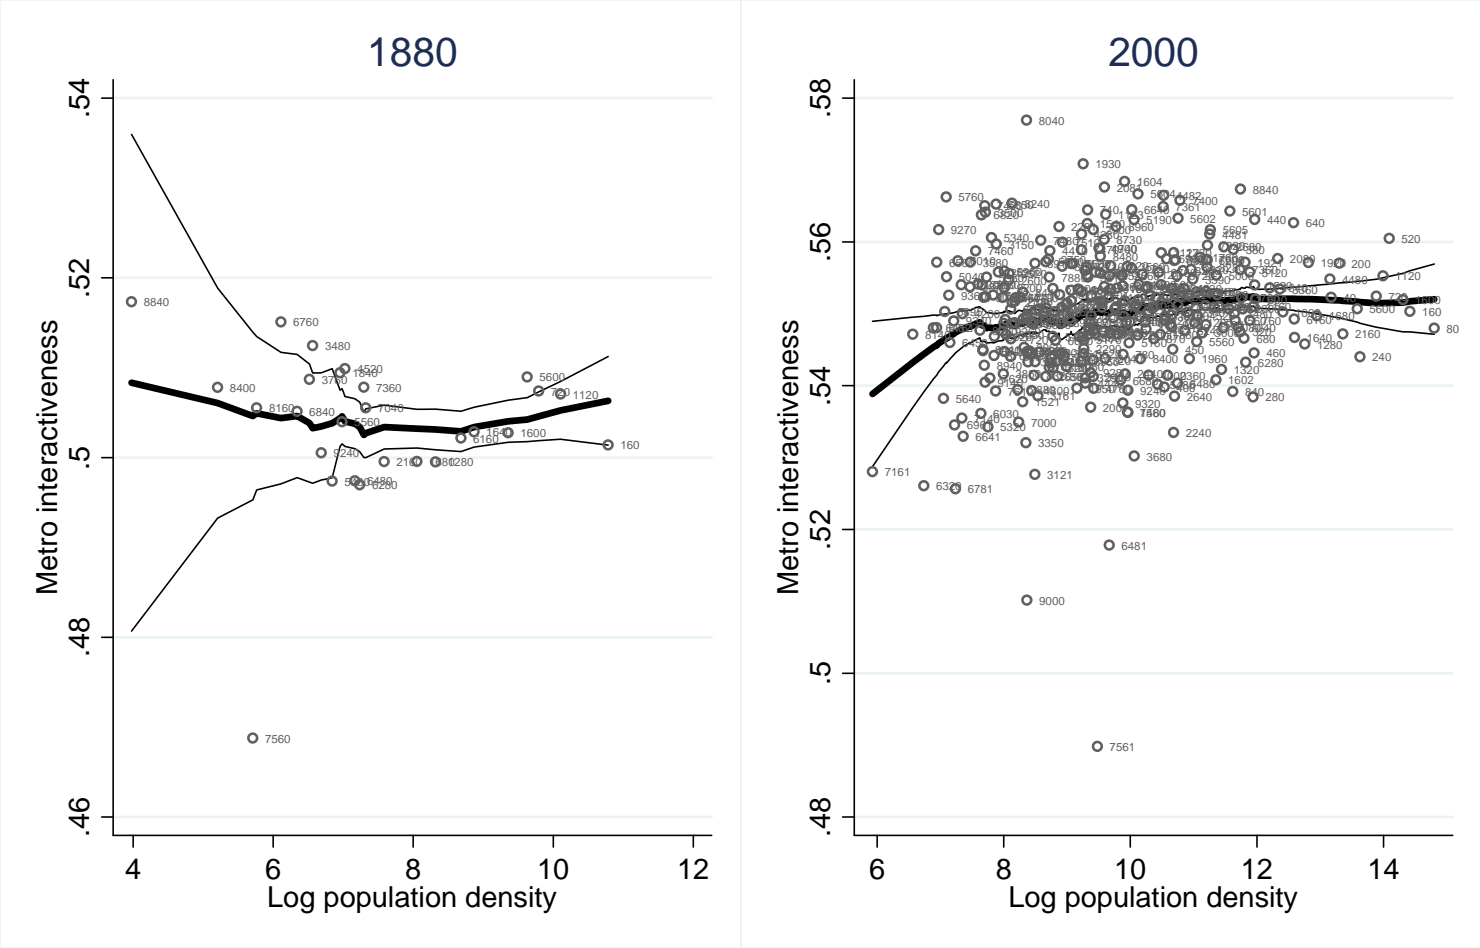
Note: Sample is a cross-section of counties from 1880-1930; Interactiveness is our baseline measure using the 1991 DOTs and Classes IV-V of the thesaurus; Highways per km is length of highways within a county in the Gallup 1931 map divided by county area; Log phones per capita is log number of residence telephones in 1935 divided by population in 1930; Log area is log county area; Log Population 1880 is log county population in 1880; Pershing highways per km is the length of highways proposed for military reasons within a county in the Pershing 1922 map divided by county area; Log remoteness from long distance outlet is the log of the sum of the distances to primary and secondary outlets on the AT&T long distance telephone network. Robust standard errors in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%.

Figure 1: Mean Interactiveness in Metro and Non-Metro Areas over Time



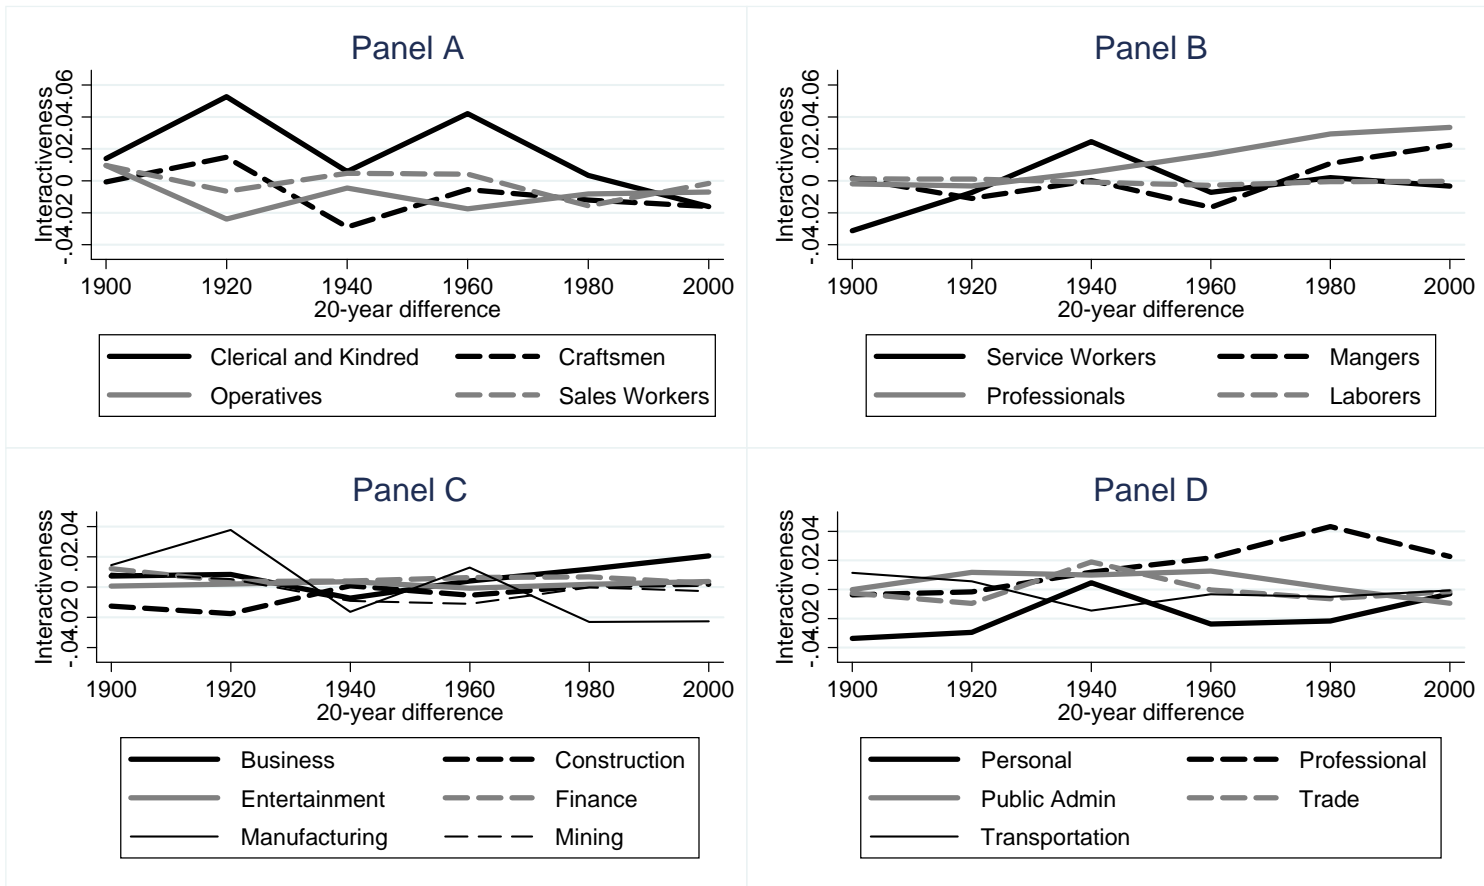
Notes: Mean interactiveness computed using time-invariant occupational descriptions from the 1991 DOTs.

Figure 2: Mean Interactiveness Across Metro Areas in 1880 and 2000



Notes: Mean interactiveness computed using time-invariant occupational descriptions from the 1991 DOTs. Thick solid line is fitted values from locally-weighted linear least squares regression. Thin solid lines are 95 percent point confidence intervals.

Figure 3: Decomposition of Difference in Change in Interactiveness Between Metro and Non-Metro Areas



Notes: Decomposition of the difference between metro and non-metro areas in the change in mean interactiveness over twenty-year intervals (equation (28) in the paper) into the contributions of two-digit occupations and sectors. Mean interactiveness based on time-invariant occupational descriptions from the 1991 DOTs.