Structural Transformation and the Rise of Information Technology

Giovanni Gallipoli, Christos A. Makridis

Appendix A. Online Appendix (Not for Print)

Appendix A.1. Supplement to Data Measurement

Appendix A.1.1. Index Measures of IT Intensity

We use a combination of input data from O*NET on the underlying tasks, knowledge, and skills workers use at a six-digit level of occupational heterogeneity. While we detail the different indices below, we have also experimented with different clustering algorithms besides our baseline approach of averaging across each of these indices.

Computers and Electronics (knowledge): Knowledge of circuit boards, processors, chips, electronic equipment, and computer hardware and software, including applications and programming Question: What level of COMPUTERS AND ELECTRONICS is needed to perform your current job? 1 (lowest level): Operate a VCR to watch a pre -recorded training tape 4 (intermediate): Use a word processor 6 (high): Create a program to scan computer disks for viruses.

Interacting with computers (work activity): Using computers and computer systems (including hardware and software) to program, write software, set up functions, enter data, or process information. Question: What level of WORKING WITH COMPUTERS is needed to perform your current job? 1 (lowest level): Enter employee information into a computer database 4 (intermediate): Write software for keeping track of parts in inventory 6 (high): Set up a new computer system for a large multinational company.

Programming (skills): Writing computer programs for various purposes Question: What level of PROGRAMMING is needed to perform your current job? 1 (lowest level): Write a program in BASIC to sort objects in a database 4 (intermediate): Write a statistical analysis program to analyze demographic data 6 (high): Write expert system programs to analyze ground radar geological data for probable existence of mineral deposits.

System (skill): Evaluation Identifying measures or indicators of system performance and the actions needed to improve or correct performance, relative to the goals of the system Question: What level of SYSTEMS EVALUATION is needed to perform your current job? 1 (lowest level): Determine why a coworker has been overly optimistic about how long it would take to complete a task 4 (intermediate): Identify the major reasons why a client might be unhappy with a product 6 (high): Evaluate the long-term performance problem of a new computer system.

Quality control analysis (skill): Conducting tests and inspections of products, services, or processes to evaluate quality or performance Question: What level of QUALITY CONTROL ANALYSIS is needed to perform your current job? 1 (lowest level): Inspect a draft memorandum for clerical errors 4 (intermediate): Measure new part requisitions for tolerance to specifications 6 (high): Develop procedures to test a prototype of a new computer system.

Operations analysis (skill): Analyzing needs and product requirements to create a design Question: What level of OPERATIONS ANALYSIS is needed to perform your current job? 1 (lowest level): Select a photocopy machine for an office 4 (intermediate): Suggest changes in software to make a system more user friendly 6 (high): Identify the control system needed for a new process production plant.

Updating and using Relevant knowledge (work activity): Keeping up-to-date technically and applying new knowledge to your job Question: What level of UPDATING AND USING RELEVANT KNOWLEDGE is needed to perform your current job? 1 (lowest level): Keep up with price changes in a small retail store 4 (intermediate): Keep current on changes in maintenance procedures for repairing sports cars 6 (high): Learn information related to a complex and rapidly changing technology.

Technology design (skill): Generating or adapting equipment and technology to serve user needs. Question: What level of TECHNOLOGY DESIGN is needed to perform your current job? 1 (lowest level): Adjust exercise equipment for use by a customer 4 (intermediate): Redesign the handle on a hand tool for easier gripping 6 (high): Create new technology for producing industrial diamonds.

Analyzing Data or Information (work activity): Identifying the underlying principles, reasons, or facts of information by breaking down information or data into separate parts. Question: What level of ANALYZING DATA OR INFORMATION is needed to perform your current job? 1 (lowest level): Determine the location of a lost order 4: (intermediate): Determine the interest cost to finance a new building 6: (high): Analyze the cost of medical care services for all hospitals in the country.

Processing Information (work activity): Compiling, coding, categorizing, calculating, tabulating, auditing, or verifying information or data. Question: 1 (lowest level): Tabulate the costs of parcel deliveries 4: (intermediate): Calculate the adjustments for insurance claims 6: (high): Compile data for a complex scientific report.

Engineering and technology (knowledge): Knowledge of the practical application of engineering science and technology. This includes applying principles, techniques, procedures, and equipment to the design and production of various goods and services. Question: What level of knowledge of ENGINEERING AND TECHNOLOGY is needed to perform your current job? 1 (lowest level): Install a door lock 4: (intermediate): Design a more stable grocery cart 6: (high): Plan for the impact of weather in designing a bridge.

Management of Material Resources: Obtaining and seeing to the appropriate use of equipment, facilities, and materials needed to do certain work. Question: What level of MANAGEMENT OF MATERIAL RESOURCES is needed to perform your current job? 1 (lowest level): Rent a meeting room for a management meeting 4: (intermediate): Evaluate an annual uniform service contract for delivery drivers 6: (high): Determine the computer system needs of a large corporation and monitor use of the equipment.

We take the sum of these constructed indices and produce a standardized z-score with a mean of zero and standard deviation of one. We classify jobs as high (low) IT if the z-score is above (below) zero.

Appendix A.1.2. Background Comparisons of Indices

Here we discuss comparisons of our IT index to alternative measures. Namely, we examine how results change when we use a close analogue that classifies jobs as high IT if they are above the median z-score of IT intensity. Figure A.1 plots the earnings and employment premia across the manufacturing and services sectors under both the baseline (mean) and alternative (median) score rule. These premia are defined as the logged earnings and employment difference between high and low IT jobs, respectively. Panels A and C illustrate that the earnings premia are very similar in both levels and trends—the correlation between the two is 0.83. For example, where the definition based on the mean implies an initial 1970 logged earnings difference around 0.50 in manufacturing, the definition based on the median implies an earnings difference of roughly 0.39.

Panels B and D show that the employment premia have similar trends. The differences in levels should not be surprising in light of the fact that the definition based on the median contains a larger number of observations since it splits occupations evenly into high and low IT, whereas the definition based on the mean includes roughly 7% fewer observations in the high IT category based on the distribution of IT intensity across jobs.

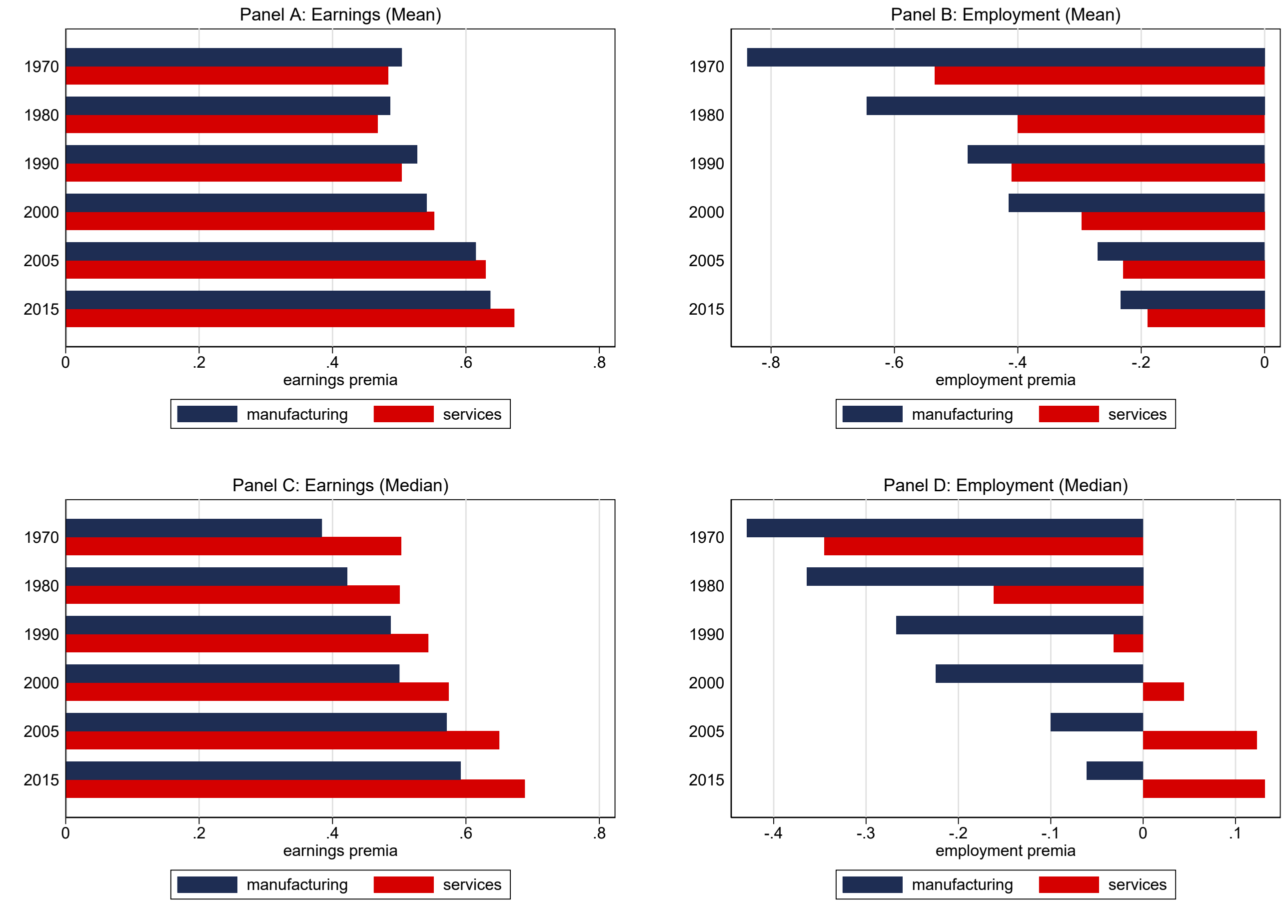

Figure A.1: Comparison of Mean and Median Based Information Technology Labels

Notes.–Sources: Census Bureau and O*NET, 1970-2015. Panels A and C plot the logged (average) annual earnings among IT workers net of logged annual earnings among non-IT workers (weighted by Census sample weights to produce the average) in the manufacturing and services sector using the two different definitions of IT jobs. Earnings is deflated using the 2010 personal consumption expenditures index and the sample is restricted to workers earning over $5,000/year, $2/hour, and working 500 hours/year. Panels B and D plot the logged total number of IT workers net of the logged total number of non-IT workers using the two different definitions of IT jobs. In the "mean" based definition, IT workers are classified as those in an occupation with an IT intensity score above the mean (z-score of zero) in the distribution of five-digit occupations, whereas, in the "median" based definition, IT workers are classified as those in an occupation with an IT intensity score above the median z-score in the distribution of five-digit occupations.

We also examine an alternative measure of IT intensity among workers, proposed by the Census (Beckhusen, 2016). Initially introduced in 1970, this classification designated only three occupational groups as IT jobs, and by 2010 it included 12 occupational groups. Figure A.2 plots both the employment count of IT workers and the employment share of IT jobs in the overall economy since 1970. This measure captures jobs that overwhelmingly use computers (e.g., database administrators and computer programmers), while omitting a wide array of occupations that regularly interact with other types of information technology. We, therefore, conclude that our approach is preferable as the O*NET provides a systematic way of surveying all occupations and their relative interaction and dependence on IT related skills, task, and knowledge.

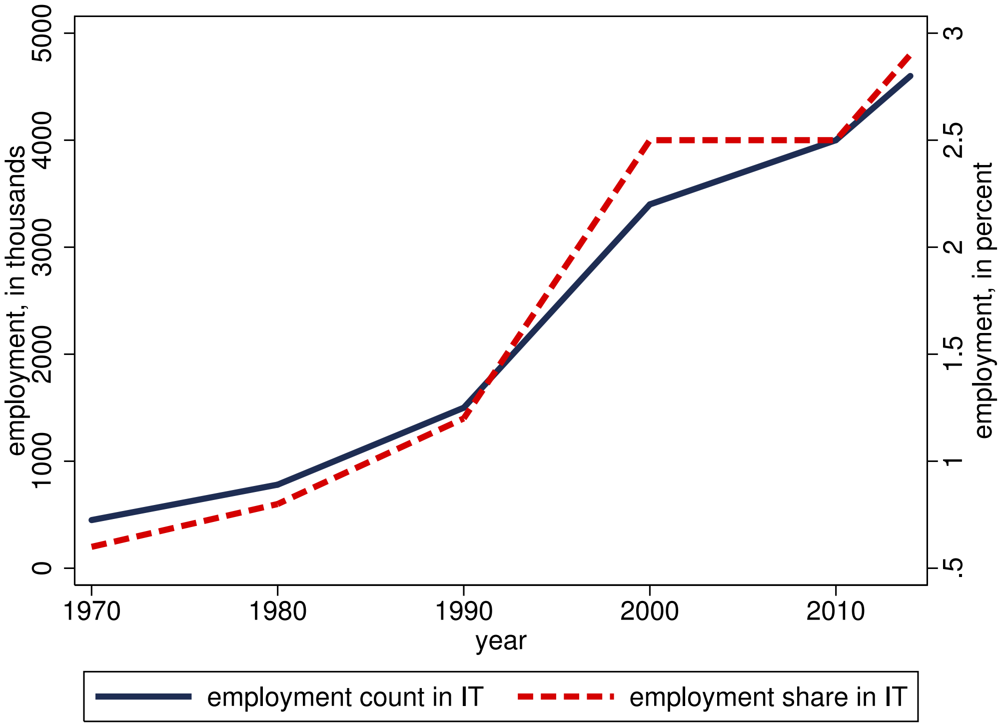

Figure A.2: Information Technology Workers, 1970-2014
Notes.–Sources: See Beckhusen (2016) for calculations. Occupation codes were harmonized to the 2010 census classification.

Appendix A.2. Supplement to Descriptive Statistics

The main text presents regressions of various outcome measures (e.g., logged hourly wage) on IT intensity using Census micro-data. While these coefficients characterize the conditional mean, we now explore the distribution of these outcomes in occupations with high and low IT using six-digit occupational variation from the Occupation and Employment Statistics (OES), displayed in Figure A.3. The benefit of these data is that they are comprehensive and contain the most granular variation available that can match with our IT intensity index. We find that there is a remarkable difference in the distribution of hourly wages between jobs with more versus less IT intensity. The difference in employment and inequality between the two are a little more subtle, although high IT jobs tend to be larger in both respects.

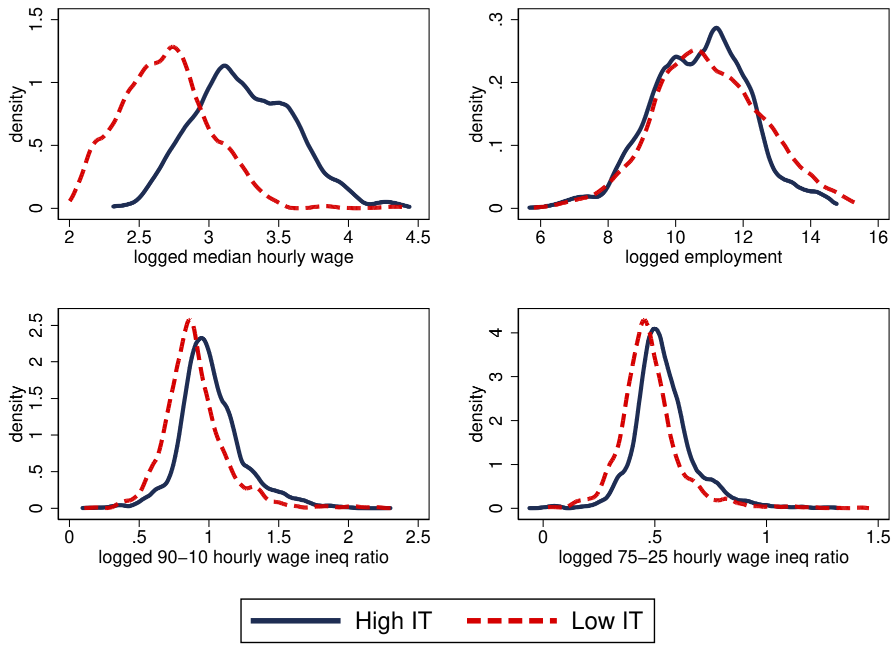

Figure A.3: Distribution of Hourly Wages, Employment, and Inequality, by IT Intensity

Notes.–Sources: Bureau of Labor Statistics Occupation and Employment Statistics (national) tables, 2000-2015. The figure plots the distribution of logged median hourly wages, employment, and hourly wage ratio inequality proxies for high and low IT-intensive occupation.

We subsequently examine the cross-sectional differences in earnings and employment premia by major industry. Figure A.4 plots logged income in IT-intensive jobs net of income in non-IT-intensive jobs, together with the IT employment share, separately by major industry, using the Census micro-data for 1980 and 2013-2015. Whereas in some industries there is a very small premium (e.g., wholesale), in others there is a high premium. For example, in finance, insurance, and real estate, IT-intensive workers earn approximately 70% more than their counterparts. While both FIRE and manufacturing sectors have similar IT income premia, the employment share of IT workers is much greater in FIRE than it is in manufacturing (e.g., 40% versus 25%).

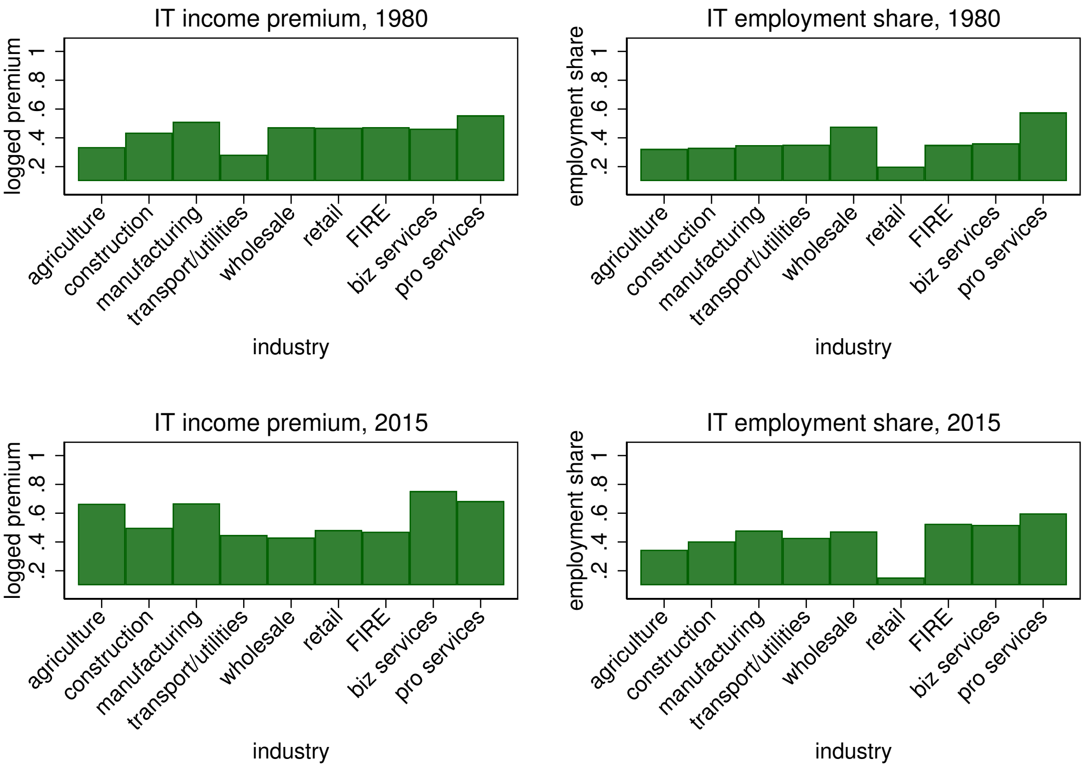

Figure A.4: Earnings and Employment Premia, by Industry and Year Group

Notes.–Sources: Census Bureau and O*NET, 1980 and 2013-2015. The figures plot the logged IT earnings premium obtained by taking logged labor income in IT-intensive jobs net of logged labor income from non-IT-intensive jobs across, and the IT employment share, both across industries. Observations are weighted by the survey sample weights.

Appendix A.3. Supplement to the Descriptive Evidence and Facts

Appendix A.3.1. Robustness of Controlling for Skill Intensity

The main text presents evidence showing that hourly wages are significantly higher in jobs with high IT intensity, relative to low intensity. However, once concern is that IT intensity is simply correlated with other valuable skills through selection channels. We address this concern by exploiting cross-sectional variation through regressions of the form

$$y_{ot} = \alpha skill_o + \gamma_t IT_o + \psi_{o'} + \epsilon_{ot} \tag{A.1}$$

where y denotes our outcome variable of interest (logged employment, inequality, and hourly wages), $skill$ denotes a vector of occupation-specific skills, IT denotes the intensity of information technology, and $\psi_{o'}$ denotes fixed effects on four-digit occupation cells. We present two main sets of estimates for Equation A.1: unconditional and conditional correlations. The conditional correlation estimates illustrate that the return to IT intensity is large, even after controlling for the skill content associated to different tasks. These fixed effect estimates also illustrate that the wage premium persists after controlling for non-random sorting of different workers into different occupations at a detailed four-digit level.

Table A.1 documents these results. Beginning with the logged hourly wage as the outcome variable, the unconditional correlation estimate suggests that a standard deviation rise in information technology is associated with a large 0.45% rise in the median logged hourly wage. The estimate is still statistically and economically significant after introducing detailed measures of skill intensity and four-digit fixed effects. Importantly, IT intensity is approximately half as large in magnitude as the association between hourly wages and non-routine & cognitive skills, which suggests that IT intensity is not merely a proxy for high skilled occupations. We also find that, in the cross-section, a standard deviation increase in IT intensity is associated with a large 0.57% decline in occupational employment and a 0.12% rise in the 90-10 logged hourly wage difference.

While the employment and inequality differences between IT and non-IT jobs are stark in the cross-section, they become statistically insignificant once we control for four-digit occupation and year fixed effects. Only an hourly wage premium remains: a standard deviation rise in IT intensity is associated with a 0.13% rise in the hourly wage. Regressions that control for occupation and year fixed effects also contain the standard set of skill intensity measures from the skill-biased technical change literature (Acemoglu and Autor, 2011). Strikingly, a standard deviation rise in non-routine and cognitive skills is associated with almost as large of an increase in the hourly wage as a proportional increase in IT intensity—0.12 for non-routine & cognitive versus 0.13 for IT. We also find that an increase in non-routine & cognitive is associated with a statistically significant decline in employment and a statistically significant rise in inequality, consistent with early evidence from Autor et al. (2003).

Table A.1: IT Intensity and Task Content of Jobs

Dep. var. =	logged hourly wage		logged employment		logged 90-10 ratio	
	(1)	(2)	(3)	(4)	(5)	(6)
information technology	.45***	.13***	-.57***	.02	.12***	.00
	[.02]	[.02]	[.10]	[.09]	[.02]	[.01]
non-routine, cognitive		.12***		-.28***		.05***
		[.02]		[.11]		[.02]
non-routine, non-cognitive		-.02		.06		-.01
		[.02]		[.08]		[.02]
routine, cognitive		.04***		-.05		.01
		[.01]		[.06]		[.01]
routine, manual		-.02		-.14		-.09***
		[.02]		[.12]		[.03]
non-routine, manual		.03		.17		.04*
		[.03]		[.12]		[.02]
R-squared	.63	.92	.13	.64	.22	.73
Sample Size	6897	6897	7516	7516	6979	6979
Controls	No	Yes	No	Yes	No	Yes
4-dig Occupation FE	No	Yes	No	Yes	No	Yes
Year FE	No	Yes	No	Yes	No	Yes

Notes.–Sources: Occupation and Employment Statistics (BLS) and O*NET, 2004-2015. The table reports the coefficients associated with regressions of logged employment, median hourly wages, and the logged 90-10 hourly wage differential on a standardized measure of IT intensity and standardized measures of skills following the strategy in Acemoglu and Autor (2011). Standard errors are clustered at the six-digit occupation level.
h

Appendix A.3.2. Time Series Variation in IT Tasks Intensity

We provide a brief characterization of the time series heterogeneity in aggregate IT intensity by plotting the evolution of each of its sub-components between 2000 and 2016, weighted by each six-digit occupation's employment from the Occupational Employment Statistics (OES) data. Table A.2 reports the mean and standard deviation of each input intensity over time. Across most of the categories, the intensity is growing and generally matches our intuition about the types of tasks that have become more common, such as processing information, analyzing data, and task related to engineering and technology. However, dispersion in the intensity is staying roughly constant across most categories.

We now provide another way of characterizing the heterogeneity more visually. In particular, looking at the mean IT intensity might confound heterogeneity in the distribution of IT intensity across occupations. While the main text provides some useful heuristics that display the average IT intensity over time, it confounds a significant amount of heterogeneity across 773 unique six-digit occupations. Figure A.5, therefore, plots the distribution of each input to the aggregate index in both 2004-2006 and 2014-2016 time periods. While a few of the distributions are relatively time invariant (e.g., quality control and programming), many distributions exhibit interesting changes. For example, both the "computers & electronics" and "interaction with computers" categories grow in both mean and skewness.

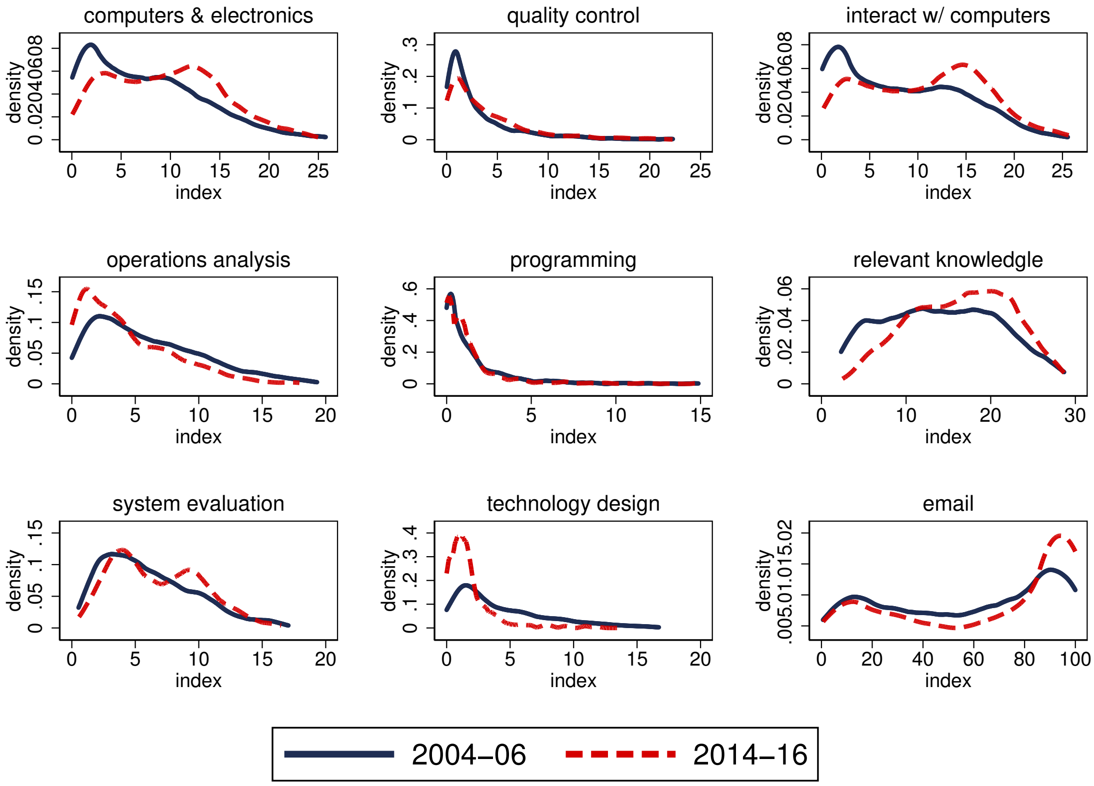

Figure A.5: Distribution of Inputs to Information Technology Intensity

Notes.–Sources: O*NET. The figure plots the intensity of the nine inputs to the aggregate information technology (IT) index for 2004 and 2016. The annual measures are weighted across six-digit occupations using average employment from the Occupation Employment Statistics (OES).

Appendix A.3.3. Robustness: Labor Productivity Measures

The main text reports time series for real labor productivity, defined as value added per hour worked, using KLEMS data from 1950 to 2010. Sectors are defined as high IT if the weighted z-score of IT intensity, obtained from matched O*NET and Census micro-data for 2000, 2006-2008, and 2013-2015, is greater than zero. Our classification implies that NAICS22, NAICS324, NAICS325, NAICS333, NAICS334, NAICS335, NAICS336, NAICS339, NAICS51, NAICS52, NAICS541, NAICS55, NAICS61, NAICS92 are high IT sectors, with the remaining sectors as low IT (NAICS11, NAICS21, NAICS23, NAICS311, NAICS314, NAICS321, NAICS322, NAICS326, NAICS327, NAICS331, NAICS42, NAICS44, NAICS451, NAICS452, NAICS453, NAICS454, NAICS48, NAICS531, NAICS532, NAICS561, NAICS562, NAICS62, NAICS71, NAICS72, NAICS81.

Table A.2: Time Series of Technology and Information Services Inputs

	2004-2005		2006-2007		2008-2009		2010-2011		2012-2013		2014-2015	
	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
computers and electronics	6.4	4.9	7.3	5.0	7.5	5.0	8.2	5.6	8.6	5.9	8.8	5.9
drafting	2.2	3.2	2.4	3.3	2.7	3.5	2.8	3.7	2.9	3.6	2.8	3.5
interacting with computers	8.3	6.1	8.5	6.0	8.9	6.0	9.6	6.1	10.1	6.2	10.4	6.2
operations analysis	4.6	4.0	5.8	4.2	4.8	3.9	3.5	3.6	3.4	3.6	3.2	3.4
programming	1.2	2.1	1.4	1.7	1.4	1.6	1.0	2.0	1.2	2.4	1.2	2.3
relevant knowledge	12.2	6.2	12.6	5.8	13.0	5.8	13.9	5.8	14.3	5.8	14.4	5.8
systems evaluation	5.3	3.8	5.1	3.6	5.3	3.7	5.7	3.6	6.0	3.6	6.0	3.6
technology and design	2.9	2.9	4.2	3.4	3.7	3.1	1.2	1.6	1.4	1.8	1.4	1.7
engineering and technology	2.8	4.5	3.2	4.6	3.4	5.0	3.7	5.0	3.7	4.9	3.7	4.8
managing resources	4.1	3.4	5.2	3.9	4.8	3.6	2.9	2.6	2.9	2.5	2.9	2.5
analyzing data	8.8	5.8	8.4	5.4	8.7	5.6	9.4	5.8	9.9	5.8	10.0	5.8
processing information	10.8	5.6	11.1	5.4	11.5	5.5	12.0	5.6	12.3	5.5	12.5	5.5
Observations	1381		1376		1389		1451		1520		1540	

Notes.–Sources: O*NET, 2000-2016. The table reports the means and standard deviations of each information technology sub-index at a six-digit level of occupation aggregation weighted by employment. Observations are weighted by employment.

One interesting result from our classification is the high value added per worker in the high IT manufacturing sector. A possible explanation behind these levels' differences is their extensive use of capital equipment, particularly since these are mostly primary metals sectors. Figure A.6 addresses this concern by netting out capital compensation from value added, which produces a high IT services time series closer to high IT manufacturing. In fact, the growth in labor productivity under this approach is 700% for high IT services, whereas it is only 275% for high IT manufacturing (versus 506% and 327%, respectively, when using the baseline definition that does not net out capital compensation).

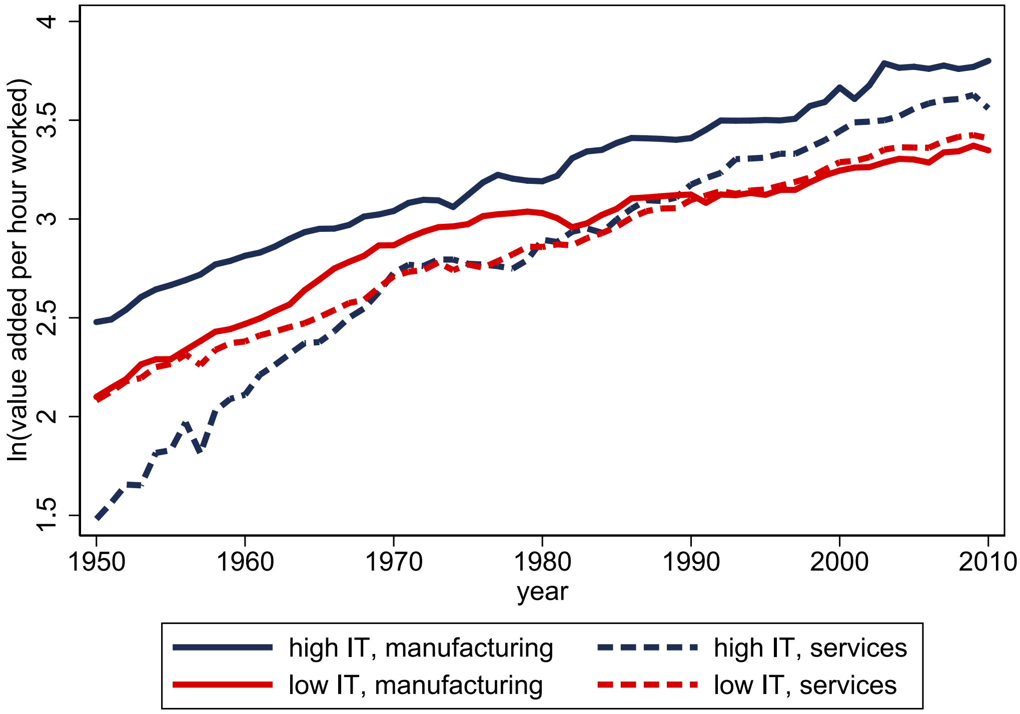

Figure A.6: Value Added per Hour Worked in Manufacturing and Services, 1950-2010

Notes.–Sources: USA KLEMS, 1950-2010. The figure plots sectoral average annual value added *net of capital compensation* per worker in constant thousands of dollars (deflated using the 2009 GDP deflator from the St. Louis Fed) for high and low information technology manufacturing and services sectors. Sectors in each high/low IT and manufacturing/services category are averaged based on the long-run employment average between 1950 and 2010. Industries are classified as high information technology if the weighted average z-score of information technology intensity within three-digit NAICS codes is greater than zero. We construct these weighted average scores using a sample of Census micro-data for 2000, 2006-2008, and 2013-2015 matched with IT intensities at a five-digit level as implemented in the baseline definition. The figure shows that logged value added per hour worked has grown primarily in the high IT manufacturing and services sector with the bulk of the growth driven by high IT services.

As an additional point of comparison, we partition sectors in a different way. This allows us to visualize how different groupings of high and low technology sectors influence our understanding of heterogeneity within the services sectors. We use an alternative classification based on a "cherry-picking approach", in which we select sectors as high tech based on whether the sector sounds like one that makes intensive use of technology. In particular, we define high tech services sectors as publishing / software, motion picture / sound recording, broadcasting / telecommunications, information / data processing, federal reserve banks / credit intermediation, securities and commodity contracts, insurance carriers, funds / trusts, real estate, rental / leasing services, legal services, computer systems design / related, whereas low tech services sectors are wholesale trade, retail trade, social assistance, performing arts, amusements gambling / recreation, accommodation, food services / drinking, and other services.

We use value added in constant dollar of labor compensation as a measure of labor productivity . We turn to this definition from the 2017 release of the KLEMS data because it holds constant potential differences in prices at a sectoral level, but it does not contain a measure of employment or total hours worked, which we would otherwise use as a closer point of comparison with our baseline. Figure A.7 plots the resulting time series for manufacturing (pooled) and both high and low tech services. We find that the bulk of the increase in labor productivity is driven by the high tech services sector. While our selection criteria for high tech services is admittedly cherry-picked, these results show that our baseline approach provides, if anything, a conservative measure of value added in high tech services.

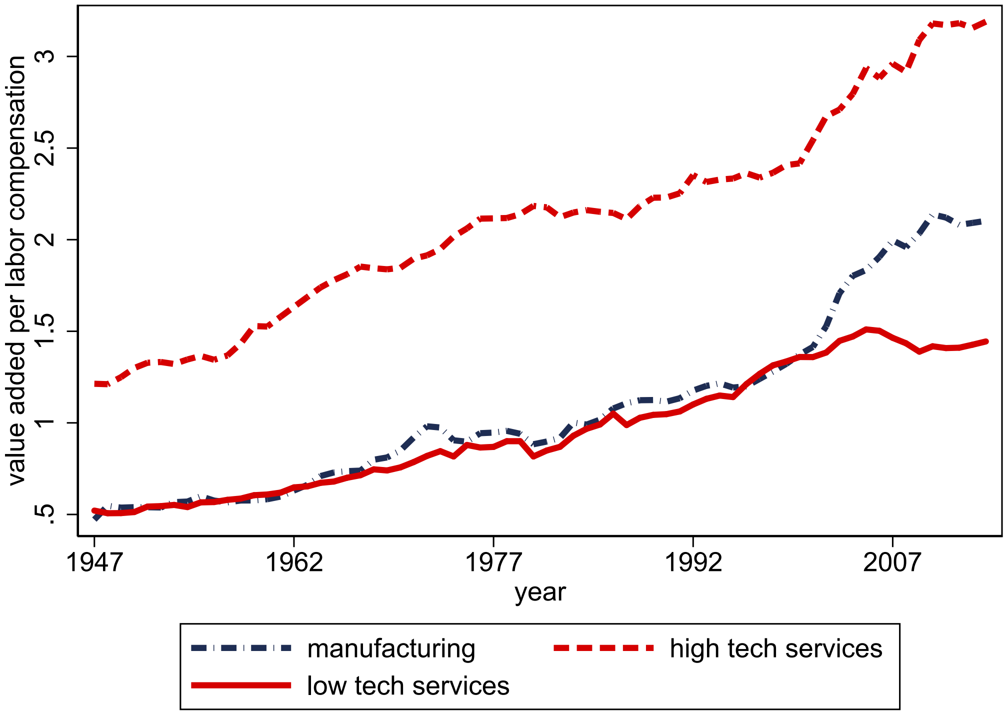

Figure A.7: Value Added / Employee Compensation in Manufacturing and Services, 1947-2014

Notes.–Sources: USA KLEMS, 1947-2014. The figure plots sectoral average annual value added per labor (employee) compensation in constant thousands of dollars for manufacturing and services sectors. Sectors in each category are averaged based on the long-run employment average between 1947 and 2014. High skilled services sectors cover information and financial sub-sectors, including: publishing industries (including software), motion picture and sound recording, broadcasting and telecommunications, information and data processing, federal reserve banks credit intermediaries, securities commodity contracts, insurance carriers, funds, trusts and other financial vehicles, real estate, rental and leasing, legal, and computer systems design. Low skilled services sectors include wholesale / retail trade, social assistance, and other services (typically repair).

Appendix A.3.4. Robustness: Earnings and Employment Premia

We begin by showing further robustness of the earnings and employment premia to alternative refinements of the definition, namely the hourly wage and total hours (across all workers) premia. Panel A in Figure A.8 shows a similar result as the main text that the hourly wage premium grew nearly identically in both manufacturing and services sectors from 1980-2013. The premium is also quite quantitatively close to the earnings premium, which suggests that hours differences play only a small role. Panel B tells a similar story. The observed total hours worked premium is similar as the employment premium in the main results, which suggests that the intensive margin differences in hours worked per worker in IT and non-IT jobs are fairly similar. There is a marginally larger quantitative narrowing of the IT total hours premium, relative to the employment premium, which follows from the fact that IT workers spend roughly 131 more hours worked per year (relative to their counterparts).[1]

[1]The difference is larger in the manufacturing sector with 143 hours worked more among IT workers versus only 125 hours worked more in the services sector. The estimate is conditional on family size, race, gender, marital status, and schooling.

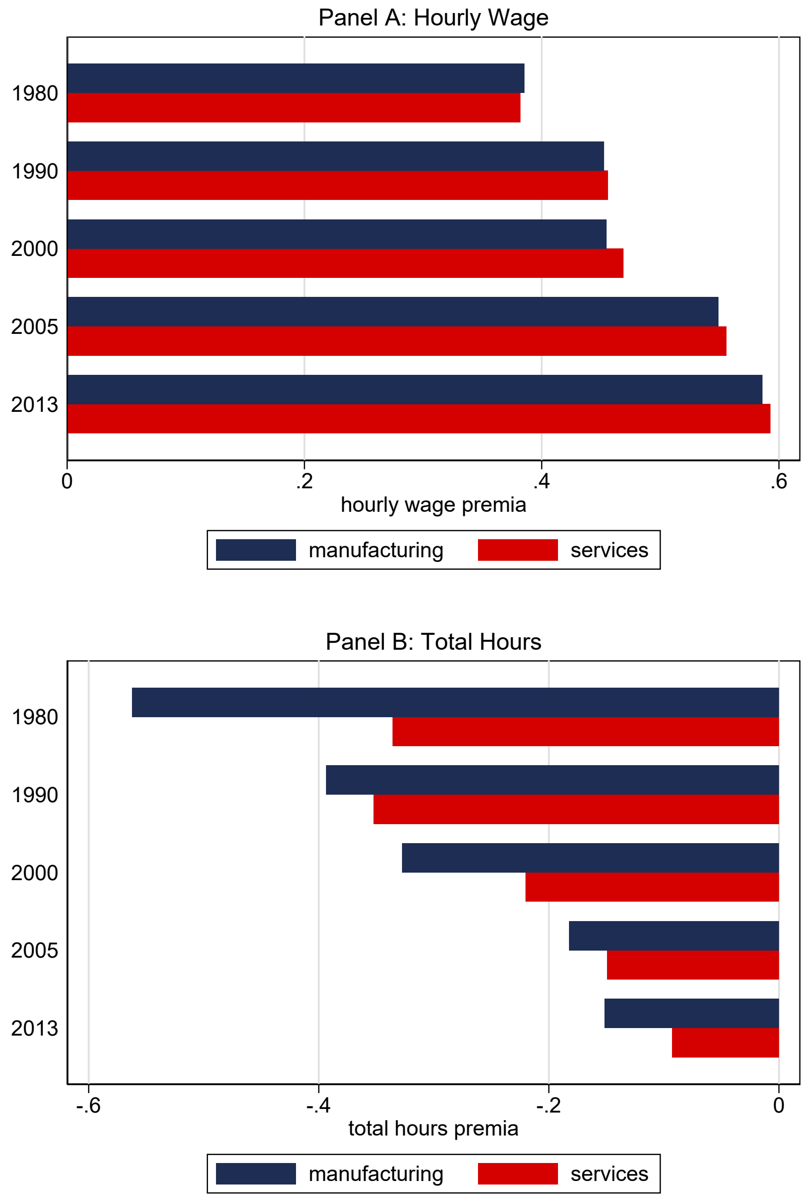

Figure A.8: Hourly Wage and Total Hours Premia to Information Technology, 1980-2015
Notes.–Sources: Census Bureau and O*NET, 1980-2015. The figure plots the hourly wage and total hours (annual hours × number employees) in high technology and information services jobs between 1980 and 2015. Earnings is deflated using the 2010 personal consumption expenditures index. We start with 1980 because hours worked is not available in 1970.

We now turn to producing similar patterns on employment and earnings premia when using the annual CPS. The primary downside is that the estimates are more noisy due to sampling variability in the workers who are surveyed at the three-digit occupation level. Nonetheless, Figures A.9 and A.10 plot these premia. The crucial observation in Figure A.9 is that, while the employment share in manufacturing is declining (thinned dotted blue line), the employment share of IT jobs in the manufacturing sector is rapidly rising (thick solid blue line). The rise of IT jobs accounts for nearly all of the growth in the services sector.

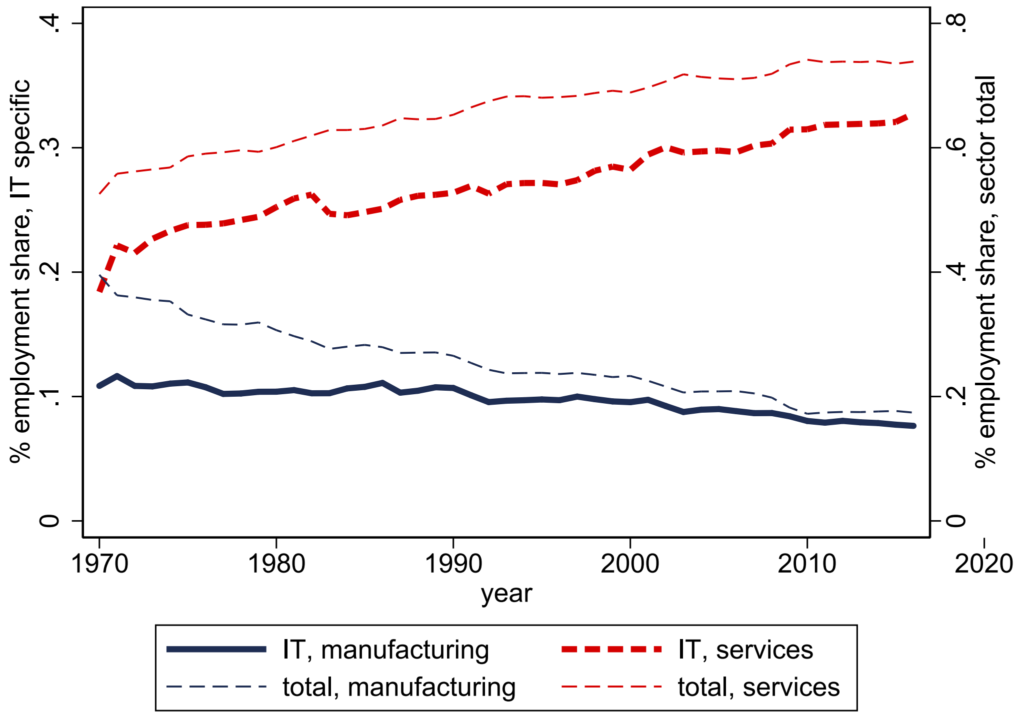

Figure A.9: Employment Shares in Technology and Information Services—Sectoral and Aggregate, CPS Robustness
Notes.–Sources: Current Population Survey (CPS) ASEC, 1970-2015. The figure plots several employment shares. The left y-axis plots the employment share of manufacturing and services information technology workers ("IT v. non-IT") relative to their sectoral total. The right y-axis plots the employment share of manufacturing and services overall ("total"). As such, the level of the right y-axis will be larger than the level of the left y-axis because it includes non-IT workers as part of the share. The purpose of providing both sets of trends is to illustrate the IT share in manufacturing and services in light of the overall trends in these two sectors.

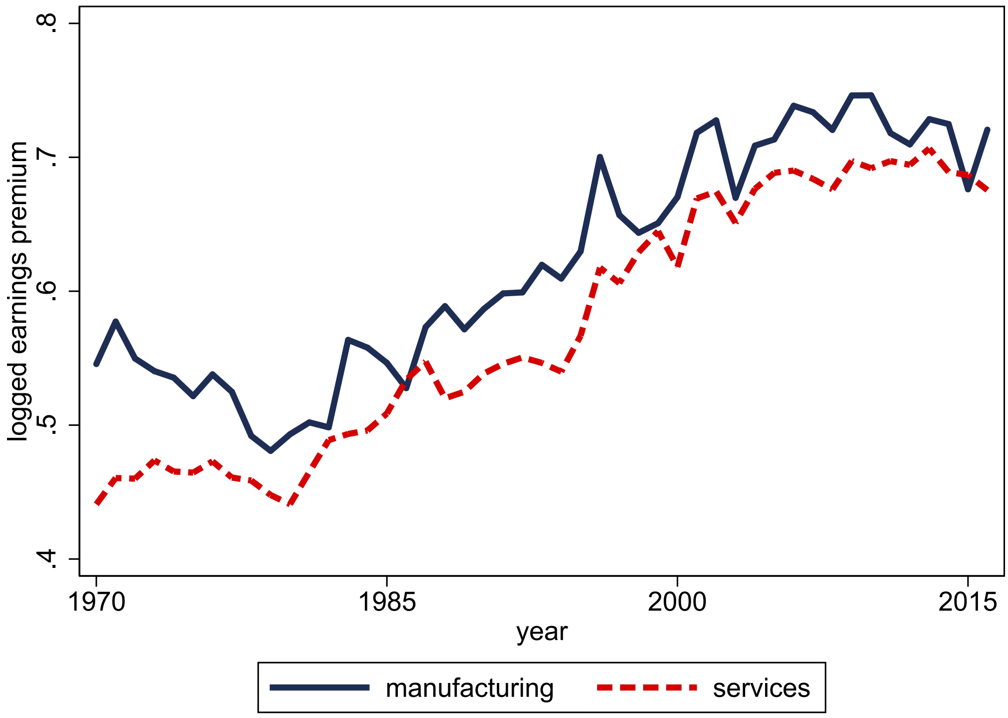

Figure A.10: Earnings in Technology and Information Services, CPS Robustness

Notes.–Sources: Current Population Survey (CPS) ASEC, 1970-2015. The figure plots earnings premia in manufacturing and services sectors where the technology and information services (IT) intensity is based on being above the median score.

Figure A.11 also examines heterogeneity in the earnings premia using two different classification strategies: K-medians and K-means with $K = 2$. K-medians produces an almost identical result as the regular medians-based approach, which is not surprising since our classification is based on the single IT score. Using K-means tends to classify more occupations as IT-intensive, which is why the earnings premia is marginally smaller—but clearly not by much by any stretch of the imagination.

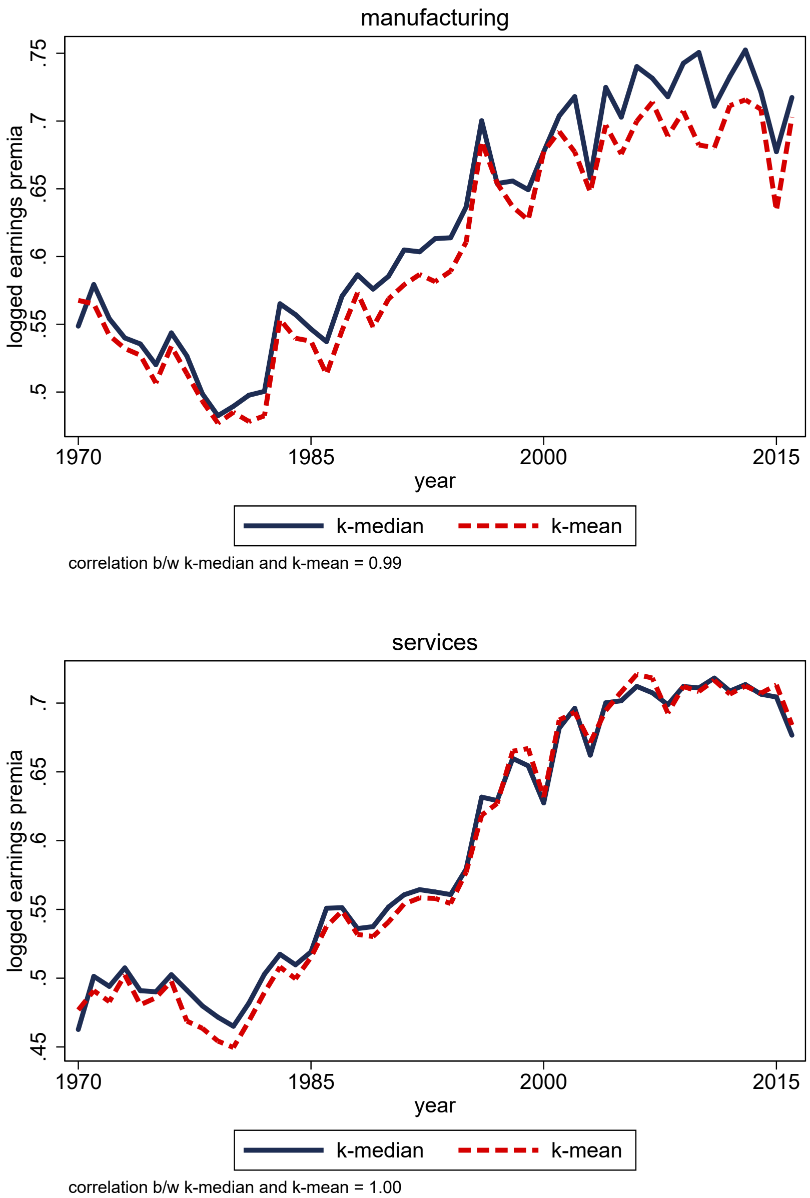

Figure A.11: Comparing CPS Earnings Premia with K-medians and K-means

Notes.–Sources: Current Population Survey (CPS) ASEC, 1970-2015. The figure plots earnings premia in manufacturing and services sectors where the technology and information services (IT) intensity is based on two approaches: using K-medians and K-means estimators for $K = 2$.

While the two series produce qualitatively similar series (e.g., the correlation between the earnings and employment premia in the two are 0.88 and 0.86, respectively), there are at least two reasons the OES data generates quantitatively different series. The first is that, since the CPS does not aim to be representative of every occupation at a detailed five-digit level, it may overstate one type of worker over another. These concerns are potentially amplified by the presence of occupational misclassification, which has been documented by Kambourov and Manovskii (2013). The second is that the CPS data only contains a five-digit occupation classification, whereas the OES data contains a six-digit classification. To the extent there is some detailed within-occupation reallocation, differences can emerge.

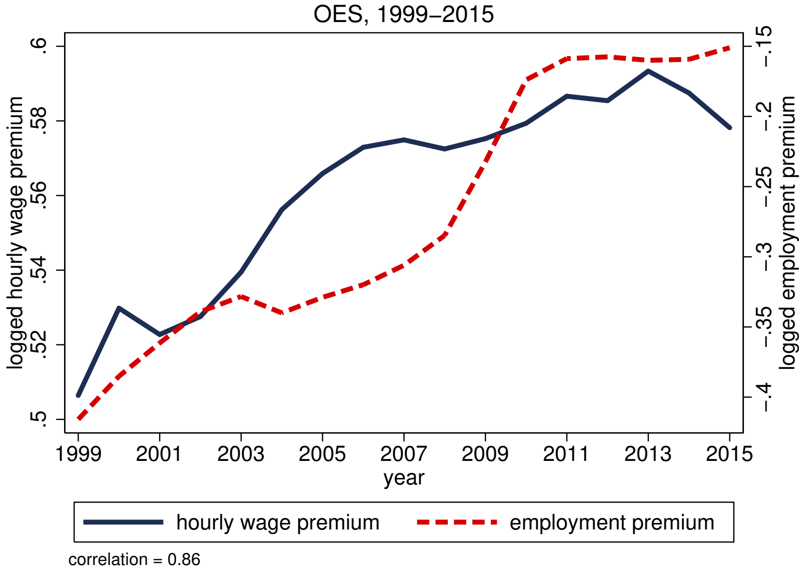

Figure A.12: Earnings and Employment in Information Technology Jobs, National
Notes.–Sources: Bureau of Labor Statistics Occupational Employment Statistics, 1999-2015. The figure plots the logged earnings premium between information technology and non-IT workers (deflated using the 2010 personal consumption expenditures index) and the logged employment premium. The employment premium is computed by summing across all IT and non-IT workers at a six-digit occupation level.

Given the large rise in both the earnings and employment IT premia, we also examine how it interacts with other documented premia in the labor market, in particular the returns to tenure (we later examine the interaction with the returns to education). Using the CPS supplement on job tenure between 1996 and 2014, we are able to non-parametrically characterize the IT premium across the tenure distribution. To do this, we regress logged hourly wages on a vector of controlling covariates (age, number of children, race, gender, marital status, education), subsequently averaging across the residualized earnings measures for each tenure bin. We include these controls to mitigate the potential effects of the composition of the labor force throughout the tenure distribution. We separately plot these returns for the 1996-2002 and 2010-2014 time periods to understand the extent to which these returns might have shifted over time.

Figure A.13 documents these results. First, and not surprisingly, the IT premium has grown over the past decade and it has shifted the earnings premium up across the entire tenure distribution. For example, between 1996-2002, IT workers with zero years of tenure would earn 30% more than their non-IT counterparts, whereas between 2010-2014 they would earn 37% more. Looking at the top of the tenure distribution, however, IT workers with 16-20 years of tenure earned roughly 24% more than their counterparts between 1996-2002, but between 2010-2014 the premium grew by only one to two percentage points.

Second, and more importantly, the IT premium is declining in tenure. While it is possible that technology companies—which employ a large share of IT workers—simply tend to have lower average tenure due to something embedded within their underlying production function, we provide evidence that the declining premium is a natural result of a career ladder where skilled workers begin in an IT-intensive job and then progress to a managerial job. Such examples are commonplace in many technology hubs, like Silicon Valley or San Francisco, where a skilled worker may begin as a data scientist or consulting out of undergraduate and then transition towards a senior managerial role after roughly a decade. The main reason for this is that, as individuals progress in the career ladder, managerial skills become increasingly important and outpace the importance of IT skills. In particular, while it is true that managers will leverage IT to broaden their span of control, their comparative advantage in management begins to outweigh their absolute advantage.

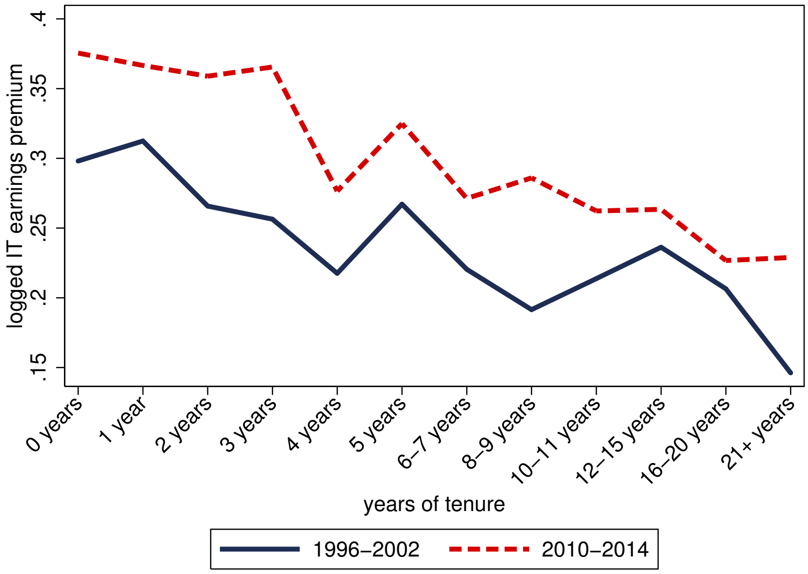

Figure A.13: Hourly Wage Premium, by Employee Tenure

Notes.–Sources: Current Population Survey Job Tenure Supplement. The figure plots the residualized logged earnings premium between IT and non-IT-intensive jobs (identified using the O*NET measurements of tasks) across the distribution of tenure using the CPS tenure supplement covering 1996, 1998, 2000, 2002, 2004, 2006, 2008, 2010, 2012, and 2014. Logged earnings are residualized by regressing them against a vector of controls, including age, education, race, gender, marital status, and number of children, taking the residual as the residualized logged earnings measure. Observations are weighted by the job tenure survey sample weights.

Appendix A.3.5. The Declining Premium for C/NR Tasks and College

The main text presents results about the earnings premium between IT and non-IT jobs restricted to the set of college degree workers, and separately for Cognitive and Non-Routine (C/NR) jobs. In both cases, the growth rate from 2000 and 2015 is greater within-group than it is across groups—that is, restricted to the set of college degree workers, the growth in the IT premium is greater than the growth in the college premium across all in the labor force. We now present complementary plots that characterize the employment premia between these jobs.

Beginning with Panel A in Figure A.14, it is remarkable how much of the employment share of college workers is accounted for by high IT workers. However, there is important longitudinal variation. In the 1970s, these IT jobs accounted for almost all of the share of college degree workers; roughly 15% had a college degree and 12% were in IT jobs. However, by 2015, nearly 37% of the labor force has a college degree and 27% are in IT jobs. The gap between the college share and the “college + IT” share reflects the surge in growth of universities and expanded access to a four-year college degree. Turning towards Panel B in Figure A.14, we see that there is not a substantial difference between the IT + C/NR and C/NR shares—that is, an even smaller difference in comparison with the college share. For example, in 1970, roughly 27% of jobs were classified as high C/NR and nearly 26% of those jobs were also classified as high IT. By 2015, the share of high C/NR jobs grew to 45% and the share of “IT + C/NR” jobs grew to nearly 40%.

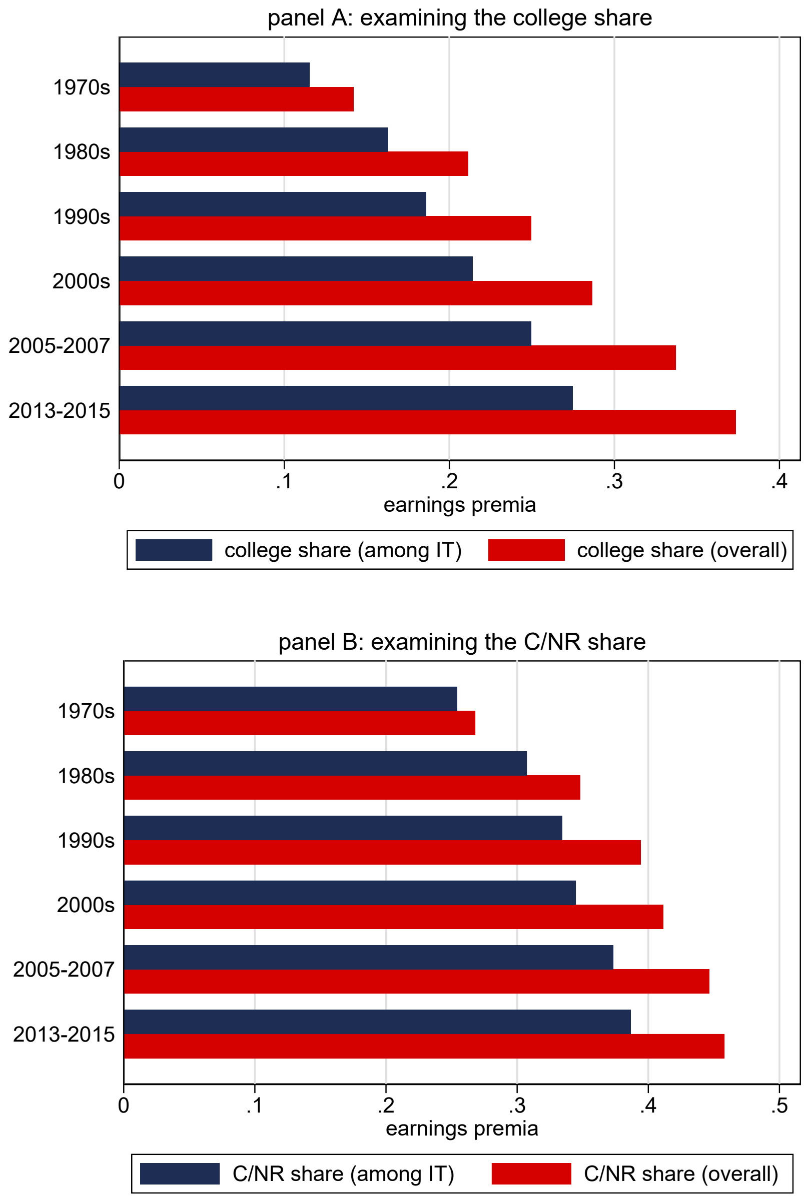

Figure A.14: Examining Heterogeneity in the Demand for Skills in High and Low IT Jobs (Employment Shares)
Notes.–Sources: Census Bureau and O*NET, 1970-2015. Panel A plots the employment share of college and high IT workers relative to the total labor force, together with the overall college share. Panel B plots the employment share of high cognitive and non-routine and high IT occupations relative to the total labor force, together with the overall share of high C/NR jobs.

We examine the evolution of employment shares across the four permutations of IT/non-IT and college/non-college jobs in Figure A.15. Importantly, the share of IT and college degree workers has grown from approximately 20% of the labor force to nearly 32%, and it dwarfs the marginal rise in non-IT and college degree workers from 2% to roughly 5% over the past 40 years. Symmetrically, the share of IT and non-college degree workers declined from 32% to 22%, which also dwarfs the decline in non-IT and non-college degree workers from 45% to 42%.

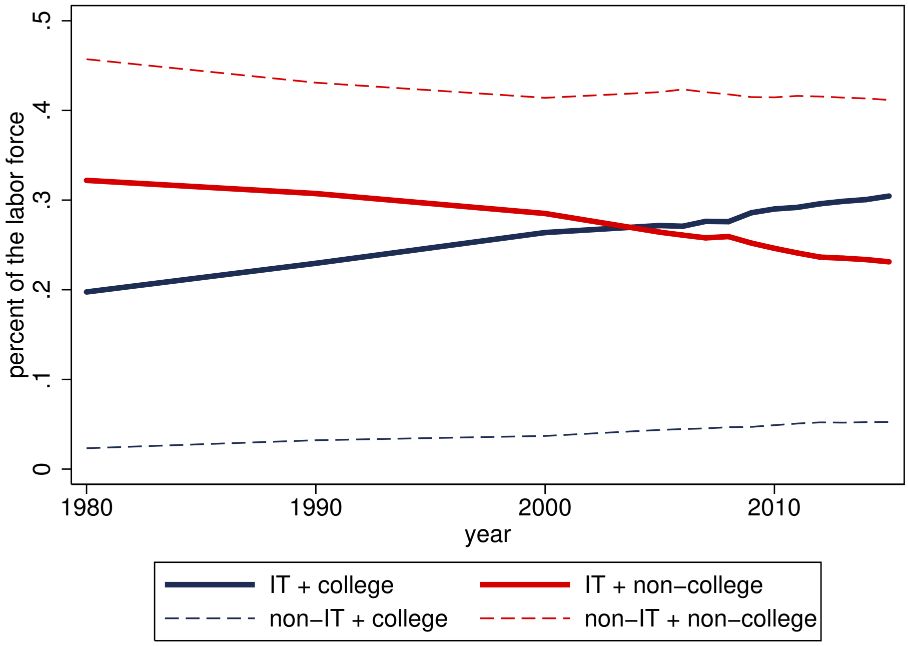

Figure A.15: Employment Shares in Information Technology Jobs, by College Attainment
Notes.–Sources: Census Bureau, 1980-2015. The figure plots the share of full-time workers across the permutations of IT versus non-IT-intensive jobs and college and non-college degree workers.

Appendix A.3.6. IT Premia using the Census Classification

In the body of the paper we present a comparison of earnings and hours premia obtained using either our IT intensity measure or the industry-based measure from Hecker (2005). While these measures capture different features of the data—in part because Hecker (2005) focuses on industries, whereas we focus on jobs—we now turn towards an additional comparison of information technology jobs discussed by Beckhusen (2016). The Census began identifying IT jobs in 1970, capturing roughly 0.6 percent of the labor market. Beckhusen (2016) discusses how the measure has expanded since then and grown to 2.9% of the population. These workers represent roughly 16% of our Census sample. We proceed to compare our measure to the Census measure of IT jobs by examining state panel regressions in first-differences for 1980, 1990, 2000, 2005, and 2015 where we regress earnings and hours worked growth on growth in the IT jobs' share, conditional on controls. Our objective is to provide additional context for our results.

Table A.3 documents our findings. Using our IT intensity measure we estimate that a one percentage point (pp) increase in the share of IT workers is associated with a 0.16pp rise in the growth rate of earnings, whereas a one pp rise in the share of Beckhusen (2016) IT jobs is associated with a 0.06pp rise. Once we add demographic controls we find that the gradient for our measure rises sharply to 0.72pp, whereas the gradient rises for Beckhusen (2016) IT jobs rises only marginally to 0.10pp. We also find that a one pp rise in the our IT intensity share is associated with a 0.12pp rise in the growth rate of hours worked, whereas we find that a one pp rise in the Beckhusen (2016) share is associated with a 0.02pp rise, conditional on controls. Regardless of the IT proxy we use, increases in the share of IT jobs are associated to increases in earnings and hours worked.

Appendix A.4. A Model of Production with Heterogeneous Tasks

To explicitly account for selection-effects due to unobserved heterogeneity among workers, we develop a model based on the original contribution of Adao (2016). The model formalizes an assignment problem

Table A.3: Comparison of Baseline and Beckhusen (2016) Measures

Dep. var. =	annual earnings growth				annual hours growth			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
IT share, growth	.16**		.72***		.21***		.12***	
	[.07]		[.11]		[.02]		[.02]	
Beckhusen share, growth		.06***		.10***		.03***		.02***
		[.01]		[.02]		[.00]		[.00]
R-squared	.02	.10	.48	.49	.38	.31	.67	.71
Sample Size	255	255	255	255	204	204	204	204
Controls	No	No	Yes	Yes	No	No	Yes	Yes

Notes.–Sources: Census Bureau, O*NET, 1980-2015. The table reports the coefficients associated with regressions of state × decade annual earnings and hours worked growth on the share of individuals who are classified as high IT using our baseline O*NET measure and an alternative classification of IT jobs from Beckhusen (2016). Controls (in the labeled specifications) include state growth rates in number of children, family size, age, gender, race, and years of schooling. Standard errors are clustered at the state-level and observations are weighted by the number of observations at a state × period cell.

of tasks and technologies in the tradition of Autor (2013).[2] The model features both worker heterogeneity and sorting across tasks, which are identified using the employment shares in different demographic groups, regions, and time periods. Estimates based on this model deliver the implied (shadow) growth rates of IT task prices, relative to non-IT-intensive tasks, as well as estimates of the changes in labor quality composition.

The endogenous selection into sectors and task groups overcomes the conventional concerns associated with the approach of Juhn et al. (1993) since residual wage variation might simply reflect measurement error or unpriced amenities (Lemieux, 2006). After estimating the model, we (i) examine the distribution of the implied price and worker quality composition effects, (ii) use them to recover, in a second step, the elasticity of substitution between IT and non-IT labor inputs, which we compare to a baseline approach of adjusting for labor services quality.

Appendix A.4.1. Task Production

Following Adao's approach, we partition jobs into groups based on geography (state) and college attainment, denoted as $g \in \Im = \{1, 2, .., G\}$, and we allow for mobility between the manufacturing and services sector. We assume that individuals enter the labor market as either college or non-college educated workers.[3] We view states as small open economies with segmented labor markets.[4] Each industry (manufacturing or services) produces its final output by combining two inputs: an IT-intensive and a non-IT-intensive labor aggregate. We assume that these intermediate labor inputs are produced by aggregating individual tasks, denoted j, supplied by different occupations. For simplicity, in the following model description we abstract from industry and we focus on tasks; however, the choice of industry is explicitly accounted for in our empirical analysis.

Workers are heterogeneous in their idiosyncratic productivity at performing different tasks. That is, each worker is endowed with a productivity vector $[L_g^{IT}(i), L_g^{NIT}(i)]$ describing her ability to perform IT and non-IT-intensive tasks.[5] To accommodate the variety of occupations that workers perform, we assume

[2] We adapt Adao's original structure to our setting, deferring readers to Adao (2016) for derivations and a more comprehensive discussion of analytical results.

[3] We do not model the option of returning to school for continuing education.

[4] One could work at a commuting zone level, but we choose not to for two reasons. First, at a lower level of aggregation, our 5% IPUMS sample begins running low on number of observations, producing imprecise shares of IT employment. Second, in talking with geographers and the Census, the classification of commuting zones before and after 2000 has changed quite considerably. While we could work at a county level, which has remained consistently defined over time, our sample size issue would be amplified even further.

[5] In our empirical implementation we allow the productivity vector to vary both by industry and task group. However, in what follows we abstract from industry-specific comparative advantage for ease of exposition.

that each aggregate occupation group $k \in \{IT, NIT\}$ is a collection of multiple (perfectly competitive) “single-task occupations” $j \in \Im^k$. Each single task output is used to produce the intermediate labor output k.

The price of the task output of occupation j is denoted as r^j. Like Adao (2016), we assume that workers have identical k-specific productivity within an occupation. That is, workers perform with identical efficiency across all occupations within k. Hence, an occupation j produces task q^j as follows:

$$q^j = Q^j(L_1^j, L_2^j, ..., L_G^j),$$

where $L_g^j = \int_{S_g^j} L_g^{IT}(i)di$ if workers are employed in an IT-intensive occupation and $L_g^j = \int_{S_g^j} L_g^{NIT}(i)di$ if workers are employed in a non IT-intensive occupation. Like Adao, we also assume that the function $Q^j(\cdot)$ is strictly increasing, concave, differentiable and homogeneous of degree one. The technology $Q^j(\cdot)$ allows for the possibility that the labor inputs of different worker groups are imperfect substitutes in production, although this is not required. The integration set S_g^j is the set of workers of type g who are employed in occupation j.

Individual efficiency units can then be used to define a notion of comparative advantage in different occupation groups, i.e. comparative advantage for an individual in the IT sector is $s_g(i) = \ln[L_g^{IT}(i)/L_g^{NIT}(i)]$, whereas absolute advantage is $a_g(i) = \ln[L_g^{NIT}(i)]$. We take individual productivity as exogenous—workers draw their efficiency from a bivariate distribution, i.e. $s_g(i) \sim F_g(s)$ and $\{a_g(i)|s_g(i) = s\} \sim \widetilde{F}_g(a|s)$.

Conditioning on product prices, the labor demand in occupation j of sector k (within group g) is given by:

$$\omega_g^k = r^j \frac{\partial Q^j}{\partial L_g^j} \quad \text{if } j \in \Im^k \tag{A.2}$$

where ω_g^k is the marginal product of workers in occupation class k and observable group g. Below we estimate how the relative productivity of IT and non-IT-intensive occupations has changed over time, and what these changes suggest about the substitutability of different jobs in the manufacturing and services industries.[6] As in Roy (1951), individuals choose the job that yields the highest utility, which in our case is merely a function of labor income. Letting $y_g^k(i)$ denote the potential logged hourly wage an individual i could earn in sector k, we write:

$$y_g^{NIT}(i) = ln(\omega_g^{NIT}) + a_g(i), \quad y_g^{IT}(i) = \omega_g^{IT} + s_g(i) + a_g(i)$$

where earnings in the IT sector are a function of both comparative and absolute advantage. Because individuals receive different wages based on their comparative advantage, they self-select into the jobs that offer a higher income. Hence, the set of individuals employed in a given sector k can be characterized as:

$$S_g^k \equiv \{i \in \mathcal{I}_g : \ k = \arg\max\{y_g^{IT}(i), y_g^{NIT}(i)\}\}$$

Markets are perfectly competitive and wages are such that the demand for labor equals the supply. Since the core part of the model relies on differences in comparative advantage based on individual productivity, we can rank individuals within each group g by their comparative advantage quantile $q \in [0, 1]$, so that $\sigma_g(q) \equiv (F_g)^{-1}(q)$ denotes an individual's efficiency in the IT sector based on their rank in the comparative advantage distribution.We can also denote the conditional distribution of absolute advantage as $\widetilde{F}_g(a|\sigma_g(q))$, with an average $\alpha_g(q)$ and variance $v_g(q)$. It follows that the logged wage schedule along the quantile range is :

[6] One potential driver of job substitution is offshoring (Schott, 2004). In fact, any force that shifts the demand for, or supply of, particular jobs is captured in the shadow price changes that we estimate. For example, if non-IT jobs are more likely to be outsourced, one might observe a rising share of IT-intensive jobs and a parallel slowdown in the growth of the shadow price of non IT-intensive jobs.

$$\overline{Y}_g^{NIT}(q) = \ln(\omega_g^{NIT}) + \alpha_g(q), \quad \overline{Y}_g^{IT}(q) = ln(\omega_g^{IT}) + \sigma_g(q) + \alpha_g(q)$$

As shown in Adao (2016), individuals sort into the IT sector if $\sigma_g(q) > \ln(\omega_g^{NIT}) - \ln(\omega_g^{IT})$, otherwise they will sort into the non IT sector. Employment composition is pinned down by marginal individuals with a comparative advantage equal to the relative efficiency-adjusted wage, i.e. $\ln(\omega_g^{NIT}) - \ln(\omega_g^{IT})$, such that $\ln(\omega_g^{NIT}) - \ln(\omega_g^{IT}) = \sigma_g(l_g^{NIT})$ where l_g^{NIT} denotes the employment share in the non IT sector. When we aggregate across the quantiles allocated to a sector, within each group g, we obtain the average logged wage in that sector:

$$\overline{Y}_g^k = \ln(\omega_g^k) + \overline{\alpha}_g^k(l_g^{NIT}) \text{ where } \overline{\alpha}_g^k(l) \equiv \begin{cases} l^{-1}\int_0^l \alpha_g(q)dq & \text{if } k = NIT \\ (1-l)^{-1}\int_l^1 (\sigma_g(q) + \alpha_g(q))dq & \text{if } k = IT \end{cases}$$

Appendix A.4.2. Estimating the Growth in Relative Task Prices: IT vs non-IT

For exposition, we abstract from the choice of industry and focus on a bivariate occupation choice (IT vs non-IT). This allows us to illustrate more clearly the way we estimate growth rates in relative returns to different tasks. In the subsequent section, we generalize the approach to account for both occupation and industry choice, so that in the empirical implementation we estimate changes in relative task returns across industries.

We index geographic region by r and let g denote an education-industry group. To obtain an empirical counterpart of the theoretical wage relationships based on comparative advantage, we use the first-order approximation from Adao (2016), which applies to quantiles of the observed wage distribution within each set (g, r, t) of observables. Letting $Y_{g,r,t}(\pi)$ denote the π-th quantile of the log-wage distribution of group g in region r in period t, we estimate

$$\Delta Y_{g,r,t}(\pi) = \Delta\ln(\omega_{g,r,t}^{IT}) + \left[\Delta\ln(\omega_{g,r,t}^{NIT}) - \Delta\ln(\omega_{g,r,t}^{IT})\right] l_{g,r,t_0}^{NIT}(\pi) + \mu_{g,r,t}X_{g,r,t}(\pi) + \Delta\nu_{g,r,t}(\pi) \tag{A.3}$$

separately for each group-by-region-by-year, where X denotes our usual set of demographic controls. As discussed in Adao (2016), $\Delta\nu_{g,r,t}(\pi)$ is a shock to the absolute advantage of workers in quantile π of the log-wage distribution. Equation A.3 identifies $\Delta\ln(\omega_{g,r,t}^{NIT})$ and $\Delta\ln(\omega_{g,r,t}^{IT})$ using the initial sectoral compositions, $l_{g,r,t_0}^{NIT}(\pi)$, and wage growths, $\Delta Y_{g,r,t}(\pi)$, across the quantiles π of the logged wage distribution. This approach leverages variation in the IT employment share across different groups at the start of the sample to identify the task prices growth, $\Delta\ln(\omega_{g,r,t}^{IT})$ and $\Delta\ln(\omega_{g,r,t}^{NIT})$. The intuition behind Equation A.3 is that, holding constant the distribution of employees at its initial value $l_{g,r,t_0}^{NIT}(\pi)$, any wage changes that accrue to workers in a given (r, g), between period $t-1$ and t, group must be attributed to changes in relative returns per efficiency unit. The estimator effectively captures that part of the price variation that cannot be attributed to worker composition effects over time.

The identifying assumption in the Adao (2016) approach is that, conditional on $X_{g,r,t}(\pi)$, pre-shock variation in sectoral employment composition is uncorrelated with variation in labor efficiency shocks among individuals with different levels of labor income within the same group-by-region-by-period. The condition is satisfied under the assumptions of a Roy model since marginal workers are indifferent between the two sectors and have similar returns in both. In other words, the assumption more generally implies that reallocation of marginal workers cannot have first-order equilibrium effects on the group's overall wage distribution. A related assumption we make is that there is sectoral reallocation between IT and non-IT jobs and between manufacturing and services jobs; see Adao (2016) for an original discussion of these assumptions.[7]

[7]One concern is that a non-IT-intensive worker in the manufacturing sector may be unlikely to switch into the services sector or re-train with new skills. We mitigate this concern by estimating our model separately for non-college and college workers, as well as controlling for a vector of demographics. More importantly, we examined this concern by computing the gradient between the employment share in manufacturing and the employment share in IT jobs at a county-level separately by five age brackets. We found that the gradients were between -0.40 and -0.46 for workers between ages 25 and 49, but -0.63 for

Having estimated $\Delta\ln(\omega^{IT}_{g,r,t})$ and $\Delta\ln(\omega^{NIT}_{g,r,t})$, one can recover the relative changes in workers' composition by taking the difference between raw wages growth and estimated task prices growth.[8] For each individual (g, r, t) triplet, we are able to recover the price growth and the change in average labor efficiency. These estimates are statistically significant for many of the triplets, although not for all. However, the real value of this procedure lies in the fact that it effectively allows us to approximate the entire distribution of price growth rates for any given group and point in time. Each individual triplet's growth rate can be considered as a (possibly noisy) estimate of a point in the distribution of growth rates for a given industry and group. For our subsequent analysis, we regard the (1980 weighted) mean and median of any such distribution as a reasonable approximation for the growth rate of prices in that industry and group.[9]

Appendix A.4.3. Accounting for Selection into Industry

In our empirical analysis, we explicitly account for the possibility that returns to tasks may vary by industry, and that workers may have a different comparative advantage across industries. To this purpose we extend the simple structure described in the previous section to let workers choose between manufacturing and services. In the simple bivariate choice problem discussed in the previous section, the marginal product growth for any given (g, r, t) triplet is a weighted average of the marginal products of workers in IT and non-IT-intensive occupations:[10]

$$\Delta Y_{g,r,t}(\pi) = \Delta\ln(\omega^{IT}_{g,r,t}) l^{IT}_{g,r,t_0}(\pi) + \Delta\ln(\omega^{NIT}_{g,r,t}) l^{NIT}_{g,r,t_0}(\pi) + \Delta\nu_{g,r,t}(\pi). \tag{A.4}$$

The equation above can be adapted to accommodate more than two types of jobs. The intuition is simple: growth in the marginal product within each job can be recovered by estimating the empirical counterpart of Equation A.4 where we consider each job against a combination of all other jobs. We specifically consider four possible jobs for workers to choose from: IT manufacturing, IT services, non-IT manufacturing and non-IT services. In the four jobs case, we can write a set of four equations indexed by industry $i \in \{man, serv\}$ and occupation $k \in \{IT, NIT\}$, as follows:

$$\Delta Y_{g,r,t}(\pi) = \Delta\ln(\omega^{k,i}_{g,r,t}) l^{k,i}_{g,r,t_0}(\pi) + \Delta\overline{\omega}^{k,i}_{g,r,t}(1 - l^{k,i}_{g,r,t_0}(\pi)) + \Delta\nu_{g,r,t}(\pi) \tag{A.5}$$

where $\overline{\omega}^{k,i}_{g,r,t}$ is a weighted average of the growths in logged marginal product of the other three jobs (that is, those that are not in occupation k and industry i). After rearranging, we obtain the following generalized estimation equation, characterizing the marginal product growth of the four job groups available:

$$\Delta Y_{g,r,t}(\pi) = \Delta\overline{\omega}^{k,i}_{g,r,t} + (\Delta\ln(\omega^{k,i}_{g,r,t}) - \Delta\overline{\omega}^{k,i}_{g,r,t}) l^{k,i}_{g,r,t_0}(\pi) + \Delta\nu_{g,r,t}(\pi) \tag{A.6}$$

Since there are two occupation groups k and two industries i, also in this case we have four estimation equations; our next section reports each equation that we use. This specification allows us to recover the growth rate of IT, relative to non-IT, tasks in each of the two industries, for all the triplets (g, r, t).

Appendix A.4.4. Generalizing the Task Prices: Industry and Task Groups

In order to extend the model to allow workers to also choose between industries, we begin with the equation 34 of Adao (2016), which enables us to recover the marginal product growth of two types of workers (IT and non-IT in our application):

$$\Delta Y_{g,r,t}(\pi) = \Delta\ln\omega^{IT}_{g,r,t} l^{IT}_{g,r,t_0}(\pi) + \Delta\ln\omega^{NIT}_{g,r,t} l^{NIT}_{g,r,t_0}(\pi) + \Delta\nu_{g,r,t}(\pi) \tag{A.7}$$

workers above the age of 50. These results somewhat mimic those by Autor et al. (2015), who find that the trade-employment elasticity in manufacturing is coming almost entirely from older workers.

[8] Specifically, let $\Delta\kappa^{IT}_{g,r,t}$ and $\Delta\kappa^{NIT}_{g,r,t}$ denote the changes in labor efficiency (composition effects) for IT and non-IT. We can quantify them by solving: $\Delta\kappa^{IT}_{g,r,t} = \Delta\ln(w^{IT}_{g,r,t})/\Delta\ln(\omega^{IT}_{g,r,t})$. To account for outliers and noise, we trim the distributions of both price and composition growth rates at the top and bottom percentiles.

[9] While we have experimented using other weights, we use the number of observations in each cell from 1980. If we, for example, allowed the weights to vary over time, we might confound and/or amplify the role of composition effects.

[10] Adao (2016) provides a derivation of the result from Equation 34 in his job market paper draft.

Now, notice that this equation can be reinterpreted to accommodate more than two types of workers. The intuition is the following: the marginal product growth of each kind of job can be recovered by estimating the equation above considering each job against a combination of the other jobs. In the four types of workers case, we can write the following four equations:

$$\Delta Y_{g,r,t}(\pi) = \Delta \ln\omega^{IT,man}_{g,r,t} l^{IT,man}_{g,r,t_0}(\pi) + \Delta\overline{\omega}^{IT,man}_{g,r,t}(1 - l^{IT,man}_{g,r,t_0}(\pi)) + \Delta\nu_{g,r,t}(\pi) \tag{A.8}$$

$$\Delta Y_{g,r,t}(\pi) = \Delta \ln\omega^{IT,serv}_{g,r,t} l^{IT,serv}_{g,r,t_0}(\pi) + \Delta\overline{\omega}^{IT,serv}_{g,r,t}(1 - l^{IT,serv}_{g,r,t_0}(\pi)) + \Delta\nu_{g,r,t}(\pi) \tag{A.9}$$

$$\Delta Y_{g,r,t}(\pi) = \Delta \ln\omega^{NIT,man}_{g,r,t} l^{NIT,man}_{g,r,t_0}(\pi) + \Delta\overline{\omega}^{NIT,man}_{g,r,t}(1 - l^{NIT,man}_{g,r,t_0}(\pi)) + \Delta\nu_{g,r,t}(\pi) \tag{A.10}$$

$$\Delta Y_{g,r,t}(\pi) = \Delta \ln\omega^{NIT,serv}_{g,r,t} l^{NIT,serv}_{g,r,t_0}(\pi) + \Delta\overline{\omega}^{NIT,serv}_{g,r,t}(1 - l^{NIT,serv}_{g,r,t_0}(\pi)) + \Delta\nu_{g,r,t}(\pi) \tag{A.11}$$

where $\overline{\omega}^{IT,man}_{g,r,t}$, for instance, is a weighted average of the marginal product growth of the three sectors that are not IT manufacturing. Rearranging the four equations above we get the following equations that enable us to estimate the marginal product growth of the four types of jobs that the workers are now allowed to choose:

$$\Delta Y_{g,r,t}(\pi) = \Delta\overline{\omega}^{IT,man}_{g,r,t} + (\Delta\ln\omega^{IT,man}_{g,r,t} - \Delta\overline{\omega}^{IT,man}_{g,r,t}) l^{IT,man}_{g,r,t_0}(\pi) + \Delta\nu_{g,r,t}(\pi) \tag{A.12}$$

$$\Delta Y_{g,r,t}(\pi) = \Delta\overline{\omega}^{IT,serv}_{g,r,t} + (\Delta\ln\omega^{IT,serv}_{g,r,t} - \Delta\overline{\omega}^{IT,serv}_{g,r,t}) l^{IT,serv}_{g,r,t_0}(\pi) + \Delta\nu_{g,r,t}(\pi) \tag{A.13}$$

$$\Delta Y_{g,r,t}(\pi) = \Delta\overline{\omega}^{NIT,man}_{g,r,t} + (\Delta\ln\omega^{NIT,man}_{g,r,t} - \Delta\overline{\omega}^{NIT,man}_{g,r,t}) l^{NIT,man}_{g,r,t_0}(\pi) + \Delta\nu_{g,r,t}(\pi) \tag{A.14}$$

$$\Delta Y_{g,r,t}(\pi) = \Delta\overline{\omega}^{NIT,serv}_{g,r,t} + (\Delta\ln\omega^{NIT,serv}_{g,r,t} - \Delta\overline{\omega}^{NIT,serv}_{g,r,t}) l^{NIT,serv}_{g,r,t_0}(\pi) + \Delta\nu_{g,r,t}(\pi) \tag{A.15}$$

Appendix A.5. Elasticity Estimates: Alternative IT Intensity Measure

In the paper we report estimates of the elasticities based on the baseline definition of IT intensity, which relies on the average z-score of IT intensity. An alternative modeling choice would be to define IT intensive occupations based on whether their IT score exceeds the median IT intensity. For ease of comparison Table A.4 reports elasticity estimates under both definitions of high IT intensity: the baseline (mean) and the median.

Estimates are not substantially different and suggest that the value we estimate are robust to reclassifying occupations in terms of their IT intensity.

	Quality-Adjusted Measures						Quality & Selection -Adjusted Measures			
	(M)	(S)	(M)	(S)	(M)	(S)	(M)	(S)	(M)	(S)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Baseline	1.72***	1.24***	1.28***	1.15***	1.31***	1.38***	1.63***	1.31***	1.77***	1.56***
	(0.138)	(0.049)	(0.066)	(0.045)	(0.115)	(0.116)	(0.056)	(0.042)	(0.092)	(0.086)
Median	1.38***	1.55***	1.24***	1.22***	1.27***	1.36***	1.68***	1.51***	1.72***	1.85***
	(0.083)	(0.077)	(0.060)	(0.051)	(0.097)	(0.112)	(0.051)	(0.058)	(0.074)	(0.118)
Time Diff.(Δ)	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Instrument	No	No	No	No	Yes	Yes	No	No	Yes	Yes
N	408	408	306	306	306	306	267	267	267	267

Table A.4: Comparison of Elasticity Estimates: Alternative Measure of IT Intensity

*Notes.–*Sources: Census Bureau, 1990-2015. The table reports the elasticity of substitution between information technology (IT) and non-IT workers in two ways. Columns 1-6 estimate the elasticity according to the following equation

$$\Delta \ln \left(\omega_i^{IT} L_i^{IT} / \omega_i^{NIT} L_i^{NIT} \right) = \Delta_t \ln \left(\alpha_i / (1 - \alpha_i) \right) + \nu_i \Delta_t \ln \left(L_i^{IT} / L_i^{NIT} \right) + \nu_i \Delta_t \left(\varepsilon_i^{IT} / \varepsilon_i^{NIT} \right)$$

where our outcome variable is the ratio of the wage bills and our right-hand-side variable is the ratio of effective labor services, which we estimate in log-levels and first-differences, using ordinary least squares and instrumental variables. We obtain quality-adjusted labor services premia by first estimating a group $\times$ year skill price obtained by regressing annual earnings on an indicator for working in an IT job, conditional on controls, and subsequently dividing total earnings by the skill price proxy and multiplying by an individual's total hours worked during the year. Controls in the skill price estimation include race, family size, marital status, gender, a quadratic in years of schooling, and age bin fixed effects (20-29, 30-39, 40-49, 50+). Controls in our estimation in the equilibrium condition include: group-specific average family size, share of males, race shares, a quadratic in schooling, and age bin shares. Columns 7-10 estimate the elasticity according to the following equation

$$\Delta \ln \left(\omega_i^{IT} / \omega_i^{NIT} \right) = \Delta_t \ln \left(\alpha_i / (1 - \alpha_i) \right) + (\nu_i - 1) \Delta_t \ln \left(L_i^{IT} / L_i^{NIT} \right) + \nu_i \Delta_t \left(\varepsilon_i^{IT} / \varepsilon_i^{NIT} \right)$$

where the difference from columns 1-6 is that our skill prices and quality-adjusted measure of labor services are obtained from our selection model based on Adao (2016). Controls in all specifications are the same and our instruments include the IT employment premia growth between 1970 and 1980 and the share of IT workers in 1970. Observations are weighted by the number of workers observed in each cell in 1980.The only difference between the "mean" and "median" rows are that the "median" elasticities are from measures of high IT defined by whether the individual is above the median IT intensity z-score, whereas the "mean" from our baseline is defined by whether the individual is above the mean (z-score of zero by construction).

Appendix B. References

Acemoglu, D., Autor, D. H., 2011. Skills, tasks and technologies: Implications for employment and earnings. Handbook of Labor Economics 4b, 1043–1171.

Adao, R., 2016. Worker heterogeneity, wage inequality, and international trade: Theory and evidence from Brazil. Working paper.

Autor, D. H., 2013. The "task approach" to labor markets: An overview. NBER working paper.

Autor, D. H., Dorn, D., Hanson, G. H., 2015. Untangling trade and technology: Evidence from local labour markets. Economic Journal 125, 621–646.

Autor, D. H., Levy, F., Murnane, R. J., 2003. The skill content of recent technological change: An empirical exploration. Quarterly Journal of Economics 118 (4), 1279–1333.

Beckhusen, J., 2016. Occupations in information technology. American Community Survey Reports.

Hecker, D., 2005. High-technology employment: A NAICS-based update. Monthly Labor Review.

Juhn, C., Murphy, K. M., Pierce, B., 1993. Wage inequality and the rise in returns to skill. Journal of Political Economy 101 (3), 410–442.

Kambourov, G., Manovskii, I., 2013. A cautionary note on using (March) CPS and PSID data to study worker mobility. Macroeconomic Dynamics 17 (1), 172–194.

Lemieux, T., 2006. Increasing residual wage inequality: Composition effects, noisy data, or rising demand for skill? American Economic Review 96 (3), 461–498.

Roy, A., 1951. Some thoughts on the distribution of earnings. Oxford Economic Papers 3, 135–146.

Schott, P., 2004. Across-product versus within-product specialization in international trade. Quarterly Journal of Economics 119 (2), 647–678.