



WP 17-30

Giovanni Gallipoli

University of British Columbia, Canada
The Rimini Centre for Economic Analysis

Christos A. Makridis

Stanford University, USA

STRUCTURAL TRANSFORMATION AND THE RISE OF INFORMATION TECHNOLOGY

This paper was part of the Carnegie-Rochester-NYU Conference Series on Public Policy (Carnegie Mellon University in November 2017).

Copyright belongs to the author. Short sections of the text, not exceeding three paragraphs, can be used provided proper acknowledgement is given.

The Rimini Centre for Economic Analysis (RCEA) was established in March 2007. RCEA is a private, nonprofit organization dedicated to independent research in Applied and Theoretical Economics and related fields. RCEA organizes seminars and workshops, sponsors a general interest journal *The Review of Economic Analysis*, and organizes a biennial conference: *The Rimini Conference in Economics and Finance* (RCEF). Scientific work contributed by the RCEA Scholars is published in the RCEA Working Papers and Professional Report series.

The views expressed in this paper are those of the authors. No responsibility for them should be attributed to The Rimini Centre for Economic Analysis.

Structural Transformation and the Rise of Information Technology

Giovanni Gallipoli and Christos A. Makridis*

November 30, 2017

PREPARED FOR THE CARNEGIE ROCHESTER AND NYU
CONFERENCE ON PUBLIC POLICY

Abstract

Has the emergence of information technology changed the structure of employment and earnings in the US? We propose a new index of occupation-level IT intensity and document several long-term changes in the occupational landscape over the past decades. Using Census micro-data between 1970 and 2015, we show that: (i) the share of workers in IT-intensive jobs has expanded significantly, with little or no pause; (ii) IT jobs enjoy a large and growing earnings premium, even after controlling for general task requirements (e.g., cognitive, non-routine); and (iii) the rise of the IT employment share is closely associated with declines in the manufacturing employment share. Although the earnings premia for college-educated and high cognitive/non-routine skilled workers have declined in the aggregate since 2000, we show that they have continued growing in IT jobs. We subsequently introduce an equilibrium model of occupational sorting based on comparative advantage between IT and non-IT jobs to quantify the contribution of IT jobs towards accelerating the pace of structural transformation. Our results suggest that technological growth among IT jobs has played a major role in accounting for the surge in high tech service labor productivity since 1980.

*Corresponding author: Christos Makridis, Department of Management Science & Engineering, Department of Economics, Huang Engineering Center, 475 Via Ortega, Stanford University, Stanford, CA 94305-4121, cmakridi@stanford.edu (<http://stanford.edu/~cmakridi/>). Giovanni Gallipoli: Vancouver School of Economics, University of British Columbia, gallipol@mail.ubc.ca. We thank Murilo Esteves de Santi, Elba Gomez Navas and Dante Zago for valuable research assistance. We are grateful to Ben Pugsley for many valuable comments on an earlier draft. We also thank David Autor, Raj Choudhury, Paul Gaggl, Robert Valletta, and participants at the Carnegie-Rochester-NYU Conference on Public Policy.

1 Introduction

Understanding large scale sectoral reallocation is essential for identifying the factors of technological change and growth (Baumol, 1967; Kuznets, 1973), such as the decline in the agricultural employment share (Gollin et al., 2002), the rise of the service economy (Buera and Kaboski, 2012), and the rise in job polarization (Barany and Siegel, 2017). Leading theories of structural transformation focus on either demand-side mechanisms, which introduce heterogeneity in income elasticities (Kongsamut et al., 2001), or supply-side mechanisms, which highlight heterogeneity in sectoral growth rates (Ngai and Pissarides, 2007; Acemoglu and Guerrieri, 2008).¹

Evidence suggests that, over the past few decades, differences in technological progress across sectors have been crucial for structural transformation in the United States (Herrendorf et al., 2015). Both labor productivity growth and the expansion in the employment share in services have continued at a relatively strong pace. Because the services sector has lower output per worker than manufacturing, models that posit wage equalization predict that increases in expenditures on labor-intensive services would eventually be reflected in lower aggregate productivity growth (Baumol, 1967).² In other words, given that the services sector has kept expanding both in terms of its employment and labor compensation shares, such models predict stagnation in labor productivity. This paper examines why this has not happened.

We show that the emergence of jobs that are intensive in information technology (IT) tasks may have helped avoid this long term stagnation; see Figure 1 for motivating evidence on the relationship between IT and the manufacturing employment share since 1970. Our argument is that technological discoveries, such as the introduction of personal computers in the 1970s and the spread of the internet in the 1990s, have led to an expansion of IT-intensive jobs. These jobs, however, only imperfectly substitute for non-IT jobs. Moreover, we show that IT-intensive jobs are closer substitutes for non IT-intensive jobs in the manufacturing sector, relative to the services sector.³ We argue that, while the rise of IT-intensive jobs accelerated the rate of structural

¹Combining both of these mechanisms has also received attention; see, for example, Buera and Kaboski (2009). Boppart (2014) and Comin et al. (2015) reconcile these mechanisms by modifying preferences.

²As Baumol wrote in his 1967 article, "...if wages and productivity in the progressive sector both go up 2 per cent per year, costs there will not rise at all. On the other hand, if in the non progressive sector productivity is constant, every rise in wages must yield a corresponding addition to costs—a two per cent cumulative rise in wages means that, year in year out, costs must be two per cent above those of the preceding year. Thus, the very progress of the technologically progressive sectors inevitably adds to the costs of the technologically unchanging sectors of the economy, unless somehow the labor markets in these areas can be sealed off and wages held absolutely constant, a most unlikely possibility."

³We formally estimate these elasticities later, but for anecdotal evidence, consider the instance where a mere

transformation in manufacturing, it did not result in an equal displacement of jobs in the non-IT-intensive services sector nor in equalization of the marginal product of different tasks. In this sense, the emergence of IT-intensive jobs may account for much of the continued rise in value added per worker of the past decades. This increased productivity is not fully reduced by shifts in employment, especially in the services sector.⁴ We also provide direct evidence that these productivity effects show up in sectoral real output and value added.

[INSERT FIGURE 1 HERE]

In the first part of the paper, we document three novel stylized facts. We begin by separately examining productivity growth in the manufacturing and services sectors and show that, while gross output per worker in manufacturing exceeds that in services, value added per worker is actually larger in services. Moreover, since the 1990s, most of the growth in value added per worker has come from high tech service industries, which exhibit nearly twice as large value added per worker on average than the average manufacturing industry.⁵

To shed light on the cause of these sectoral differences, we introduce a new measurement strategy to identify jobs with greater intensity of information technology and associated tasks. Using detailed six-digit occupation classifications from O*NET, we study the cross-section and time series evolution of IT-intensive jobs and their earnings premium. We combine the Decennial Census and the Current Population Survey (CPS) to document that the share of IT-intensive workers has grown from 34.4% to 43.7% between 1970 and 2015. Similarly, the raw earnings premium in these jobs has grown from 47.7% to 66.4% over the same period. These results are consistent with a broader literature about information technology and firm productivity (Bresnahan et al., 2002; Brynjolfsson and Hitt, 2003; Bloom et al., 2012; Tambe and Hitt, 2012). We also compare our results with recent evidence about the flattening of the college premium (Valletta, 2016) and the slowing demand for cognitive and non-routine (“C/NR”) skills (Beaudry et al., 2016). Despite the flattening of returns in the broader cross-section, there continues to be a steady rise in the education and C/NR premium among IT workers.

14 IT-intensive jobs in a manufacturing plant in Austria are enough to produce 500,000 tons of steel a year:

<https://www.bloomberg.com/news/articles/2017-06-21/how-just-14-people-make-500-000-tons-of-steel-a-year-in-austria>.

⁴These results are consistent with anecdotal evidence from, for example, a LinkedIn survey of employers in 2016 that found 22 out of the 25 skills that recruiters are most interested in are technology and information technology related; <https://blog.linkedin.com/2016/01/12/the-25-skills-that-can-get-you-hired-in-2016>

⁵We follow the approach suggested by Herrendorf et al. (2013) by focusing on value added through occupational tasks, rather than gross output.

We also exploit within-county variation between 1980 and 2013 and find that a percentage point (pp) rise in the IT employment share is associated with a non-trivial 0.35pp decline in the manufacturing employment share under our preferred specification. Our conditional correlations are identified from decadal within-county variation in employment shares, after controlling for demographic shifts and average wages. The results are robust to instrumenting for the contemporaneous IT employment share between 1990 and 2013.⁶ The fact that these conditional correlations hold even after controlling for average wages points towards the possibility that differences in productivity are a driver of structural transformation (Ngai and Pissarides, 2007).⁷

Of course, unobserved differences among individual workers in high and low IT jobs might confound our descriptive results. In the second part of our paper, therefore, we develop a model featuring task and worker heterogeneity based on the original approach of Adao (2016). Rather than controlling for worker heterogeneity, the model explicitly incorporates individuals' comparative advantage in a Roy-like setting. Under the assumption that workers have comparative advantages over IT and non-IT tasks, we estimate key model parameters and use these estimates to decompose the contribution of price variation, hours worked and workers' quality (composition) to structural transformation. This approach also allows us to quantify the relative substitution of IT and non-IT labor inputs separately for the manufacturing and the services sectors.⁸ In particular, we find that these labor inputs are less substitutable in the services sector, with an elasticity of 1.31 versus 1.62 in the manufacturing sector.⁹

The model allows us to derive estimable relationships linking price changes to quantity changes. We use these empirical relationships to construct counterfactual price changes and hours changes. In particular, we show that: (i) price adjustments play a critical role in mitigating shifts in the employment structure; (ii) quantity responses to price changes are more pronounced in the manufacturing sector, which explains the crowding out of non-IT jobs and the continuing decline of the manufacturing employment share; (iii) the services sector exhibits a more stable employment

⁶We also use the 1980 IT employment share as an instrument, which is pre-determined with respect to both contemporaneous outcomes and to the "China shock".

⁷Alternative theories that focus instead on income effects emphasize structural changes in per capita income.

⁸While traditional SBTC contributions have focused on college attainment as a proxy for skill (Katz and Murphy, 1992; Autor et al., 2006) and/or heterogeneity between non-routine & cognitively intensive occupations (Autor et al., 2008; Autor and Dorn, 2013), an important feature of the SBTC intellectual foundation is that it can accommodate alternative, relevant sources of heterogeneity, such as age (Card and Lemieux, 2001) and gender (Acemoglu et al., 2004). See Acemoglu and Autor (2011) for a survey.

⁹We test and reject the null hypothesis that these elasticities are the same, with a F -statistic of 33. Our estimates of the elasticities are robust to the use of different IV procedures to deal with the potential concern of reverse causality and endogeneity of aggregate labor inputs.

structure, as well as large and persistent productivity growth in IT jobs relative to non-IT jobs.

Our work also connects with two broad, recent debates about the role of technology for labor markets and productivity. The first such debate examines the decline in productivity growth over the past fifteen years. While some prominent researchers have argued that productivity might simply be mismeasured due to the difficulty of quantifying output in technology-intensive activities (Brynjolfsson and McAfee, 2011; Mokyr, 2014; Bryne et al., 2013; Feldstein, 2015; Hatzius and Dawsey, 2015; Bryne et al., 2016), Syverson (2017) suggests that mismeasurement is an unlikely explanation and Brynjolfsson et al. (2017) argue that new technologies (e.g., artificial intelligence) might simply have more delayed effects on the real economy since it takes time for the benefits to diffuse.¹⁰ Acemoglu et al. (2014) examine productivity growth in the manufacturing sector, finding that it was in fact rather low in IT-using *industries*. Rather than arguing that productivity growth is mismeasured and/or higher than measured in the manufacturing sector, our work simply points towards an important dimension of heterogeneity—namely, that significant productivity growth may be more recently coming from IT-intensive *jobs* in the services sector.

The second literature we contribute towards highlights growing concerns about automation and income inequality. For example, Harrigan et al. (2016) use administrative data from France to show that information technology and technological change have played a major role in facilitating job polarization.¹¹ However, they focus specifically on technical managers, engineers and technicians (the “techies”), rather than information technology workers as defined by occupational task descriptions, to study polarization in the distribution of jobs. In a similar vein, while some prominent experts have argued that a whopping 47% of U.S. employment is at risk of being displaced by automation (Frey and Osborne, 2013), Acemoglu and Restrepo (2017b) show that continuing automation is consistent with balanced growth, creating new tasks and raising the demand for heterogeneous skills. Even in the presence of short-run displacement of jobs (Acemoglu and Restrepo, 2017a), which clearly matter for welfare and distributional considerations (Eden and Gaggl, 2017), what matters most for structural transformation is the long-run demand for skills. In this sense, our paper contributes towards a better understanding of how information technology shapes the sorting of workers across sectors and the role of price effects in moderat-

¹⁰See an interesting discussion of related issues here: <https://www.wsj.com/articles/silicon-valley-doesnt-believe-u-s-productivity-is-down-1437100700>

¹¹Their contribution relates to findings in Autor et al. (2003) and Autor and Dorn (2013) in that information and communications technology largely complement cognitive and non-routine skills. Below we also provide evidence suggesting that this may be true, as we show that the earnings premium for college and cognitive/non-routine skills after 2000 has been increasing for information technology workers (but not for non-IT workers).

ing reallocation. Our findings are also consistent with those by Duernecker et al. (2017). These authors suggest that, while Baumol’s disease may have somewhat slowed aggregate GDP growth, its effects are incrementally smaller and will be limited in the future because most of the growth will continue to come from sustained demand for the output of high productivity services sectors. Our focus on the growing importance of IT jobs highlights a specific supply-side mechanism that drives higher productivity growth in the services sector. Identifying these occupational shifts will be important in further developing the current models (e.g., Buera and Kaboski (2012), Herrendorf et al. (2013), Boppart (2014), and Comin et al. (2015)) and characterizing the next stage of structural transformation—the move towards information services.

2 Data and Measurement

2.1 Sources

Occupational Task and Network.—O*NET (<http://www.onetcenter.org/questionnaires.html>) is the main source of occupation task, skill, and work environment characteristics. To construct our measure of information technology (IT) intensity, we select several measurements in their database, including: knowledge about computers and electronics, interacting with computers on the job, programming tasks, systems management and performance tasks, quality control analysis tasks, operations analysis tasks, updating and using relevant knowledge on the job, technology design, analyzing data and information on the job, processing information on the job, engineering and technology tasks, management of material resources on the job; see Appendix A.1.

These tasks are rated with an index at the six-digit occupation level over the intensity and frequency with which the skill, knowledge, or activity is applied on the job; we sum these indices together and produce an aggregate IT index, which we standardize to a mean of zero and standard deviation of one.¹² We also construct measures of skill intensity that have been commonly applied in the skill-biased technical change literature—including routine manual, non-routine manual/physical, routine cognitive, non-routine cognitive/analytical, and non-routine cognitive/interpersonal—and show that IT intensity predicts significant variation in hourly wages even after controlling for these standard skill measures.¹³

¹²While we implement several validation exercises, one simple approximation is to compare our highest IT intensity occupations with those from the Census Bureau’s list (see their ACS report “Occupations in information technology”). Appendix A.1 also provides a comparison.

¹³The broad occupation category that is closest to IT intensity jobs is that including non-routine cognitive

While O*NET allows us to produce a time-varying measure of intensity between 2004-2016, we primarily use a five-digit time-invariant indicator for high and low IT intensity occupations matched with Census micro-data between 1970 and 2015 (harmonized to 2010 SOC codes). We gauge our assumption that occupations that are IT-intensive from 2004 to 2016 are IT-intensive before 2004 by examining how IT intensity between 2004-2006 predicts IT intensity between 2014-2016. We find significant persistence: a coefficient of 1.069 (p -value = 0.00) and R -squared of 0.79. The fact that within-occupation IT intensity in the early 2000s predicts within-occupation IT intensity a decade later does not guarantee that the same would hold over the 1980s and 1990s, but we view the result simply as a diagnostic.¹⁴ We also experiment with an alternative index based on Hecker (2005) who classified NAICS codes into high-technology sectors.

Occupation and Employment Statistics.—The Bureau of Labor Statistics releases annual data at a 6-digit occupation level on employment and wages. The vast majority of these occupations can be matched with O*NET data on information technology intensity. However, after 2012 the level of granularity in the OES tables begins to surpass O*NET in some dimensions. There are 667 occupations between 2000 and 2003, 714 between 2004 and 2009, 734 between 2010 and 2011, and 769 between 2012 and 2015 that match over. We use these data as a validation source for a few exercises because of its granularity and comprehensive span.

Individual Micro-data.—We primarily draw from the and Census Bureau’s American Community Survey (ACS) and Decennial Census between 1970 and 2015 accessed through the Integrated Public Use Microdata (IPUMS) data portal at the University of Minnesota, as well as the annual Current Population Survey (CPS) for several robustness exercises that allow us to hone in on within-decade variation. To mitigate concerns about partial attachment to the labor market we restrict our samples to full-time workers between age 20 and 65, with over \$5,000 in annual labor income, at least 20 weeks worked per year, and over \$2 hourly wages. We deflate nominal variables using the 2010 real personal consumption expenditure index.

Industry Classifications.—We define the services sector based on the following industries: utilities, transportation and warehousing, wholesale and retail trade, information, finance and insurance, real estate, professional / scientific / technical services, management of companies and enterprises, administrative and support / waste and remediation services, educational services,

 occupations, which consist of thought processes requiring absorption and decision-making of abstract information. We refer readers to Aedo et al. (2013) for further details.

¹⁴Other approaches are emerging to classify jobs too. For one clever and resourceful example, see Atalay et al. (2017) who use textual analysis to recover measures of skill and task intensity between 1960 and 2000 to help explain the rise in earnings inequality.

healthcare and social assistance, arts / entertainment / recreation, accommodation and food services, and other personal services.¹⁵ We also consider several specifications where we further partition the services sector into high and low technology sub-sectors. The high tech set includes finance, insurance, real estate, and business services, while the wholesale / retail trade, utilities, transportation, and personal / other services are in the low-tech sector. We define the manufacturing sector based on NAICS 31-33 codes (non-durable and durable manufacturing) and include agriculture and mining.¹⁶

2.2 Measuring IT-Intensity across Jobs

One contribution of our analysis is the introduction of a task-based measure of the intensity of information technology (IT) tasks within a job.¹⁷ Since IT jobs are not a well-defined category, we pursue several strategies that appear to converge towards similar results as our baseline. Our primary approach uses several O*NET measurements—a combination of task, knowledge, and skill intensities—that are broadly related with technology and data, and subsequently takes an unweighted average of them to produce an index.¹⁸ Using the sum of these indices, we define jobs as high or low in information technology (IT) intensity if they are above or below the median index value, respectively.¹⁹ Using the aforementioned definition as the baseline, Figure 2 shows that there is a significant relationship between earnings and IT intensity at the five-digit occupation level: a standard deviation rise in IT intensity is associated with a 0.83% increase in annual earnings.

[INSERT FIGURE 2 HERE]

We also pursue two other strategies to help validate our baseline measure. First, we draw on the classification of high-technology occupations introduced by the Census Bureau (Hecker, 2005).

¹⁵<https://www.census.gov/econ/services.html>

¹⁶Our results are robust to including the education and healthcare sectors as low skilled in the services sectors and/or excluding agriculture and mining from the manufacturing sector. These latter sub-sectors make little difference because employment shares are quite small in agriculture and mining. Similarly, while education and healthcare clearly contain groups of highly skilled workers, many are not. Our classification on high and low tech services sectors is influenced by the Census (<https://www.bls.gov/opub/mlr/2005/07/art6full.pdf>), but we also have examined college attainment across these sectors.

¹⁷To put the benefit of our task-based measure into context, Acemoglu et al. (2014) remark that: “A second category of explanation for these unexpected results is that our measure of IT investment, constructed by averaging computer investment data from 1977–2007, misses the mark.”

¹⁸See Appendix A.1 for a list of the full questions used on each index.

¹⁹We have experimented with alternative clustering algorithms. Treating the unweighted sum index as the outcome feature, *K*-means and *K*-medians did not produce classifications that were significantly different from our median-based approach. We defer to the median-based approach for its transparency and stability.

We expand the definition to include jobs that are broadly related to technology, including, for example, media and communications specialists since 2000.²⁰ To verify robustness we implement the following exercises: we report the gradient of high IT jobs from regressions of logged earnings and logged hours worked on the high IT indicator, conditional on controls, for both of our IT-intensity measures. We recover a raw (conditional) earnings premium of 60% (36.9%) using our baseline classification, whereas the premium is 55% (31.2%) when we use the Hecker (2005) classification. We also recover a raw (conditional) hours premium of 9.96% (7.77%) and 9.24% (6.48%), for each IT-intensity measure respectively.²¹ The fact that the conditional correlations are similar suggests that our classification is broadly capturing comparable sets of workers, despite Hecker’s classification being only based on industry. Second, we draw on the Bureau of Economic Analysis (BEA) measures of IT capital expenditures from 2000 to 2015 and correlate them with the occupation-weighted industry measure of IT intensity implied by our baseline measure. Figure 3 plots logged IT capital and the IT employment share at the three-digit occupation level, producing a correlation of 0.51. While we would not expect the two variables to be perfectly correlated unless capital and labor IT are perfect complements, the positive correlation is reassuring.

[INSERT FIGURE 3 HERE]

2.3 IT-Intensity and Heterogeneity: Descriptive Statistics

We begin by documenting the set of characteristics that are associated with IT-intensity by looking at both jobs and workers over time and across space. These are documented in Table 1. There are stark gender and college attainment differences between workers in high and low IT-intensity jobs. Moreover, these differences have been changing considerably over time. For example, in 1970 the shares of males and college degree workers in IT-intensive jobs were, respectively, 76% and 33%. This can be contrasted to a share of only 59% males and 4% college grads in jobs with lower IT intensity. However, these differences narrowed: for 2010-2015 the shares of males and college grads were 54% and 58% in IT-intensive jobs, versus 52% and 18% in non-IT-intensive jobs.

We also see a large increase in the hourly wage premium, growing from \$6.68/hour in 1980 to \$12.55/hour in 2013 (an 87% increase). Controlling for college attainment reduces the premium to \$5.05/hour in 1980 and \$9.11/hour in 2010-2015 (a 65% increase), consistent with the com-

²⁰We present a formal definition in Appendix A.1 that tabulates all the occupation titles.

²¹We use family size, number of children, age, gender, race, and dummies for years of schooling as controls in producing the conditional correlations. We use survey sample weights in each specification.

plementarity between skills and computer technologies (Autor et al., 1998). However, the hourly wage premium still remains large. Interestingly, the increase in the hourly wage premium occurred concurrently with, and in spite of, an increase in the hours-worked IT premium. For example, high IT job workers allocated 8.3% more time to market activity in 1980 than workers in low IT jobs, but this gap grew to 11% by 2010-2015.²² Over this same time period, the IT earnings premium grew from 46.4% to 66.4%. The fact that the earnings premium is rising faster than the hours premium for IT workers suggests that demand is outpacing supply—a question we return to later in the context of our structural model. Appendix A.2 examines how the distribution of employment, hourly wages, and inequality varies across jobs with high and low IT intensity. While the distribution of the hourly wage for IT jobs is shifted significantly to the right, there is interestingly little statistical difference in within-group inequality in high IT and low IT jobs.

[INSERT TABLE 1 HERE]

We briefly explore differences in skill content across jobs with high and low IT intensity, displayed below in Figure 4 using data from the Occupation Employment Statistics (OES) between 2000 and 2015 at the six-digit level. There is a remarkable difference in the distribution of cognitive, technical, social, and general skills across high and low IT jobs. There are, however, no meaningful differences across manual and service skill intensities. Crucially, there continues to be a significant hourly wage IT premium even when controlling for differences in skill intensities at the six-digit occupation level, suggesting that IT intensity is not simply capturing heterogeneity in non-routine and cognitive skills. These exercises are documented in Appendix A.3.1. We also present descriptives in Appendix A.3.2 on the time series of IT intensity across jobs.

[INSERT FIGURE 4 HERE]

3 Structural Change and the Emergence of Technology and Information Services

²²We have also cross-walked our IT index into the American Time Use Survey (ATUS) and found that, using their more reliable diary-based time use reporting method, IT workers allocate roughly 21 more minutes per day to work activities, which totals nearly 128 hours worked more per year.

3.1 Heterogeneity in Industry Labor Productivity

How has productivity evolved over time in the manufacturing and services sectors? Although the manufacturing sector exhibits greater total factor productivity (TFP) (Herrendorf et al., 2015), the purpose of this section is to document the presence of significant heterogeneity in labor productivity (in constant dollars) within the services sector. In particular, we distinguish high-tech and low-tech sectors based on the Hecker (2005) Census classification, rather than introducing our IT intensity measure, to first establish a benchmark for our subsequent analysis. After documenting these productivity facts, we later apply our IT intensity indicator, which allows for far richer variation across the entire set of six-digit occupation classifications and avoids the possible mis-classification of employee sorting across tasks that occurs at an industry-level.

Using the U.S. KLEMS between 1947 and 2014, we begin by plotting average sectoral value added and gross output per worker in Figure 5, partitioning services into high and low technology sectors. Starting with value added per worker, we point out two observations. First, average sectoral value added per worker is nearly twice as large in high technology services sectors, relative to even the manufacturing sector. Second, the bulk of the increase in value added has been generated from high technology services sectors.²³ Although the level of average sectoral gross output per worker is marginally lower in the high tech services sector, relative to manufacturing, similar patterns hold over time.²⁴

We will later show that one of the major reasons for the difference in productivity between high and low tech services industries is their heterogeneous composition of IT-intensive jobs. In this sense, while there may be large components of the services sector that exhibit low productivity growth (e.g., trade and other services), other services industries have been growing their productivity remarkably fast over the past decades. These results are consistent with work by Duernecker et al. (2017) who argue that one should distinguish between high and low productivity services sectors when examining the potential role of Baumol’s cost disease.

²³It is important to net out the contribution of intermediate goods and normalize by employment. Since the intermediates share of output differs between sectors, netting out its contribution towards gross output helps us focus on variation in labor input productivity between the two sectors. As we discuss below, we use a value added approach to national accounting for structural transformation as described by Herrendorf et al. (2013).

²⁴We find similar results when using TFP, constructed as the residual of logged gross output regressed on logged employment, capital, and intermediates from USA KLEMS data. Our use of the term “average sectoral” refers to the fact that we average productivity per worker across each sub-sector weighting by their long term employment. While aggregate productivity per worker in manufacturing or services generates similar patterns, we do not want to confound changes in employment (e.g., the overall size of manufacturing and services) with other fundamental changes taking place within each sub-sector.

[INSERT FIGURE 5 HERE]

In what follows, we will focus on value added per worker as an important measure for understanding aggregate productivity growth. This approach follows from our emphasis on the role of different occupations in production and on the way broad tasks (occupation inputs) are aggregated in different industries.²⁵ As we will discuss in greater detail, by focusing on the categories of IT-intensive and non-IT-intensive jobs as task inputs we are able to highlight the relative contribution of these broad occupation inputs to value added after netting out the role of intermediate goods. As shown by Herrendorf et al. (2013), whether one uses value added or gross output would also influence the measurement of consumption in general equilibrium models of structural transformation. The distinction is especially important in our setting because IT jobs exhibit lower intermediate intensity, relative to manufacturing. If the expansion of IT jobs does, at least partly, affect the relative elasticity of final output to intermediates, then it is preferable to purge out intermediates from productivity measures.

3.2 Heterogeneity in Occupational Employment and Earnings

Having motivated our emphasis on high-tech services, we next focus on occupation heterogeneity within both the manufacturing and services sectors. Using micro-data between 1970 and 2015 from the Census Bureau, we begin our analysis by plotting the earnings and employment premia between high and low IT-intensive jobs. Figure 6 documents that the earnings premium for IT-intensive jobs in both manufacturing and services sector. Whereas the earnings premium grew from 50.4% to 63.6% in the manufacturing sector between 1970 and 2013 (a 26.2% rise), it grew from 48.3% to 67.2% in the services sector (a 40% rise). We find very similar patterns when working with hourly wages and residual hourly wages, which purges variation in demographic characteristics and educational attainment.

Turning towards differences in employment, Figure 6 also documents the employment premium for IT-intensive jobs. While manufacturing had only a few IT-intensive jobs in the 1970s, with an employment premium of -0.838, the share of such jobs grew dramatically over time to an employment premium of -0.233 (a 70% rise) by the end of our sample period. Over the same

²⁵If the primary metric was, for example, gross output, one might want to model more explicitly how capital stocks and expenditures on materials enter the production of final goods and services. Our empirical approach to estimating labor substitutability is consistent with alternative, and possibly heterogeneous, ways capital stock may affect added value across industries.

period the economy also experienced a significant growth in the share of IT-intensive jobs in the services sector, from an employment premium of -0.534 to -0.184 (a rise of 64.6%). These results suggest that, in the face of the well-known decline in the manufacturing employment share (Herrendorf et al., 2014), any job growth (or retention) in the manufacturing sector must have been primarily in IT-intensive jobs. Figure 20 in Appendix A.3.3 presents evidence consistent with this corollary: using the Current Population Survey (CPS), we show that the employment share of IT jobs in the manufacturing sector remained quite stable, despite the rapid decline of the overall manufacturing employment share.

[INSERT FIGURE 20 HERE]

In Appendix A.3.3 we replicate these earnings premia using a combination of the CPS and OES datasets to illustrate that these premia are comparable in other standard sources. While there are some differences in levels since the level of aggregation differs (three-digit occupation in CPS and six-digit in OES), the qualitative patterns are the same. We also present robustness exercises using a different classification of IT-intensive jobs, namely a more restrictive scheme introduced by the Census Bureau that focuses heavily on science and engineering jobs (Hecker, 2005). We have also estimated these premia controlling for demographic characteristics, such as age and education, and found that our patterns still stand strong. We also find similar results using hourly wages, but defer to annual earnings since hours worked are not reported in 1970.

We now examine how these patterns relate with recent results from Beaudry et al. (2014; 2016) about the decline in demand for non-routine and cognitive (“C/NR”) skills and from Valletta (2016) about the flattening of the college premium.²⁶ In particular, has the flattening taken place for all workers, or is the flattening driven by non-IT jobs? After replicating their aggregate results using the Census, we find evidence of the latter. Beginning with educational attainment, Panel A in Figure 7 plots the overall college premium (red) with the college premium for IT workers (blue). Importantly, while all are larger and increasing, the premium is increasing especially for IT workers. For example, between 1970 and 1990, the college premium grew by 9% overall, whereas it grew by 16% for IT workers. Between 1990 and 2015, the college premium grew by 19% overall, but it grew by a remarkable 26% for IT workers. During these periods, the supply of IT workers

²⁶Beaudry et al. (2016) argue that, because of the housing boom in the early 2000s, the rise of employment in the construction and other complementary sectors masked these downward employment trends (Charles et al., 2016). It was not until the Great Recession that these labor market changes became apparent. Autor (2017) points out, however, that the decline in the earnings premium actually began in the 1990s, which is even more of a puzzle since IT investment was still large at the time.

was also rapidly increasing, suggesting that the supply of labor was not keeping up with the increasing demand for IT skills in the labor market (see Appendix A.3.4 and Figure 25 for these employment shares).

Turning towards C/NR skills, Panel B in Figure 7 plots the overall C/NR premium (red) with the C/NR premium for IT workers (blue). Unlike the college premium, there is a closer association between the overall C/NR and IT-C/NR premia since the majority of occupations considered high C/NR are also high IT. Despite similarity between the two premia in 1970 bordering around 40%, the premia began diverging in the 1990s. In particular, the IT premium among C/NR jobs begins outpacing the overall premium by over 5%. For example, while the overall C/NR premium grew by 14% between 1970 and 1990, it grew by 18% among IT workers. Moreover, the overall C/NR grew by 40% between 1990 and 2015, whereas it grew by 45% among IT workers.

[INSERT FIGURE 7 HERE]

While these plots may appear inconsistent with the evidence in Beaudry et al. (2016) and Valletta (2016) on the declining return to C/NR and college skills post-2000, they are not—they merely highlight an important source of heterogeneity. Turning away from the longer-run phenomena captured in Figure 7, we now introduce annual data from the American Community Survey between 2005 and 2015 to study the returns to IT jobs among the set of college degree workers and high C/NR jobs in more recent years. Using these data, we estimate regressions of logged annual earnings on an indicator for high IT, college attainment (and separately high C/NR skills), and their interaction, conditional on controls, for each year. Figure 8 plots the estimated interaction between IT and college attainment in Panel A and IT and C/NR skills in Panel B. While the levels of the two interactions vary, the overall pattern points towards an unambiguous rise in the premium associated with college and C/NR skills for high IT workers. In this sense, although the average return to college and C/NR skills might be declining overall as Beaudry et al. (2016) and Valletta (2016) point out, they are increasing for IT-intensive jobs.

[INSERT FIGURE 8 HERE]

3.3 IT-Intensive Jobs and Structural Change: Some Evidence

Motivated by these results about the dramatic change in the labor market following the emergence of IT-intensive jobs, we next provide some reduced-form evidence about their association with structural transformation. In particular, we ask whether increases in the employment share of IT

have contributed to the hollowing out of manufacturing employment, after accounting for a variety of confounding factors. Using our decadal Census micro-data to produce county weighted averages between 1980 and 2015, we restrict the sample to counties with over 2,000 survey respondents and run regressions of the form:

$$m_{ct} = \beta X_{ct} + \gamma IT_{ct} + \psi_c + \lambda_t + \epsilon_{ct} \quad (1)$$

where subscripts c and t denote county and decade, m denotes the manufacturing employment share, X denotes a vector of local controlling covariates, IT denotes the IT employment share, and ψ and λ are county and decade fixed effects. A possible endogeneity problem in estimating Equation 1 arises from the fact that higher productivity workers might sort into IT jobs (and thus locations with more of these jobs). To the extent that IT jobs are comparatively less likely to be concentrated in the manufacturing sector, γ may be biased downwards.

While we do not take a strong causal stance on γ , we provide evidence that this relationship is not spurious or entirely contaminated by endogeneity problems. In particular, we begin by introducing several controls that help reduce concerns about omitted variables bias, including: the logged hourly wage, the share of college degree workers, the share of whites, and the share of males. The inclusion of both wages and college shares helps ensure that we are not comparing areas with systematically different labor markets. The share of males addresses the concern of heterogeneous entry among females into different labor markets. The inclusion of county and year fixed effects removes all time-invariant differences across these labor markets, such as exposure to trade shocks or local human capital networks. To address the potential endogeneity arising from reverse causality and unobserved time-varying shocks, we also instrument the contemporaneous IT employment share with its 1980 value, which is pre-determined for $t > 1980$.

Table 2 documents our results. Column 1 presents the simple unconditional correlation, which suggests that a percentage point rise in the IT share is associated with a 0.67 percentage point decline in the manufacturing share. We subsequently introduce several controls in column 2, lowering the point estimate on the IT share to -0.226. Here, we find that the college share is negatively correlated with the manufacturing share, which is consistent with existing evidence on growing polarization across industries and the skill premium (Autor and Dorn, 2013). However, while these controls help address contemporaneous omitted variables, any time-invariant source of spatial heterogeneity may create attenuation bias. Column 3 introduces county and year fixed effects without controls, which suggests that a percentage point rise in the IT share is associated

with a 0.38pp decline in the manufacturing employment share. Recognizing the presence of time-varying unobservables, we once again add our controls, on top of county and year fixed effects, and find that the gradient only declines marginally in magnitude to an associated 0.35pp.

We finally examine the potential concern of reverse causality by instrumenting the contemporaneous IT-intensive share between 1990 and 2015 with the historical 1980 IT-intensive share. Our identifying variation emerges from the fact that counties with greater exposure to IT-intensive jobs in 1980 are more likely to experience an expansion during the 1990s technology boom. At the very least, the robustness of our estimates shows that contemporaneous factors, like China’s entry onto the global stage (Autor et al., 2013), are unlikely to reverse the main result.

[INSERT TABLE 2 HERE]

4 A Model of Production with Heterogeneous Tasks

In order to explicitly quantify the contribution of different factors to the emergence of IT-intensive services sectors, we develop a model based on the original and novel contribution of Adao (2016). The model formalizes an assignment problem of tasks and technologies in the tradition of Autor (2013).²⁷ The model contains 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).²⁸ 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, and (iii) perform counterfactual exercises using the model’s structure.

4.1 Task Production

Following Adao’s approach, we partition jobs into groups based on geography (state) and college attainment, denoted as $g \in \mathfrak{S} = \{1, 2, \dots, G\}$, and we allow for mobility between the manufactur-

²⁷We adapt Adao’s original structure to our setting, deferring readers to Adao (2016) for derivations and a more comprehensive discussion of analytical results.

²⁸Under the original approach, residual wage variation might simply reflect measurement error or unpriced amenities (Lemieux, 2006).

ing and services sector. We assume that individuals enter the labor market as either college or non-college educated workers.²⁹ We view states as small open economies with segmented labor markets.³⁰ 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.³¹ To accommodate the variety of occupations that workers perform, we assume that each aggregate occupation group $k \in \{IT, NIT\}$ is a collection of multiple (perfectly competitive) “single-task occupations” $j \in \mathfrak{S}^k$. 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 p^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:

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

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

²⁹We do not model the option of returning to school for continuing education.

³⁰One 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.

³¹In 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.

$s_g(i) = \ln[L_g^{IT}(i)/L_g^{NIT}(i)]$, 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 $\{a_g(i)|s_g(i) = s\} \sim \tilde{F}_g(a|s)$.

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

$$w_g^k = p^j \frac{\partial Q^j}{\partial L_g^j} \quad \text{if } j \in \mathfrak{S}^k \quad (2)$$

where w_g^k is the marginal product of workers in occupation class k and observable group g . Later in the paper 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.³² As in Roy (1951) model, 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:

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

where $\omega_g^k = \ln(w_g^k)$ and 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 $\tilde{F}_g(a|\sigma_g(q))$, with an average $\alpha_g(q)$ and variance $v_g(q)$. It follows that the logged wage schedule along the quantile range is :

³²One 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.

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

As shown in Adao (2016), individuals sort into the IT sector if $\sigma_g(q) > \omega_g^{NIT} - \omega_g^{IT}$, 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. $\omega_g^{NIT} - \omega_g^{IT}$, such that $\omega_g^{NIT} - \omega_g^{IT} = \sigma_g(l_g^{NIT})$ 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:

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

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

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

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 3 identifies $\Delta \omega_{g,r,t}^{NIT}$ and $\Delta \omega_{g,r,t}^{IT}$ 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, $\Delta \omega_{g,r,t}^{IT}$ and $\Delta \omega_{g,r,t}^{NIT}$.

The intuition behind Equation 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.³³

Having estimated $\Delta\omega_{g,r,t}^{IT}$ and $\Delta\omega_{g,r,t}^{NIT}$, one can recover the relative changes in workers' composition by taking the difference between raw wages growth and estimated task prices growth.³⁴ 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.³⁵

³³One 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 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.

³⁴Specifically, let $\Delta\kappa_{g,r,t}^{IT}$ and $\Delta\kappa_{g,r,t}^{NIT}$ denote the changes in labor efficiency (composition effects) for IT and non-IT. We can quantify them by solving: $\kappa_{g,r,t}^{IT} = \Delta w_{g,r,t}^{IT} / \Delta\omega_{g,r,t}^{IT}$. To account for outliers and noise, we trim the distributions of both price and composition growth rates at the top and bottom percentiles.

³⁵While 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

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:³⁶

$$\Delta Y_{g,r,t}(\pi) = \Delta \omega_{g,r,t}^{IT} l_{g,r,t_0}^{IT}(\pi) + \Delta \omega_{g,r,t}^{NIT} l_{g,r,t_0}^{NIT}(\pi) + \Delta \nu_{g,r,t}(\pi). \quad (4)$$

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 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:

$$\Delta Y_{g,r,t}(\pi) = \Delta \omega_{g,r,t}^{k,i} l_{g,r,t_0}^{k,i}(\pi) + \Delta \bar{\omega}_{g,r,t}^{k,i} (1 - l_{g,r,t_0}^{k,i}(\pi)) + \Delta \nu_{g,r,t}(\pi) \quad (5)$$

where $\bar{\omega}_{g,r,t}^{k,i}$ is a weighted average of the growths in marginal product for 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:

$$\Delta Y_{g,r,t}(\pi) = \Delta \bar{\omega}_{g,r,t}^{k,i} + (\Delta \omega_{g,r,t}^{k,i} - \Delta \bar{\omega}_{g,r,t}^{k,i}) l_{g,r,t_0}^{k,i}(\pi) + \Delta \nu_{g,r,t}(\pi) \quad (6)$$

Since there are two occupation groups k and two industries i , also in this case we have four estimation equations.³⁷ 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) .

effects.

³⁶Adao (2016) provides a derivation of the result from Equation 34 in his job market paper draft.

³⁷In Appendix A.4.1, we report each individual equation used for estimation.

4.4 Results by Industry and Demographics

We begin by plotting the distributions of growth rates for marginal products and composition effects in Figure 9. We do this for high and low IT jobs separately by industry, pooling across college \times year \times state. Several interesting facts become apparent from this analysis. First, mean price growth is larger in IT jobs, relative to their non-IT counterparts. For example, IT price growth is 5% on average in the manufacturing sector and 3% in the services sector, whereas non-IT prices are actually declining by 9% and 3% in manufacturing and services, respectively. The fact that non-IT prices are declining in both manufacturing and services is consistent with a relatively increasing demand for information services and IT tasks—an observation confirmed when we estimate the aggregate production technology below. Moreover, the fact that price growth is mostly negative for non-IT jobs in the manufacturing sector partly reflects an ongoing employment shift as workers are reallocated away from non-IT jobs and replaced with increased automation. Figure 27 in Appendix A.4.2 documents additional descriptive patterns for price growth in manufacturing and services across each decade; price growth was largest in both sectors between 1990 and 2000, consistent with the timing of the technology boom.

A natural implication of greater price growth among IT relative to non-IT jobs is that their average worker quality composition may be slower or even negative. Although lower quality composition growth may at first appear counter-intuitive, the relationship follows from our model setup and estimating equations. Intuitively, the employees transitioning from non-IT jobs into IT jobs are likely lower quality (lower comparative advantage). In this sense, through the cleansing effect usually associated to large sectoral shocks (Caballero and Hammour, 1994), reallocation leaves the higher quality workers in the non-IT jobs.

Second, IT price growth is slightly larger in the manufacturing sector than in services (by two percentage points), whereas the non-IT price growth is six percentage points lower in the manufacturing sector. While the difference is not massive, this differential growth cumulatively adds up over several decades. The fact that IT price growth is larger in the manufacturing sector, whereas non-IT price growth is lower, reflects the exodus of non-IT workers in the manufacturing sector as they get replaced by automation (Autor, 2014). However, as we will show shortly, IT workers are a weaker substitute for non-IT workers in the services sector, consistent with the discussion in Autor (2015). It is not surprising, then, that differences exist also in the change of workers' quality composition: average efficiency growth among non-IT workers in the manufacturing sector is much larger than that in the services sector (by nine percentage points) as the marginal workers

who moved to the services sector had less of a comparative advantage in that sector.

[INSERT FIGURE 9 HERE]

4.5 Results Across Space

We now turn to characterizing spatial heterogeneity in the IT price premium across the manufacturing and services sectors over the three time intervals of 1980-1990, 1990-2000, and 2000-2015 (pooling together college and non-college workers). That is, we consider the IT price growth net of non-IT price growth, and we plot it across states for each decade in Figure 10. Some important observations stand out. First, the manufacturing sector experienced greater IT price premium growth throughout every decade. This is due to the differential growth of IT and non IT task prices across industries, that we highlighted above: IT price growth is 5% on average in the manufacturing sector and 3% in the services sector, while non-IT prices are actually declining by 9% and 3% in manufacturing and services, respectively. This means that the growth gap between IT and non-IT task prices is much larger in the manufacturing sector, as visible in the heat maps.

Consistent with recent evidence on the increasing role of automation that crowds out routine and manual jobs (Autor et al., 2003; Autor and Dorn, 2013), nearly all the jobs that have been lost in manufacturing are in non-IT occupations. This phenomenon has been accompanied by a significant increase in the return to IT tasks, especially large in manufacturing.

Second, an acceleration in the relative IT price growth occurred during the technology boom of the 1990s. California, Massachusetts, and Texas were big beneficiaries of the IT price run-up in both manufacturing and services during that decade. The relative IT price growth in manufacturing has been strong also between 2000 and 2013, including also various locations in the South. This was accompanied by a progressive crowding out of non-IT workers from the combination of automation and IT workers employment. We show below that these results are consistent with our estimates of a higher elasticity of substitution between IT and non-IT workers in manufacturing, which means that workers quantities adjust more flexibly to technology shocks that raise the relative productivity of IT tasks.

In contrast, the slowdown, or outright decline, in the service sector IT price premium between 2000 to 2013 is indication that states where the IT revolution took hold earlier (e.g. California and Washington state) may be going through a catch up phase in terms of relative employment in IT task for services. There is evidence that the labor market did respond to the 1990s technology

boom by providing incentives for greater IT training among workers (e.g., through targeted college attainment as discussed by Acemoglu and Autor (2011)). This, in turn, raised the supply of workers who could perform IT tasks. In this sense, even though we see that the hourly wage premium has continued growing among IT workers, the growth in the IT tasks marginal product (relative to non-IT jobs) has slowed down and, in some cases, even reversed. Crucially, we show later that the elasticity of substitution between IT and non-IT tasks is significantly smaller in the services sector, which implies a faster cooling down of the IT price premium growth in response to growing IT-intensive employment.

We also examine the cross-sectional dispersion of the IT price growth premia by considering a potential driving force. One possibility is that labor market distortions in some locations restrict mobility and returns to skill accumulation, driving down the demand along the lines of Ljungqvist and Sargent (1998) and Ljungqvist and Sargent (2008). More rigid labor markets can create inertia and slow down the process of adapting to technological shocks and to production side changes. To examine the potential role of these labor market distortions, we use the Current Population Survey (CPS) partitioned into year groups for 1990, 2000, 2005-2007, and 2013-2015 to produce measures of union density and logged employment by industry \times college \times state cells.

We regress the IT price premium growth, in manufacturing and services, on these variables, conditional on demographic controls (e.g., race and age). We find that a 10% rise in union density is associated with a 5.3% acceleration (p -value = 0.087) in the growth of the IT price premium in the services sector, but a 1.4% decline (p -value = 0.373) in the manufacturing sector. One interpretation of these results is that the strengthening of unions in the services sector depresses the demand for non-IT jobs where unionization is greater, thereby contributing to a faster rise in the relative price for IT skills. However, since unions are relatively stronger in the manufacturing sector and exist relatively broadly in both IT and non-IT jobs, there is a less significant effect on the growth of the manufacturing IT premium. While these results are purely descriptive, they underscore the potential role of labor market institutions in either amplifying or restricting the expansion of the IT price premium.

[INSERT FIGURE 10 HERE]

4.6 Accounting for Structural Transformation

We can use our estimates of the growth in the marginal product (shadow price) of different tasks, and their average labor efficiency, to account for the changing factor shares of IT and non-IT labor services in the manufacturing and services sectors. We do so by imposing restrictions on the production structure and explicitly model the way marginal prices are determined. Crucially, in what follows we assume that individuals can choose both industry and task group.

We posit that output (value added) in industry i , denoted Φ_{it} , is generated using aggregate IT and non-IT efficiency-weighted labor services, denoted L_{it}^{IT} and L_{it}^{NIT} :

$$\Phi_{it} = \left[\alpha_{it} \left(L_{it}^{IT} z_{it}^{IT} \right)^{\nu_i} + (1 - \alpha_{it}) \left(L_{it}^{NIT} z_{it}^{NIT} \right)^{\nu_i} \right]^{\frac{1}{\nu_i}} \quad (7)$$

where z denotes task-specific productivity shocks. The random variables $z_{it}^j = \exp(\varepsilon_{it}^j)$ are independent across industries i and task groups j , with $\varepsilon_{it}^j \sim iid(0, \sigma_{ij}^2)$.³⁸ Equation 7 captures the fact that there is imperfect substitutability of IT and non-IT labor services, and that this substitutability varies across industries. The efficiency-weighted inputs are the product of total hours worked (raw hours) and efficiency per hour worked (quality):

$$L_{it}^{IT} = H_{it}^{IT} E_{it}^{IT}, \quad L_{it}^{NIT} = H_{it}^{NIT} L_{it}^{NIT} \quad (8)$$

where α_{it} denotes the factor share of IT-intensive labor services of workers in industry i , H denotes their work hours, and E denotes their average efficiency per hour worked. Taking derivatives, the marginal products of *effective* labor services for each industry i correspond to the shadow prices w^{IT} and w^{NIT} defined in equation (2). That is, the marginal products of IT and non-IT-intensive labor services can be written as,

$$MP_{it}^{IT} = w_{it}^k = \frac{\partial \Phi_{it}}{\partial L_{it}^{IT}} = \Phi_{it}^{1-\nu_i} \left[\alpha_{it} (L_{it}^{IT})^{\nu_i-1} (z_{it}^{IT})^{\nu_i} \right]$$

³⁸We purposefully abstract from modeling capital for various reasons. First, our empirical analysis focuses on the ratio of marginal products of IT and non-IT labor, thereby allowing us to partial out capital. In this sense our empirical approach implicitly accommodates alternative specifications of capital across industries and education groups (including a simple Cobb-Douglas technology in capital and aggregate labor). Second, our primary object of interest is added value produced using IT and non-IT labor inputs. We can capture the degree of substitutability of labor inputs in added value across industries and education groups without making strong assumptions about the role of capital. Finally, modeling capital is complex and beyond the scope of our analysis. In fact, we know of no obvious way of measuring and linking IT-specific capital and IT-intensive labor inputs. However, if capital were more complementary with IT tasks, our results would be strengthened.

$$MP_{it}^{NIT} = w_{it}^{NIT} = \frac{\partial \Phi_{it}}{\partial L_{it}^{NIT}} = \Phi_{it}^{1-\nu_i} \left[(1 - \alpha_{it}) (L_{it}^{NIT})^{\nu_i-1} (z_{it}^{NIT})^{\nu_i} \right]$$

The ratio of these marginal products implies the following equilibrium condition for each industry i :

$$\frac{w_{it}^{IT}}{w_{it}^{NIT}} = \left(\frac{\alpha_{it}}{1 - \alpha_{it}} \right) \left(\frac{L_{it}^{IT}}{L_{it}^{NIT}} \right)^{\nu_i-1} \left(\frac{z_{it}^{IT}}{z_{it}^{NIT}} \right)^{\nu_i} \quad (9)$$

Taking logarithms $\omega_g^k = \ln(w_g^k)$ and first-differencing over time produces the following equation:

$$\Delta_t (\omega_i^{IT} - \omega_i^{NIT}) = \Delta_t \ln \left(\frac{\alpha_i}{1 - \alpha_i} \right) + (\nu_i - 1) \Delta_t \ln \left(\frac{L_i^{IT}}{L_i^{NIT}} \right) + \nu_i \Delta_t \left(\frac{\varepsilon_i^{IT}}{\varepsilon_i^{NIT}} \right) \quad (10)$$

We can then recover the time varying factor shares, $\Delta_t \ln[\alpha_i/(1 - \alpha_i)]$, and the aggregate elasticity of substitution, $1/(1 - \nu_i)$, using the variation in marginal products, hours and efficiency that we have previously estimated using the Adao approach. The estimated growth of shadow price premia, $\Delta_t (\omega_i^{IT} - \omega_i^{NIT})$, vary by location and group.³⁹ The estimation delivers the average growth, $\Delta \ln[\alpha_{it}/(1 - \alpha_{it})]$, between 1980 and 2013, which is needed to quantify the long term changes; we call this ratio the drift since it describes what would have occurred with $\Delta_t \ln(L_i^{IT}/L_i^{NIT}) = 0$.⁴⁰

Estimation results. We next present our estimated elasticities from Equation 10.⁴¹ Having estimated $\nu - 1$, we recover the elasticity of substitution $1/(1 - \nu)$. We find that manufacturing exhibits a larger elasticity of substitution (1.63) relative to the services sector (1.31). This difference is statistically significant (p -value well below 1%) and is consistent with the more pronounced crowding out of non-IT jobs by IT jobs in the manufacturing sector over the past 40 years. To address the concern that our elasticity estimates may suffer from endogeneity of the quality-adjusted labor services premia, we experiment with instrumenting for them using two pre-determined variables: the historical 1970 share of IT workers and the growth rate between 1970 and 1980 in the share of IT workers. Using this approach we find, again, that the elasticity of substitution in the manufacturing sector is larger (1.68) than that in the services sector (1.39). These estimates are

³⁹An alternative approach would have been to collapse all observations across different locations into a unique growth rate. However this would have resulted in an unnecessary loss of information.

⁴⁰We have also estimated specifications in which we gauge how the factor shares change over individual decades. This information does not add anything to our long term analysis and is available. Elasticities are not sensitive to allowing factor shares to change at different rates over time.

⁴¹We use a simultaneous equation estimator that allows us to (i) flexibly combine information for different decades, (ii) experiment with and without restricting parameters to be the same across decades, (iii) control for endogeneity through first-stage IV regressions, when necessary.

summarized in Table 3.⁴²

Assessing model behavior. Using our baseline estimates we implement some simple diagnostic exercises to assess the model’s fit and behavior. First, we gauge the quality of model predictions by comparing actual and predicted changes in human capital premium between IT and non-IT jobs, separately by decade and industry. These comparisons are displayed in Figure 11. The model clearly generates much of the variation in human capital inputs, accounting for 80% or more of the observed variation across decade-industry pairs. Second, we relate our estimates of task productivity growth to proxies of aggregate productivity. A caveat is in order: our model focuses on relative task price changes, hence it is not designed to examine the properties of aggregate multi-factor productivity (TFP). Nonetheless, we contrast estimates of TFP⁴³ with model-based estimates of exogenous changes in IT tasks relative productivity growth. This allows us to verify whether relatively larger growth in IT task productivity is in any way associated with larger industry-level residual TFP. Crucially, our theoretical framework does not impose any restrictions on the size and direction of these relationships.

We measure changes in relative IT-productivity by projecting the growth in estimated IT task price premia on the growth in IT task input premia, using a specification with no intercept.⁴⁴ The omission of the constant implies that, by design, the residuals estimated from the latter specification include also changes in the ‘technological drift’ term, $\Delta \ln[\alpha_{it}/(1-\alpha_{it})]$. Our objectives are to establish how these residuals correlate with TFP estimates from BEA data and to assess the extent to which technological drift in the returns to IT tasks co-moves with standard proxies of (residual estimate) TFP. We examine the correlation between aggregate TFP proxies and the IT productivity drift in, respectively, manufacturing and services. As discussed in Appendix A.5,

⁴²We choose the OLS results to serve as our baseline estimates. We report the IV results to highlight that possible reverse causality does not alter our key findings. Estimated elasticities are similar and the drift in $\alpha/(1-\alpha)$ is larger in the services sector than in the manufacturing sector, as was the case in our OLS analysis.

⁴³Given the lack of reliable state-by-industry data about capital stocks, we measure TFP using two alternative proxies of capital input. First, we use capital income at the level of state \times two-digit industry. Second, we simply use data on aggregate capital and equipment. In either case we regress logged real output on logged employment, two-digit industry fixed effects and the logged proxy for capital input. We then take the residuals of these regressions as our aggregate total factor productivity (TFP) measures.

⁴⁴Specifically, we estimate the following regression for each industry:

$$\Delta_t (\omega_i^{IT} - \omega_i^{NIT}) = (\nu_i - 1)\Delta_t \ln(L_i^{IT}/L_i^{NIT}) + \underbrace{\Delta_t \ln(\alpha_i/(1-\alpha_i)) + \nu_i \Delta_t (\varepsilon_i^{IT}/\varepsilon_i^{NIT})}_{\text{changes in relative IT task productivity}}.$$

The variation due to $\Delta_t (\varepsilon_i^{IT}/\varepsilon_i^{NIT})$ is, by definition, iid in any given period and industry.

each observation is at the state-year level and our findings suggest positive correlation between residual industry-specific TFP and the growth in relative IT tasks productivity (see Figure 29). This positive correlation indicates that in state-years in which residual TFP is relatively larger we also observe relatively faster growth in IT task productivity, relative to non-IT tasks.⁴⁵

[INSERT TABLE 3 HERE]

[INSERT FIGURE 11 HERE]

4.7 Counterfactual Exercises

Using our estimates of the elasticities of substitution between IT and non-IT labor inputs, together with our growth rates in shadow prices and labor quality, we are able to implement several counterfactual exercises. First, we examine how much of the growth in the IT job shadow price premium (relative to the non-IT job shadow price) can be explained by: (i) labor quality composition effects, (ii) changes in aggregate hours worked, and (iii) variation in factor shares $\Delta_t \ln \left(\frac{\alpha_i}{1-\alpha_i} \right)$. To gauge the relative importance of each of these channels, we conduct counterfactuals where we assume that there is no growth in composition effects, hours, or factor shares. Under each scenario, we then compute the implied growth in the IT shadow price premium, by industry. Once we have the counterfactual price growth rates, together with an initial baseline 1980 estimate of the IT price premium level, we can extrapolate the IT price premium forward to the year 2013.⁴⁶

Table 4 documents these results.⁴⁷ To see how we compute the implied role of labor quality composition, hours, and factor share growth, consider the following example for the manufacturing sector. Counterfactual growth in the IT price premium, absent any composition premium growth, was 1.16% for 1980-1990, 6.18% for 1990-2000, and 8.85% for 2000-2013. Fixing the level of the price premium to its approximated level in 1980, which was 9.2%, these growth rates compound to

⁴⁵As mentioned before, our simple model is silent about the direction of this co-movement. Hence this correlation cannot be interpreted in a causal way. There is no theoretical implication in our model that would suggest a particular direction of causality between IT productivity growth and aggregate TFP.

⁴⁶Because we do not have a direct measure of the 1980 IT job price premium (in levels), we approximate it by taking the residual from a regression of the logged hourly wage on a vector of demographics and education fixed effects. The approximated IT jobs price premium (relative to non-IT jobs) in 1980 is 9.2% for the manufacturing sector and 8.9% for the services sector. These premia can be contrasted with, respectively, 48.5% and 49% premia in hourly wages in those two sectors in the same year. The fact that demographic controls, specifically education attainment, help bring down the magnitude of the premium provides us with some confidence that our approximation is at least getting closer to the level of the IT price premia. One might experiment with an alternative price normalization in 1980. While this would affect the level of price premia in each decade, their proportional changes over time would, by construction, be identical.

⁴⁷The percent contributions do not sum up to one since our estimated models do not have an R -squared of one.

a value of the IT price premium in 2013 equal to 27.7%, that is $[1.092 \times (1.016 \times 1.0618 \times 1.0885) - 1]$. This value can be contrasted with the actual change in the IT price premium obtained by applying the historical changes of 4.01%, 16.03%, and 24.12% in 1980-1990, 1990-2000, and 2000-2013, respectively, which compound to a premium of 63.6%.⁴⁸ It follows that changes in labor quality composition account for just below 1/2 of the increase in the overall IT price premium between 1980 and 2013. The economic intuition is that average labor efficiency in IT jobs has been deteriorating relative to non-IT jobs, implying an increase in the IT price premium (when the relative IT input goes down, because of lower average efficiency, its price goes up).

Table 4 documents some striking changes in the historical price ratio w^{IT}/w^{NIT} , which has grown significantly in both manufacturing and services sectors between 1980 and 2013. Specifically, in manufacturing w^{IT}/w^{NIT} jumped from 1.092 to 1.636 (a sixfold increase in the IT premium), while in services it went from 1.089 to 1.313 (a threefold increase in the IT premium).

For manufacturing, we find that factor shares played a major role in the growth of the IT price premium. If one removed the effect of technological change captured by the factor shares, the IT premium in 2013 would drop from 63.6% to just about 2%. However, this change pales in comparison to the effect of factor shares in services. When we remove the effect of exogenous technological change captured by factor shares, the w^{IT}/w^{NIT} for services in 2013 drops from 1.313 to 0.040, meaning that the IT premium drops from positive 31% to negative 96%. This reversal highlights the enormous pace of productivity changes in the services IT sector over the past four decades. As we discuss below, exogenous productivity growth biased towards IT jobs accounts for the bulk of the trends in structural transformation.

The of growth in the factor shares resonates with our reduced-form evidence. We showed earlier that the IT employment and earnings premia have both grown significantly over time, especially in services. The overall drift in the relative productivity of IT jobs has, therefore, helped shape the intensity and pace of structural transformation. In particular, both manufacturing and services had a progressive increase in IT employment driven by growth in the relative productivity of IT jobs. For example, nearly all of the decline in the employment share of manufacturing took place in non-IT jobs (see Figure 20 in Appendix A.3.3). The productivity drift of IT-intensive jobs translates into large factor share growth and its corresponding effects on relative prices. In fact, our results suggest that, absent growth in these factor shares, Baumol's cost disease would have likely overtaken the services sector as price effects would not have led to an increase in IT-intensive

⁴⁸The premium of 63.6% is obtained as $1.089 \times (1.0401 \times 1.1602 \times 1.2411) - 1$.

jobs, which have been integral for growth in labor productivity (Figure 5).

We also find that the changes in relative aggregate hours contributed negatively to the IT price premium in manufacturing, and positively in services. When we remove the effect of observed hours changes the manufacturing IT price premium in 2013 goes up to 80% from the observed price premium of 63.6%. The intuition behind this is simple: as the relative share of IT workers increases in manufacturing, the relative IT premium is dragged down. Historically, the ratio of IT job hours worked in manufacturing (relative to non IT job hours) increased by almost 16% between 1980 and 2013, implying a negative equilibrium effect on the IT price premium. Interestingly, the opposite is true for services: the ratio of hours worked in IT jobs to hours worked in non-IT jobs actually decreased by roughly 20% between 1980 and 2013. This was not due to a lack of growth in IT employment, which increased significantly, but rather to the fact that the growth was accompanied by a marginally faster growth in hours worked in non IT-intensive services jobs. Using the same equilibrium argument, this implies an increase in the IT price premium in services, so that in our counterfactual exercise we find that the observed 31.3% IT price premium would have been only a 10% premium in the absence of changes in relative hours worked.

These counterfactuals highlight the way in which the price of labor adjusts through different channels—that is, through the quantity and quality of labor services and broader macroeconomic trends that affect the drift of task-specific prices. There are two key observations we can draw from the findings so far. First, technological change has been driving the growth of IT in the recent past, with exogenous productivity growth being especially fast in IT service jobs. Second, this productivity growth has been accompanied by different outcomes for non IT-intensive jobs in manufacturing and services. In particular, disappearing non IT jobs are the main reason for the shrinking manufacturing share, while in the services sector non IT jobs have been growing steadily, together with IT jobs. In part, this is due to the fact that IT and non IT jobs are more easily substitutable in the manufacturing sector, as indicated by the significant difference in estimated elasticity. This is, however, only a part of the story: in fact, it is plausible that increasing demand for skilled IT services jobs outpaced the historical supply of workers that can cover them. These results relate with both an academic and popular press literature on the presence of skill mismatch (Güvenen et al., 2015),⁴⁹ especially for IT tasks.⁵⁰ This skill gap might be one of the reasons we observe a significant run-up in both IT tasks prices and employment.

⁴⁹<https://hbr.org/2014/08/employers-arent-just-whining-the-skills-gap-is-real>

⁵⁰<https://www.usatoday.com/story/tech/talkingtech/2017/03/28/tech-skills-gap-huge-graduates-survey-says/99587888/>

[INSERT TABLE 4 HERE]

Our second set of counterfactual exercises examines the role of IT jobs for the transformation of the employment structure (that is, hours worked in IT and non-IT jobs). We begin by computing the counterfactual changes in aggregate hours that would be consistent with no change in the relative price of IT-intensive jobs relative to non IT jobs (no change in the IT task price premium). This scenario is interesting because it provides a partial equilibrium counterpart to our previous results: if productivity changes in IT tasks had not been buffered by price adjustments, what would have been the change in the employment structure of IT and non-IT jobs?

Historically, in the manufacturing sector, IT hours-worked have gone up from just 36.4% of total hours in 1980 to roughly 40% of total hours worked in 2013. Using our technology estimates, we show in Table 5 that in the absence of price adjustments total hours worked in IT jobs in 2013 would have been over 55% of total manufacturing hours. This is significantly higher than the 40% share observed in 2013 and suggests that increases in the IT price premium have mitigated the speed of structural employment change in manufacturing. Whether this is due to frictions that do not allow for costless reallocation in production or for demand in available IT positions outstripping the supply of workers who can fill them, the result suggests that increasing IT task price premia have been instrumental in buffeting the crowding out of non IT jobs in manufacturing.

Looking at services, the IT hours share actually went down from 41.7% in 1980 to 36.2% in 2013. This drop did not occur because of a reduction in the number of IT jobs but rather for a more than proportional increase in the number of non IT positions. Services grew enormously over the past decades, in both IT and non IT jobs. However, when we consider a scenario with no price adjustments, the counterfactual share of IT jobs in 2013 jumps to 42.1%. This means that, just like in the case of manufacturing, price adjustments have weakened the shift to substitute non-IT jobs with IT jobs. In the absence of these price changes, the employment structure would have led to more IT jobs in 2013. Large structural changes in the price structure have reduced the need for large shifts in the prevalence of IT jobs, suggesting that higher wage inequality has helped mitigate even more extreme changes in the structure of employment over the past decades.

The last counterfactual exercise looks at the role of changes in the composition of workers' labor efficiency in IT and non-IT jobs. In particular, we restrict the worker quality composition in 2013 to be the same as the one prevailing in 1980. The bottom panel of Table 5 shows that under this counterfactual scenario the share of IT hours worked, out of total hours worked, in 2013 would be significantly lower both in manufacturing (28%) and services (30%). This result indicates that

the relative quality of IT workers has been deteriorating over time, as employers presumably tried to fill new IT positions by hiring workers with relatively less effective skills. This in turn pushed up the number of IT hours while lowering the average worker efficiency. Restricting workers quality to be constant would have required, all else equal, a lower share of IT-intensive hours. In this counterfactual we allow prices to adjust, keeping relative IT price growth to its historical values. The exercises in Table 5 suggest that both price adjustments and workers' selection have been of critical importance for observed changes in the structure of prices and employment.

[INSERT TABLE 5 HERE]

Our results point towards the importance of price adjustments over time. For example, the decline in the manufacturing employment share would have been much more severe and rapid if prices were not able to adjust. Similarly, absent growth in the productivity of IT services, prices would not have led to a reallocation of workers to IT jobs in the services sector, which we argue may have prevented Baumol's cost disease from taking root in the services sector. While our model offers insight into the competing roles of price and composition effects across sectors, we recognize that it only provides an approximation of the full suite of dynamic general equilibrium feedback that were present in reality. Future work should use these counterfactuals as heuristics to help discipline the inclusion of occupational sorting and productivity growth across sectors.

5 The Emergence of IT-Intensive Services and Solow's Paradox

Traditional arguments about the increasing pervasiveness of IT point towards greater productivity and automation in the work-place. However, as Acemoglu et al. (2014) clearly show for the U.S. manufacturing sector, there is little systematic evidence that IT-using industries have actually become more productive over time. While there is some evidence of faster growth in output per worker in IT-using industries, these differences are driven by a more rapid decline in labor employed in those industries and a relatively constant, or even declining, level of output (rather than higher output, as one might have expected). For these reasons, Acemoglu et al. (2014) conclude that Solow's famous paradoxical statement that "...you can see the computer age everywhere but in the productivity statistics" is far from being resolved. We revisit these issues through the lens of our model and make two contributions: (i) we use our measure of the IT employment share based

on occupational heterogeneity to replicate Acemoglu et al. (2014) and compare our results with those obtained using their industrial heterogeneity measure, and (ii) we apply our IT measure to distinguish between labor productivity growth in high and low IT industries within both the manufacturing and services sectors.

We begin by constructing the employment share of IT workers across each three-digit industry, comparing it with the baseline IT capital intensity measure in Acemoglu et al. (2014). The (weighted) correlation between these alternative measures of IT prevalence is 0.46 (see Figure 12).⁵¹ One reason our measure differs from industry-based proxies is that we focus on the underlying tasks that workers perform at a five-digit occupation level, rather than measuring capital assets as in prior work (e.g., Stiroh (2002)). While we do not argue that looking at a measure of IT-using labor is better than a measure of IT-using capital, we simply point out that they capture different features of reality. Using our employment share of IT by industry, we subsequently estimate regressions identical to Acemoglu et al. (2014) at the three-digit industry \times year level:

$$y_{jt} = \sum_t \delta^t (IT_j \times \lambda_t) + \eta_j + \lambda_t + \epsilon_{jt} \quad (11)$$

where y denotes the logged real value of shipments in our first result and logged real output in our second result (since services does not have shipments), IT denotes a measure of the IT intensity, η denotes three-digit industry fixed effects, and λ denotes year fixed effects. The primary coefficients of interest in Equation 11 are the δ^t values. We normalize the coefficient on the IT variable to zero in the base year δ^0 . Hence the series $\{\delta^1, \delta^2, \dots\}$ may be read as the level of the coefficient on IT in each subsequent year, relative to the base year. These coefficients characterize the deviations from trend in the return to IT intensity across industries: hence, only their time variation (rather than their absolute value) is of interest for our analysis.

As mentioned above, we use the baseline measure of IT intensity taken from Acemoglu et al. (2014), which is the share of investment in computers (relative to total investment), together with our employment share of IT, which denotes the share of workers in IT-intensive jobs within a given industry.

[INSERT FIGURE 12 HERE]

⁵¹We also implemented a different comparison by partitioning industries into high and low IT based on our IT-intensity measure, as well as the measures from Acemoglu et al. (2014) based on computer investment, the SMT dummy, and computer investment between 1977-82. The results imply that the first measure matches our high and low IT industries in 72% of the cases, the second measure only matches ours in 48% of the cases, and the third measure matches ours in 64% of the cases.

Figure 13 plots the estimated δ^t coefficients for each year. We follow Acemoglu et al. (2014) as closely as possible by crosswalking industries from our CPS micro-data to the three-digit industry level, which we match into their SIC codes. We also drop the sample of industries that falls within their “medium computer intensity” measure since those industries expand mechanically over time due to the increasing pervasiveness of computers in the economy.⁵² The blue line in Figure 13 replicates their main result that IT use at the industry level is negatively associated with real shipments over time. However, the red line in Figure 13 illustrates a near opposite pattern when using our measure of IT intensity (correlation = -0.79)—that is, rising IT intensity is associated to positive changes in the real value of shipments. As we discussed earlier, one advantage of our measure is that it captures variation in the utilization of different tasks within each industry, which is important since IT content is inherently occupation-specific.

Why do these measures produce such heterogeneous productivity profiles? Our measure focuses on tasks (and associated human capital among the labor force) in IT jobs, whereas Acemoglu et al. (2014) focus on a industry-level measure of physical IT capital. Hence these two indices are capturing systematically different phenomena. Given our earlier evidence that the manufacturing sector exhibits greater substitutability between IT and non-IT workers, it is precisely the industries that invested more in physical capital at the expense of human capital that appear to have had lower output values. In sum, these differences underscore that the measurement strategy matters.

[INSERT FIGURE 13 HERE]

Motivated by these differences between our two measures, we next examine how the IT employment share affects our understanding of productivity growth in the manufacturing versus services sectors. That is, we now estimate Equation 11 using logged real output (deflated with 2009 price indices) as the outcome variable and the IT employment share as the IT intensity index, plotting the estimated δ^t coefficients separately for the manufacturing and services sectors.⁵³ Figure 14 highlights two important observations. First, and perhaps most importantly, whereas IT intensity explains a relatively constant share of real output productivity over the past 40 years in the manufacturing sector, IT intensity is growing rapidly in explaining the surge of productivity in the services sector. Second, the positive effect of IT intensity in the services sector has continued unabated also after the year 2000. In Appendix A.5, we also plot results for value added, sepa-

⁵²These industries in the “medium intensity” range of the classification include: SIC 3570-3579, SIC 3660-3669, and SIC 3670-3679. They also consider a narrower and a broader definition from Houseman et al. (2013).

⁵³We use real output since it most closely matches the real value of shipments used in Acemoglu et al. (2014).

rately for high versus low IT and manufacturing versus services sectors. Not surprisingly, we find that the bulk of the rise in value added is accounted for by the high IT services sector, consistent with Figure 5.

[INSERT FIGURE 14 HERE]

6 Conclusion

Recent debates about the role of technology and automation for labor productivity across different jobs have highlighted the need for new and alternative ways to describe and quantify the emergence of information technologies in production. Important manifestations of the rise in information technology (IT) is its impact on the labor market (Autor et al., 1998), the organization of firms (Bresnahan et al., 2002), productivity (Brynjolfsson and Hitt, 2003), and managerial practices (Bloom et al., 2012). However, much less is known about how the emergence of these jobs has influenced the intensity and rate of structural transformation—that is, the move from manufacturing to services employment. A key contribution of our work is to provide a new measurement of IT jobs in the labor market and assess how technological change has shaped sorting across jobs and the resulting sectoral distribution of task returns and employment.

We begin by documenting three facts. First, while the manufacturing sector has higher output per worker than the services sector, it also exhibits much lower value added per worker than the “high-tech” services sector, which has accounted for much of the services wage and employment growth in the past few decades. Second, after introducing our new measure of IT intensity, we classify jobs into high and low IT status and show that there are large and growing earnings and employment premia among these job categories over the past 40 years. In fact, we show that despite a flattening skill premium in college and cognitive/non-routine jobs after the year 2000, there has been a continued rise in the premium among IT jobs. Third, we show that increases in the IT employment share are associated with declines in the manufacturing share, even after controlling for a number observable features and unobserved fixed effects.

Motivated by these observations we develop and estimate a model of production (value added) featuring occupation heterogeneity. The model closely mimics the structure originally proposed by Adao (2016) in a different context, explicitly allowing for self-selection into IT and non-IT occupations. Using variation in hourly wages, and the employment share of IT in 1980, we use the model to estimate changes in the IT premium in both prices and relative workers’ efficiency, across

different locations and industries. Using these estimated price and composition effects, we recover key parameters in a task based value-added production technology that enable us to implement several counterfactual experiments. We use these experiments to quantify the relative importance of productivity growth in IT jobs and how it has shaped structural transformation since 1980.

We find that the elasticity of substitution between IT and non-IT jobs in the manufacturing sector is around 1.6, whereas in the services sector the elasticity is lower, around 1.3. Using the estimated elasticities, and growth in the factor shares of IT and non-IT jobs, we document that price changes played a fundamental role in shaping the rate and pace of structural transformation. In particular, we find that holding fixed growth in the factor share of IT jobs would have effectively stopped the reallocation of workers from manufacturing to IT-intensive jobs in the services sector. In other words, absent growth in the productivity of IT jobs since 1980, any reallocation that would have occurred in the services sector would have been directed towards non IT-intensive jobs, which would likely have led to much lower growth in aggregate productivity and a variant of Baumol’s cost disease.

Our results are not meant to provide a definitive conclusion on the relationship between occupation-level information technology and structural transformation. Rather, we have presented the first evidence on their relationship to discipline and motivate additional research questions about the ways in which technological changes have altered labor market sorting, average efficiency and inequality across sectors. Nonetheless, we have shown that the emergence of IT jobs has had a big impact on employment and the distribution of wages both within and across sectors. These changes may in principle account for a non-trivial part of the increasing polarization across jobs documented by Autor and Dorn (2013). We leave further analysis of these implications for future work.

References

- ACEMOGLU, D., D. AUTOR, D. DORN, G. H. HANSON, AND B. PRICE (2014): “Return of the solow paradox? IT, productivity, and employment in US manufacturing,” *American Economic Review: Papers & Proceedings*, 104, 394–399.
- ACEMOGLU, D. AND D. H. AUTOR (2011): “Skills, tasks and technologies: Implications for employment and earnings,” *Handbook of Labor Economics*, 4b, 1043–1171.
- ACEMOGLU, D., D. H. AUTOR, AND D. LYLE (2004): “Women, war and wages: The effect of female labor supply on the wage structure at midcentury,” *Journal of Political Economy*, 112, 497–551.
- ACEMOGLU, D. AND V. GUERRIERI (2008): “Capital deepening and nonbalanced economic growth,” *Journal of Political Economy*, 116, 467–498.
- ACEMOGLU, D. AND P. RESTREPO (2017a): “Robots and jobs: Evidence from US labor markets,” *Working paper*.
- (2017b): “The race between man and machine: Implications of technology for growth, factor shares and employment,” *Working paper*.
- ADAO, R. (2016): “Worker heterogeneity, wage inequality, and international trade: Theory and evidence from Brazil,” *Working paper*.

- AEDO, C., J. HENTSCHEL, J. LUQUE, AND M. MORENO (2013): "From occupations to embedded skills: A cross-country comparison," *Background Paper for the World Development Report*.
- ATALAY, E., P. PHONGTHIENGTHAM, S. SOTELO, AND D. TANNENBAUM (2017): "The evolving U.S. occupational structure: A textual analysis," *Working paper*.
- AUTOR, D. (2014): "Skills, education and the rise of earnings inequality among the 'other 99 percent'," *Science*, 344, 843–851.
- (2017): "How long has this been going on?" *NBER Conference on Research and Income in Wealth*.
- AUTOR, D. AND D. DORN (2013): "The Growth of Low Skill Service Jobs and the Polarization of the U.S. Labor Market," *American Economic Review*, 103, 1553–1597.
- AUTOR, D. H. (2013): "The 'task approach' to labor markets: An overview," *NBER working paper*.
- (2015): "Why are there still so many jobs? The history and future of workplace automation," *Journal of Economic Perspectives*, 29, 3–30.
- AUTOR, D. H., D. DORN, AND G. H. HANSON (2013): "The China syndrome: Local labor market effects of import competition in the United States," *American Economic Review*, 103, 2121–2168.
- (2015): "Untangling trade and technology: Evidence from local labour markets," *Economic Journal*, 125, 621–646.
- AUTOR, D. H., L. F. KATZ, AND M. S. KEARNEY (2006): "The Polarization of the U.S. Labor Market," *American Economic Review*, 96, 189–194.
- (2008): "Trends in U.S. wage inequality: Revising the revisionists," *Review of Economics and Statistics*, 90, 300–323.
- AUTOR, D. H., L. F. KATZ, AND A. B. KRUEGER (1998): "Computing inequality: Have computers changed the labor market?" *Quarterly Journal of Economics*, 113, 1169–1213.
- AUTOR, D. H., F. LEVY, AND R. J. MURNANE (2003): "The skill content of recent technological change: An empirical exploration," *Quarterly Journal of Economics*, 118, 1279–1333.
- BARANY, Z. AND C. SIEGEL (2017): "Job polarization and structural change," *American Economic Journal: Macroeconomics*, forthcoming.
- BAUMOL, W. J. (1967): "Macroeconomics of unbalanced growth: Anatomy of an urban crisis," *American Economic Review*, 57, 415–426.
- BEAUDRY, P., D. A. GREEN, AND B. M. SAND (2014): "The declining fortunes of the young since 2000," *American Economic Review*, 104, 381–386.
- (2016): "The Great Reversal in the demand for skill and cognitive tasks," *Journal of Labor Economics*, 34, S199–S247.
- BECKHUSEN, J. (2016): "Occupations in information technology," *American Community Survey Reports*.
- BLOOM, N., R. SADUN, AND J. VAN REENEN (2012): "Americans do IT better: US multinationals and the productivity miracle," *American Economic Review*, 102, 167–201.
- BOPPART, T. (2014): "Structural change and the Kaldor Facts in a growth model with relative price effects and non-Gorman preferences," *Econometrica*, 82, 2167–2196.
- BRESNAHAN, T., E. BRYNJOLFSSON, AND L. HITT (2002): "Information technology, workplace organization and the demand for skilled labor: firm-level evidence," *Quarterly Journal of Economics*, 117, 339–376.
- BRYNE, D. M., J. G. FERNALD, AND M. B. REINSDORF (2016): "Does the United States have a productivity slowdown or a measurement problem?" *Brookings Papers on Economic Activity*, 109–182.
- BRYNE, D. M., S. OLINER, AND D. E. SICHEL (2013): "Is the information technology revolution over?" *International Productivity Monitor*, 25, 20–36.
- BRYNJOLFSSON, E. AND L. M. HITT (2003): "Computing productivity: Firm-level evidence," *Review of Economics and Statistics*, 85, 793–808.
- BRYNJOLFSSON, E. AND A. MCAFEE (2011): "Race against the machine: How the digital revolution is accelerating innovation, driving productivity, and irreversibly transforming employment and the economy," *Lexington, MA: Digital Frontier Press*.
- BRYNJOLFSSON, E., D. ROCK, AND C. SYVERSON (2017): "Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics," *NBER working paper*.

- BUERA, F. J. AND J. P. KABOSKI (2009): “Can traditional theories of structural transformation fit the data?” *Journal of the European Economic Association*, 7, 469–477.
- (2012): “The rise of the service economy,” *American Economic Review*, 102, 2540–2569.
- CABALLERO, R. J. AND M. L. HAMMOUR (1994): “The cleansing effect of recessions,” *American Economic Review*, 84, 1350–1368.
- CARD, D. AND T. LEMIEUX (2001): “Can falling supply explain the rising return to college for younger men? A cohort based analysis,” *Quarterly Journal of Economics*, 116, 705–746.
- CHARLES, K. K., E. HURST, AND M. J. NOTOWIDIGDO (2016): “The masking of the decline in manufacturing employment by the housing bubble,” *Journal of Economic Perspectives*, 30, 179–200.
- COMIN, D. A., D. LASHKARI, AND M. MESTIERI (2015): “Structural change with long-run income and price effects,” *Econometrica*, *R&R*.
- DUERNECKER, G., B. HERRENDORF, AND A. VALENTINYI (2017): “Structural change within the service sector and the future of Baumol’s disease,” *Working paper*.
- EDEN, M. AND P. GAGGL (2017): “On the welfare implications of automation,” *Working paper*.
- FELDSTEIN, M. (2015): “The United States underestimates growth,” *Wall Street Journal*.
- FREY, C. AND M. A. OSBORNE (2013): “The future of employment: How susceptible are jobs to computerisation?” *Working paper*.
- GOLLIN, D., S. PARENTE, AND R. ROGERSON (2002): “The role of agriculture in development,” *American Economic Review*, 92, 160–164.
- GUVENEN, F., B. KURUSCU, S. TANAKA, AND D. WICZER (2015): “Multidimensional skill mismatch,” *NBER working paper*.
- HARRIGAN, J., A. RESHEF, AND F. TOUBAL (2016): “March of the techies: Technology, trade, and job polarization in France, 1994-2007,” *Working paper*.
- HATZIUS, J. AND K. DAWSEY (2015): “Doing the sums on productivity paradox v2.0,” *Goldman Sachs U.S. Economics Analyst*, 15/30.
- HECKER, D. (2005): “High-technology employment: A NAICS-based update,” *Monthly Labor Review*.
- HERRENDORF, B., C. HERRINGTON, AND A. VALENTINYI (2015): “Sectoral technology and structural transformation,” *American Economic Journal: Macroeconomics*, 7, 104–133.
- HERRENDORF, B., R. ROGERSON, AND A. VALENTINYI (2013): “Two perspectives on preferences and structural transformation,” *American Economic Review*, 103, 2752–2789.
- (2014): “Growth and structural transformation,” *Handbook of Economic Growth*, 2B, 855–941.
- HOUSEMAN, S. N., T. BARTIK, AND T. STURGEON (2013): “Measuring manufacturing: Problems of interpretation and biases in the U.S. statistics,” *Upjohn Institute*.
- JUHN, C., K. M. MURPHY, AND B. PIERCE (1993): “Wage inequality and the rise in returns to skill,” *Journal of Political Economy*, 101, 410–442.
- KAMBOUROV, G. AND I. MANOVSKII (2013): “A cautionary note on using (March) CPS and PSID data to study worker mobility,” *Macroeconomic Dynamics*, 17, 172–194.
- KATZ, L. F. AND K. M. MURPHY (1992): “Changes in relative wages, 1963-1987: Supply and demand factors,” *Quarterly Journal of Economics*, 107, 35–78.
- KONGSAMUT, P., S. REBELO, AND D. XIE (2001): “Beyond balanced growth,” *Review of Economic Studies*, 68, 869–882.
- KUZNETS, S. (1973): “Modern economic growth: Findings and reflections,” *American Economic Review*, 63, 247–258.
- LEMIEUX, T. (2006): “Increasing residual wage inequality: Composition effects, noisy data, or rising demand for skill?” *American Economic Review*, 96, 461–498.
- LJUNGQVIST, L. AND T. J. SARGENT (1998): “The European unemployment dilemma,” *Journal of Political Economy*, 106, 514–550.
- (2008): “Two questions about European unemployment,” *Econometrica*, 76, 1–29.
- MOKYR, J. (2014): “Secular stagnation? Not in your life,” *London: CEPR Press*.
- NGAI, L. R. AND C. A. PISSARIDES (2007): “Structural change in a multisector model of growth,” *American Economic Review*, 97, 429–443.

- 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, 647–678.
- STIROH, K. J. (2002): “Information technology and the U.S. productivity revival: What do the industry data say?” *American Economic Review*, 92, 1559–1576.
- SYVERSON, C. (2017): “Challenges to mismeasurement explanations for the US productivity slowdown,” *Journal of Economic Perspectives*, 31, 165–186.
- TAMBE, P. AND L. M. HITT (2012): “The productivity of information technology investments: New evidence from IT labor data,” *Information Systems Research*, 23, 599–617.
- VALLETTA, R. G. (2016): “Recent flattening in the higher education wage: premium: Polarization, skill downgrading, or both?” *Education, skills, and technical change: Implications for future U.S. GDP growth*.

7 Figures and Tables




Figure 1: Structural Transformation and the Rise of Information Technology

Notes.—Sources: St. Louis Federal Reserve Board. The figure plots the ratio of manufacturing employment to services employment in the US economy vs a well-known technology index. The Federal Reserve Bank of San Francisco collects the technology index (“the tech pulse index”) and is defined (from the FRED site) as follows: “a coincidence index of activity in the U.S. information technology sector. The index interpreted as the health of the tech sector”. The indicators they use to construct the index are: investment in IT goods, consumption of personal computers and software, employment in the IT sector, industrial production of the technology sector, and shipments by the technology sector.




Figure 2: Employment Share and Earnings in Information Technology Jobs

Notes.—Sources: O*NET and Census Bureau, 1980-2015. The figure plots logged annual earnings (weighted by the Census sample weights) with the standardized information technology services intensity (weighted by employment from the BLS) at the five-digit occupation level. Earnings are deflated using the 2010 personal consumption expenditure index and the sample is restricted to workers earning over \$5,000/year, \$2/hour, and working 500 hours/year. The figure shows a strong relationship between IT intensity and earnings in the cross-section of occupations.




Figure 3: Validating Occupation Information Technology (IT) with Capital IT

Notes.—Sources: Census Bureau, O*NET, Bureau of Economic Analysis, 2000-2015. The figure plots the IT employment share at a three-digit industry level with logged capital IT expenditures (weighted by employment from the BLS). The figure shows that there is a 0.51 correlation between capital and labor IT inputs in the cross-section of three-digit industries (weighted by employment).

Table 1: Descriptive Statistics on High and Low Information Technology Workers

	pooled											
	1970s		1980s		1990s		2000s		2010-2015			
	High	Low	High	Low	High	Low	High	Low	High	Low	High	Low
Demographics												
family size	2.9	3.1	3.4	3.4	3.0	3.1	2.9	3.1	2.9	3.0	2.9	3.0
number of children	1.0	0.9	1.4	1.3	1.0	1.0	0.9	0.9	0.9	0.9	0.9	0.8
age	41.5	39.8	40.2	40.3	38.7	37.6	39.1	37.7	40.7	39.2	42.7	40.7
male	0.58	0.53	0.76	0.59	0.67	0.54	0.59	0.53	0.57	0.52	0.54	0.52
married	0.65	0.54	0.80	0.72	0.72	0.64	0.68	0.59	0.66	0.55	0.61	0.46
white	0.83	0.76	0.95	0.87	0.92	0.86	0.88	0.80	0.84	0.76	0.79	0.72
black	0.08	0.13	0.04	0.12	0.06	0.12	0.07	0.12	0.07	0.12	0.09	0.14
schooling	14.5	12.2	13.4	10.6	14.1	11.7	14.4	12.3	14.5	12.4	14.9	12.6
college	0.52	0.14	0.33	0.04	0.41	0.08	0.47	0.11	0.51	0.13	0.58	0.18
Work												
hours worked	2105	1933	.	.	2017	1878	2058	1923	2110	1948	2123	1933
earnings	54658	30395	44308	27390	43121	27121	47347	29010	54023	31859	58486	30820
hourly wage	27.1	15.9	.	.	21.4	14.7	23.7	15.1	26.0	16.4	28.3	15.8
manufacturing	0.23	0.27	0.33	0.40	0.30	0.37	0.27	0.29	0.24	0.26	0.20	0.21
services	0.69	0.69	0.57	0.51	0.61	0.59	0.64	0.67	0.69	0.70	0.73	0.75
Observations	8123943	11141496	377515	719532	1408330	2183323	1617786	2372254	1723147	2297422	1546484	1806519

Notes.—Sources: Census Bureau and O*NET, 1970-2015. The table reports summary statistics associated with major demographic and labor market variables for high and low information technology (IT) workers by decade in our micro-data. The variables “manufacturing” and “services” refer to the employment share of workers. We refer to high IT workers as those classified as above the median based on our five-digit IT intensity z -score. Earnings are deflated using the 2010 personal consumption expenditure index and the sample is restricted to workers earning over \$5,000/year, \$2/hour, and working 500 hours/year. The manufacturing sector includes both the standard NAICS 31-33 codes on durables and non-durables and both agriculture and mining. The services sector includes finance, insurance, real estate, business services, wholesale and retail trade, utilities, transportation, personal and other services.

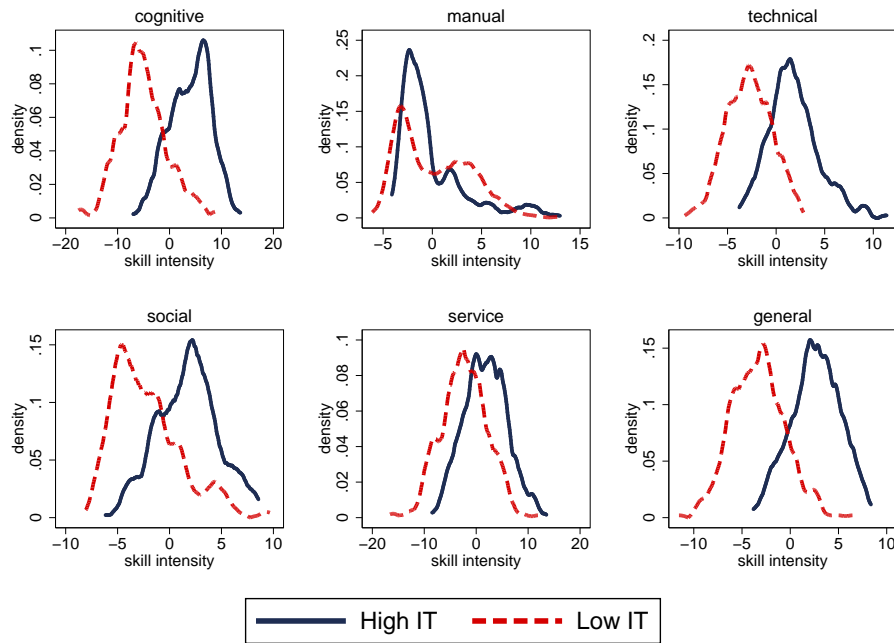


Figure 4: Distribution of Skills Across Technology & Information Services Jobs

Notes.—Sources: Occupation and Employment Statistics (national) and O*NET, 2004-2014. The figure plots the distribution of the six major categories of skills between high and low information technology intensive jobs. The skill groups are as follows: (1) cognitive skills (decision making, learning strategies, listening, learning, problem solving, coordination, and critical thinking), (2) manual (repairs, equipment maintenance, equipment selection, installation, instruction), (3) technical (programming, quality control analysis, systems analysis, systems evaluation, technology design), (4) social (persuasion, social, speaking, negotiation), (5) service (management of financial resources, of material resources, of personnel resources, monitoring, service, operations control, operations monitoring, operations analysis, troubleshooting), and (6) general (math, writing, time management, reading, science). The ONET skill data is available from 2010-2014 and is made to have a mean zero and variance of 1. Occupations are harmonized to 2010 SOC codes.




Figure 5: Value Added & Gross Output / Employee in Manufacturing and Services, 1947-2014

Notes.—Sources: USA KLEMS, 1947-2014. The figure plots average value added and gross output per worker in constant dollars for manufacturing and services sectors. 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). To produce these national time series, we weight by average long-term employment between 1950-2014.




Figure 6: Earnings and Employment Premia to Information Technology, 1970-2015

Notes.—Sources: Census Bureau and O*NET, 1970-2015. Panel A plots 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. 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. Panel B plots the logged total number of IT workers net of the logged total number of non-IT workers. IT workers are classified as those in an occupation with an IT intensity score above the median in the distribution of five-digit occupations. The figure shows the growing earnings premium and declining employment premium (i.e., growing employment share) of IT workers.

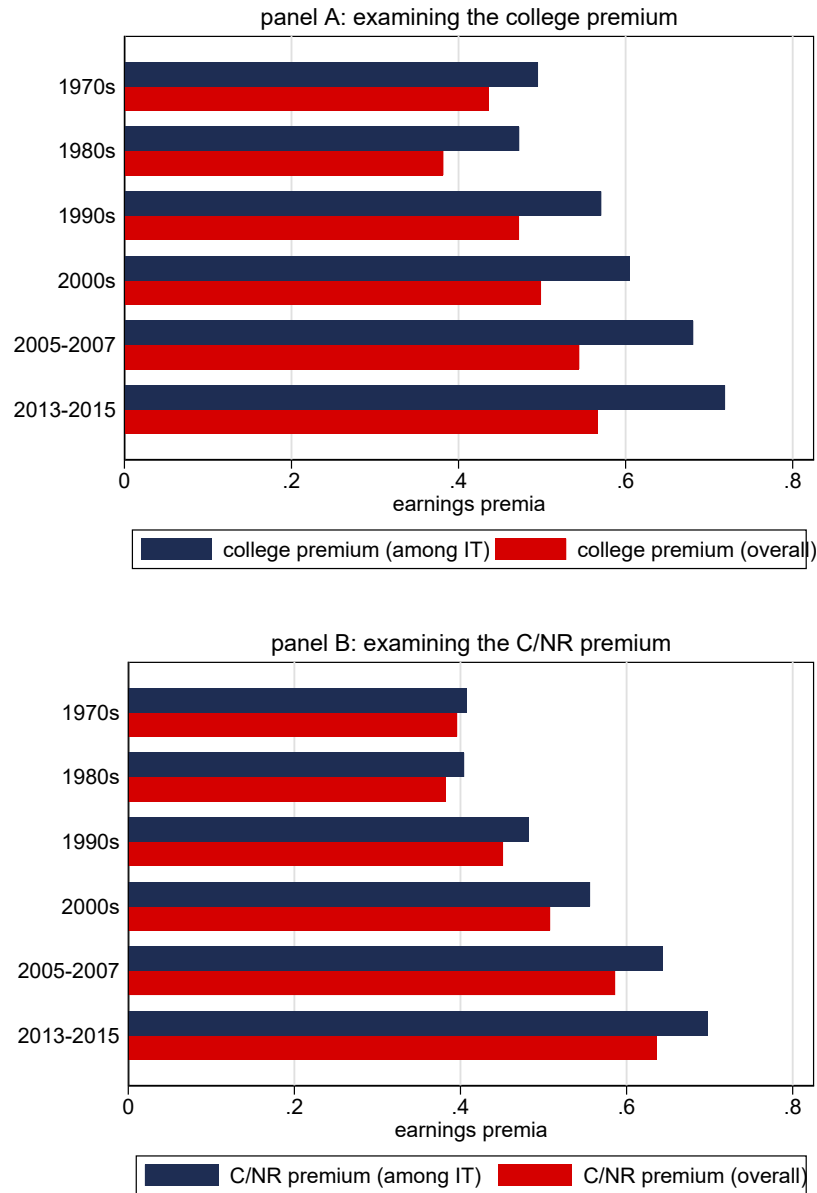


Figure 7: Examining Heterogeneity in the Demand for Skills in High and Low IT Jobs

Notes.—Sources: Census Bureau and O*NET, 1970-2015. Panel A plots the logged (average) earnings among IT and college workers net of earnings among non-college workers (in blue) with the logged (average) earnings among college workers net of logged earnings among non-college workers (in red), showing a more rapid increase among the college+IT premium, relative to the pure college premium. Panel B implements an analogous comparison, but uses cognitive and non-routine skills classified at the five-digit occupation level based on whether the intensity is above or below the median. Earnings are deflated using the 2010 personal consumption expenditure index and the sample is restricted to workers earning over \$5,000/year, \$2/hour, and working 500 hours/year. All aggregations from the Census are produced using sample weights.




Figure 8: Return to IT and College / C/NR Skills, 2005-2015

Notes.—Sources: Census Bureau and O*NET, 2005-2015. Panel A plots the coefficient associated with the college \times high IT status interactions from regressions of logged annual earnings on college attainment, high IT status, and their interaction separately by year. Panel B implements an analogous comparison, but uses cognitive and non-routine skills classified at the five-digit occupation level based on whether the intensity is above or below the median. Earnings are deflated using the 2010 personal consumption expenditure index and the sample is restricted to workers earning over \$5,000/year, \$2/hour, and working 500 hours/year. Regression estimates are weighted using sample weights and standard errors are clustered at the five-digit occupation level.

Table 2: The Rise of the Information Technology Employment Share and Structural Change

Dep. var. =	manufacturing employment share				
	(1)	(2)	(3)	(4)	(5)
IT share	-.67*** [.06]	-.23** [.09]	-.38*** [.06]	-.35*** [.09]	-.66*** [.19]
ln(wage)		.12*** [.03]		.11*** [.03]	.15*** [.03]
male		.95*** [.07]		.72*** [.09]	1.05*** [.13]
white		.14*** [.02]		-.00 [.03]	.18*** [.02]
college		-.37*** [.05]		-.23*** [.06]	-.11 [.10]
R-squared	.22	.61	.92	.94	.52
Sample Size	1434	1434	1422	1422	1090
Controls	No	Yes	No	Yes	Yes
County FE	No	No	Yes	Yes	No
Year FE	No	No	Yes	Yes	No
Instrument?	No	No	No	No	Yes

Notes.—Sources: Census Bureau, 1980-2015. The table reports estimated coefficients obtained by regressing the manufacturing employment share at the county-level on the IT employment share, conditional on the logged hourly wage, share of males, share of whites, and share of college degree workers. There are roughly 250 counties in each year of our sample. Fixed effects are included for county and year. The instrumental variable specification in column (5) uses the pre-determined 1980 share of IT workers to predict future IT growth. The sample is restricted to workers earning over \$5,000/year, \$2/hour, and working at least 500 hours/year. Standard errors are clustered at the county-level.

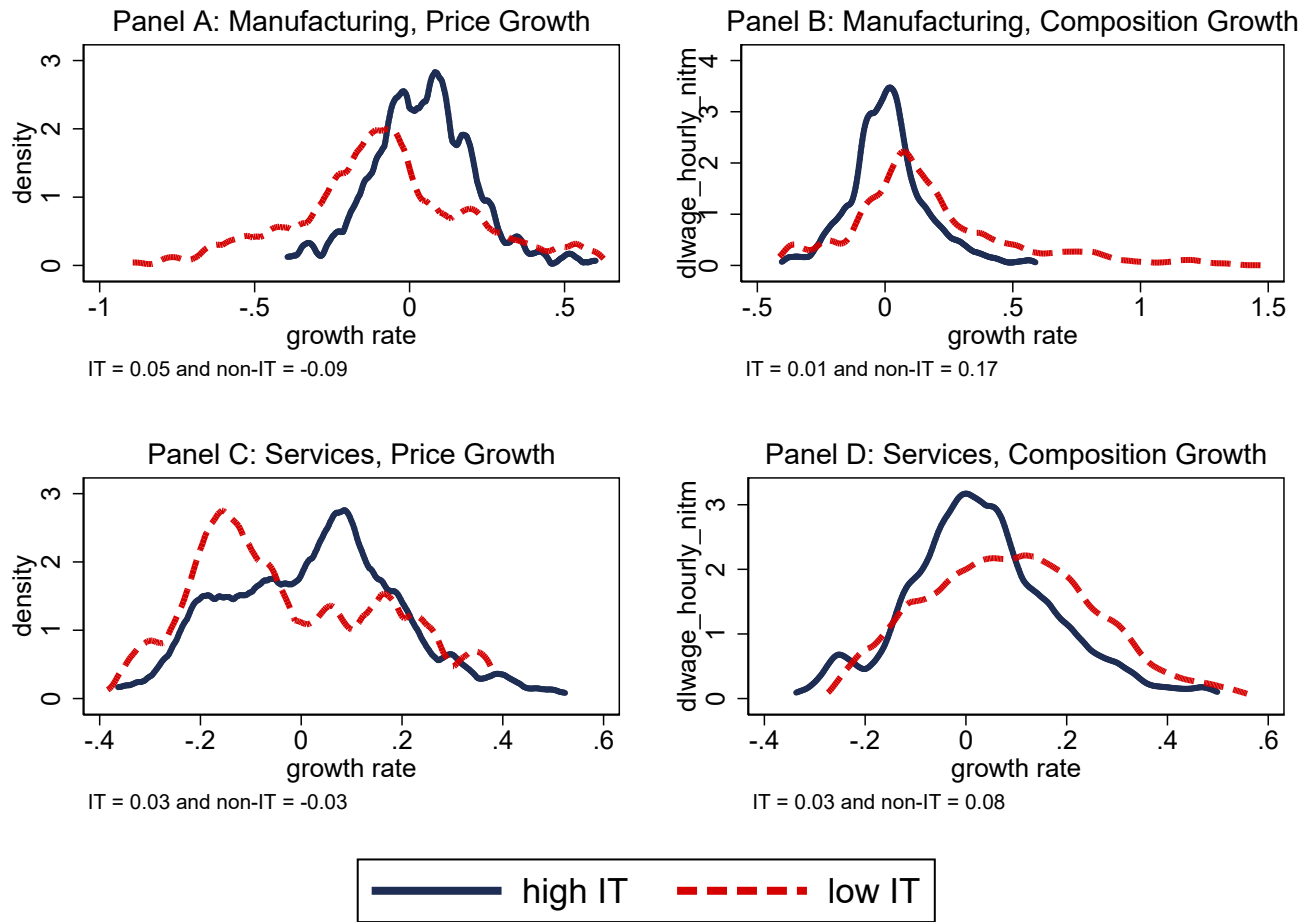


Figure 9: Distribution of Labor Price and Efficiency Growth in IT and non-IT Jobs

Notes.—Sources: Census Bureau, 1990-2015. The figure plots the distributions of the price and composition effects pooling across group, and region and year separately for manufacturing / services sectors. Price effects are estimated from the following regression $\Delta Y_{g,r,t}(\pi) = \Delta \omega_{g,r,t}^{IT} + [\Delta \omega_{g,r,t}^{NIT} - \Delta \omega_{g,r,t}^{IT}] I_{g,r,t_0}^{NIT}(\pi) + \mu_{g,r,t} X_{g,r,t}(\pi) + \Delta \nu_{g,r,t}(\pi)$ for $t_0 = 1980$. Controls include: the age distribution (fraction of people between ages 20-29, 30-39, 40-49, 50-65), marital status, race, gender, and dummies on the bottom, middle, and top of the wage distribution. Composition effects are obtained by taking the ratio of the gross wage growth and the price effects.

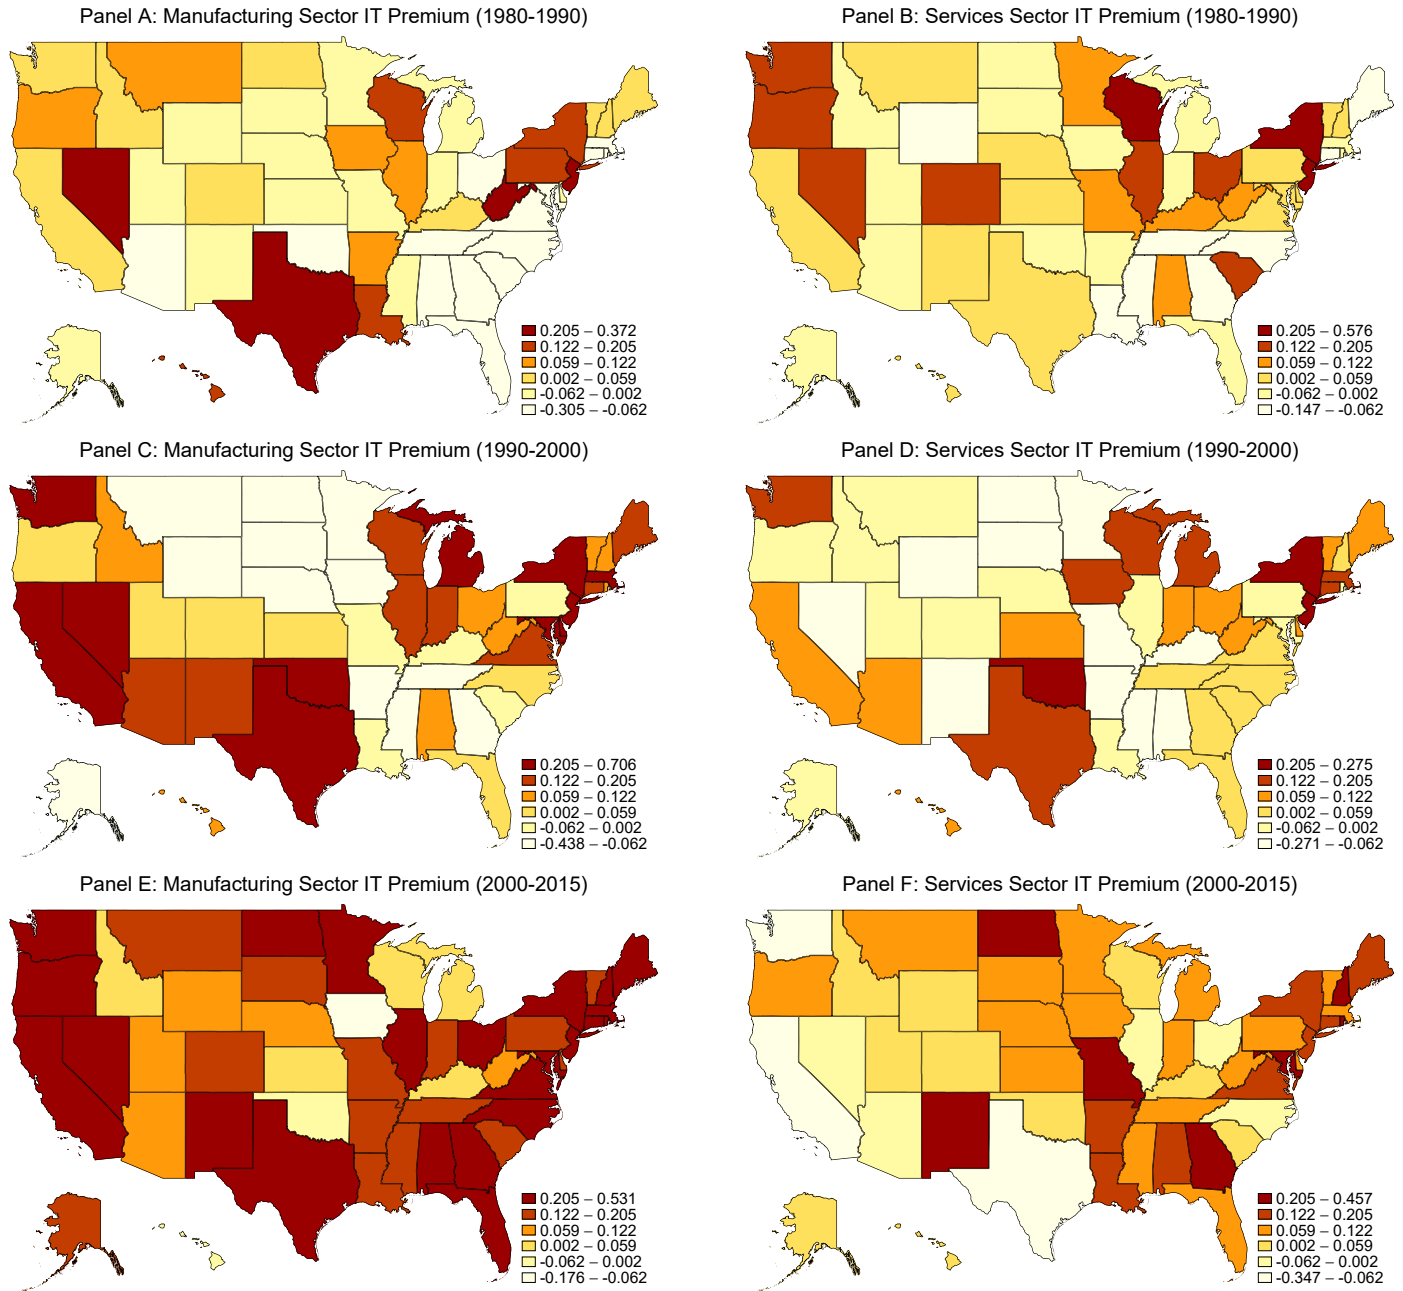


Figure 10: Spatial Heterogeneity in the Growth in the IT Price Premium

Notes.—Sources: Census Bureau, 1990-2015. The figure plots the state-level averages of the differences in price growth for labor services (the difference between $\Delta\omega^{IT}$ and $\Delta\omega^{NIT}$) estimated from the Adao (2016) model separately for manufacturing and services. Averages are weighted by the number of observations in each of the cells.

	Services	Manufacturing	Services	Manufacturing
$1/(1 - \nu_i)$	Elasticities of substitution (OLS)		Elasticities of substitution (IV)	
	1.310 (0.042)	1.627 (0.056)	1.390 (0.046)	1.682 (0.065)
$\Delta \ln[\alpha_i/(1 - \alpha_i)]$	Value Shares (OLS)		Value Shares (IV)	
	0.731 (0.280)	0.163 (0.323)	0.515 (0.240)	0.036 (0.319)

Table 3: Elasticity of Substitution between IT and non-IT Labor

Notes.—Sources: Census Bureau, 1990-2015. The table reports the coefficients associated with projections of the growth in the price of IT labor relative to non-IT labor on the growth of the ratio of quality-adjusted IT labor relative to non-IT labor:

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

where we restrict the elasticity to be time-invariant under both OLS and IV specifications. We also control for two wage dummies for where individuals lie in the wage distribution, age dummies for the share of workers in different (g, r, t) cells in the age distribution, the share who are male, married, and white. Our IV specification uses the share of IT workers in a college \times state cell from 1970 and the growth in the share of IT workers from 1970 to 1980 as instruments for the changes in quality-adjusted labor inputs. Our sample is restricted to the growth rates between 1980 to 1990, 1990 to 2000, and 2000 to 2015 and observations are weighted by the number of observations observed in each cell in 1980.

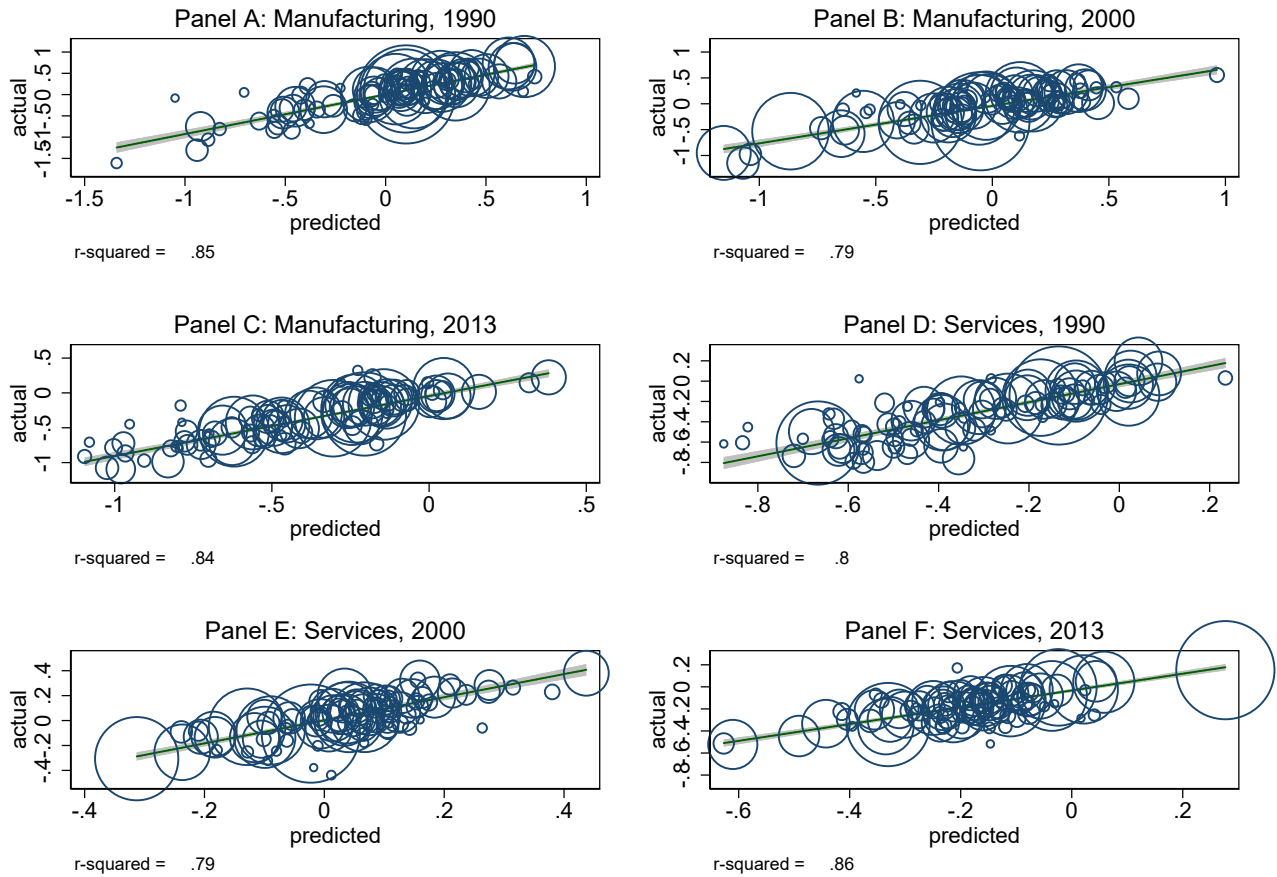


Figure 11: Comparison Human Capital Premium Growth, Data and Model

Notes.—Sources: Census Bureau and O*NET, 1990-2014. The figure plots the actual and predicted values for the ratio of marginal product growth rates between IT and non-IT workers in the manufacturing and services sector where the predicted values are obtained from the following projection:

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

where we restrict the elasticity to be time-invariant. We also control for two wage dummies for where individuals lie in the wage distribution, age dummies for the share of workers in different (g, r, t) cells in the age distribution, the share who are male, married, and white. We subsequently use these estimated coefficients to recover the human capital (the product of composition and aggregate hours) premium growth. Observations are weighted by the number of observations observed in each cell in 1980.

Year	Manufacturing		Services	
	Price ratio (IT vs non IT)	Price ratio (IT vs non IT)	Price ratio (IT vs non IT)	Price ratio (IT vs non IT)
	Actual	Counterfactual	Actual	Counterfactual
No Change in Worker Efficiency	w^{IT}/w^{NIT}	w^{IT}/w^{NIT}	w^{IT}/w^{NIT}	w^{IT}/w^{NIT}
1980	1.092	—	1.089	—
2013	1.636	1.277	1.313	1.136
No Change in Hours Worked	w^{IT}/w^{NIT}	w^{IT}/w^{NIT}	w^{IT}/w^{NIT}	w^{IT}/w^{NIT}
1980	1.092	—	1.089	—
2013	1.636	1.800	1.313	1.100
No Change in Factor Shares	w^{IT}/w^{NIT}	w^{IT}/w^{NIT}	w^{IT}/w^{NIT}	w^{IT}/w^{NIT}
1980	1.092	—	1.089	—
2013	1.636	1.019	1.313	0.040

Table 4: Counterfactual Changes to Assess the Gross IT Price Premia

Notes.—Using estimates of our comparative advantage model, we simulate three counterfactuals, which hold fixed the change in worker efficiency, hours worked, and factor shares. We simulate the counterfactual price effects in each of these cases and fix the baseline 1980 price premium level, extrapolating forward using subsequent decadal growth rates.

	Share of IT hours in total hours worked (%)	
	Manufacturing	Services
Set $\Delta(\omega^{IT} - \omega^{NIT}) = 0$ (no price changes)		
1980 (actual)	36.4%	41.7%
2013 (actual)	39.8%	36.2%
2013 (counterfactual)	55.3%	42.1%
Set $\Delta \ln(E^{IT}/E^{NIT}) = 0$ (no quality changes)		
1980 (actual)	36.4%	41.7%
2013 (actual)	39.8%	36.2%
2013 (counterfactual)	28.1%	30.9%

Table 5: Counterfactual Changes in the Information Technology Employment Share

Notes.—Using estimates of our comparative advantage model, we simulate two counterfactuals. The first computes the counterfactual changes in aggregate hours that would be consistent with no change in the relative price of IT to non-IT jobs. The second computes the counterfactual changes when holding the composition of worker quality in 2013 fixed according to 1980.




Figure 12: Comparison of Physical Computer Investment and Employment IT Shares

Notes.—Sources: Current Population Survey and Acemoglu et al. (2014), 1970-2009. The figure plots the share of computer investment relative to total (using Acemoglu et al. (2014)) with the employment share of IT (using the baseline O*NET measure) at a three-digit industry level of aggregation. Observations are weighted by the number of individuals in each three-digit industry cell obtained from the CPS. The figure shows a positive correlation between computer investment and IT workers in the cross-section of industries.




Figure 13: IT Intensity and Real Value of Shipments, AADHP and O*NET Measure

Notes.—Sources: Current Population Survey (1980-2009) and Acemoglu et al. (2014). The figure plots the coefficient on the interaction between the IT intensity measure and year dummies (normalized to 1980 as zero) from regressions of logged real value added and real payroll expenditures on year dummies, three-digit census industry dummies, and the interaction between IT and year dummies. Standard errors are clustered at the industry level and observations are weighted by the industry employment share.




Figure 14: IT Intensity and Real Output, Manufacturing and Services

Notes.—Sources: Current Population Survey and Bureau of Economic Analysis, 1970-2016. Consider the following regression: $y_{jt} = \alpha + \lambda_t + \sum_t \delta^t (IT_j \times \lambda_t) + \eta_j + \epsilon_{jt}$. The figure plots the coefficients associated with the interaction, δ^t , separately for manufacturing and services sectors. The outcome variable is real gross output, which is deflated using 2009 price indices. IT denotes the employment share of IT at a three-digit NAICS level of aggregation between 1970 and 2015. Standard errors are clustered at the industry level and observations are weighted by average employment between 1998-2015.

A Online Appendix (Not for Print)

A.1 Supplement Data Measurement Approach

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 (in-

cluding 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.

For completeness, we also compare our measure with an alternative approach from the Census (see Beckhusen (2016)). While his approach is much more restrictive in its definition of occupations, it shows a stark trend upwards in the number and shares of workers in IT-intensive occupations; see Figure 15. We report these time series merely as a heuristic for comparison.




Figure 15: Information Technology Workers, 1970-2014

Notes.—Sources: See Beckhusen (2016) for calculations. Occupation codes were harmonized to the 2010 census classification.

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 16. 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.

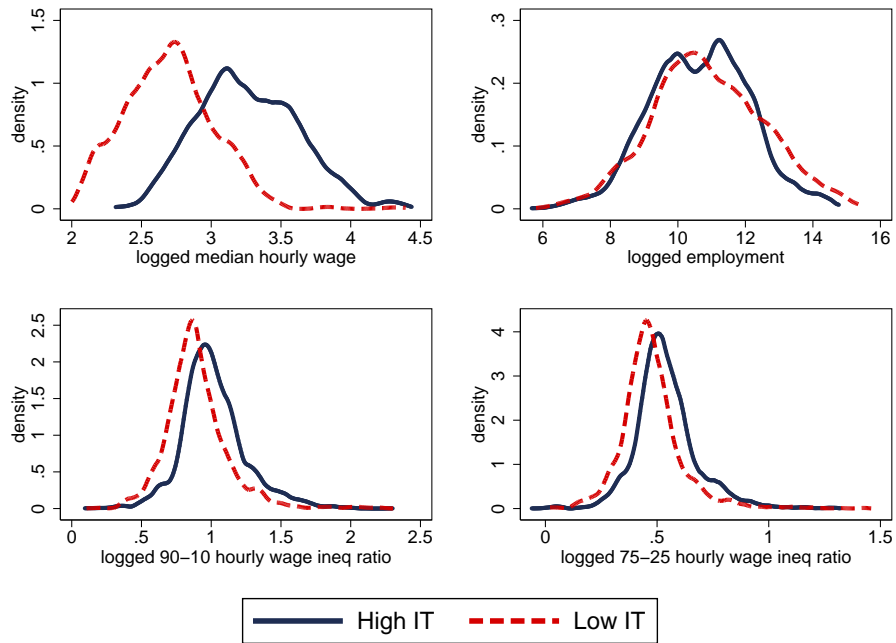


Figure 16: Distribution of Hourly Wages, Employment, and Inequality, by IT Intensity

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 17 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 17: 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.

A.3 Supplement to the Descriptive Evidence and Facts

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, one 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 hedonic regressions of the form

$$y_{ot} = \alpha skill_o + \gamma_t IT_o + \psi_{o'} + \epsilon_{ot} \quad (12)$$

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 12: 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 6 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. We also underscore the fact that these 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).

A.3.2 Time Series Variation in IT 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 7 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

Table 6: Hedonic Skill and Information Technology Regressions

Dep. var. =	logged hourly wage		logged employment		logged 90-10 ratio	
	(1)	(2)	(3)	(4)	(5)	(6)
information technology	.45*** [.02]	.13*** [.02]	-.57*** [.10]	.02 [.09]	.12*** [.02]	.00 [.01]
non-routine, cognitive		.12*** [.02]		-.28*** [.11]		.05*** [.02]
non-routine, non-cognitive		-.02 [.02]		.06 [.08]		-.01 [.02]
routine, cognitive		.04*** [.01]		-.05 [.06]		.01 [.01]
routine, manual		-.02 [.02]		-.14 [.12]		-.09*** [.03]
non-routine, manual		.03 [.03]		.17 [.12]		.04* [.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.

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 18, 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.

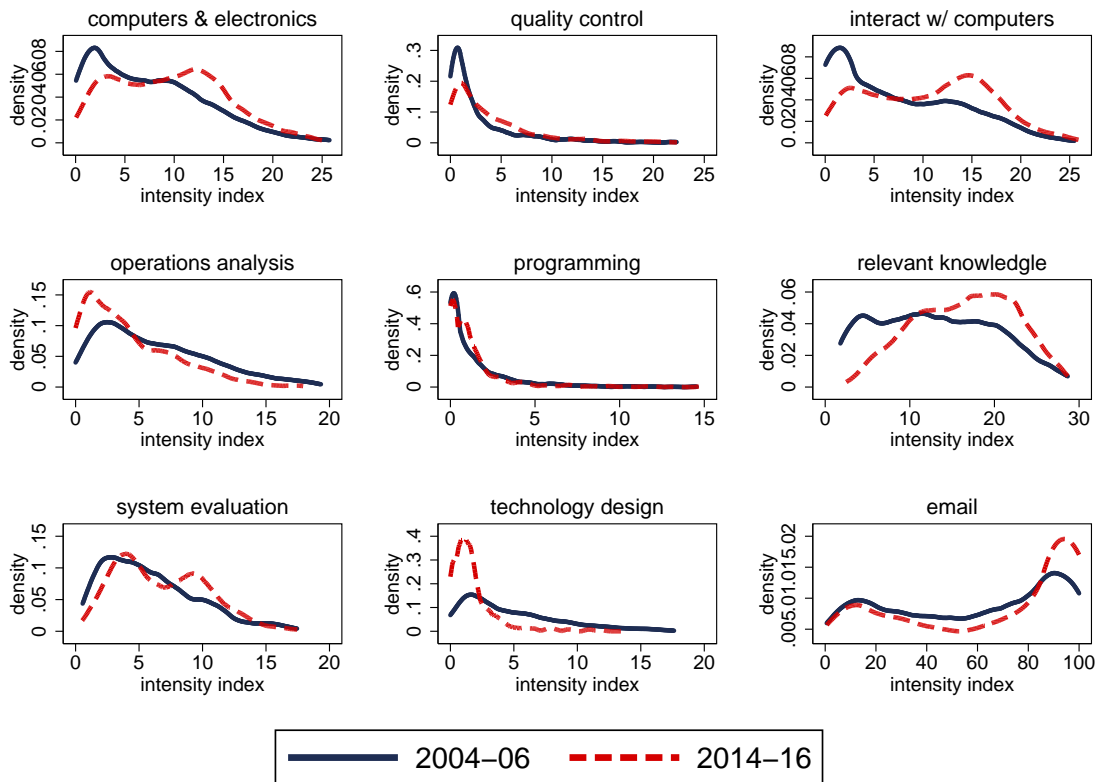


Figure 18: Distribution of Inputs to Information Technology Intensity

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

Table 7: 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.

A.3.3 Robustness of 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 19 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).⁵⁴

⁵⁴The 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 19: 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 \times 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 20 and 21 plot these premia. The crucial observation in Figure 20 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 20: 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 21: 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 22 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.

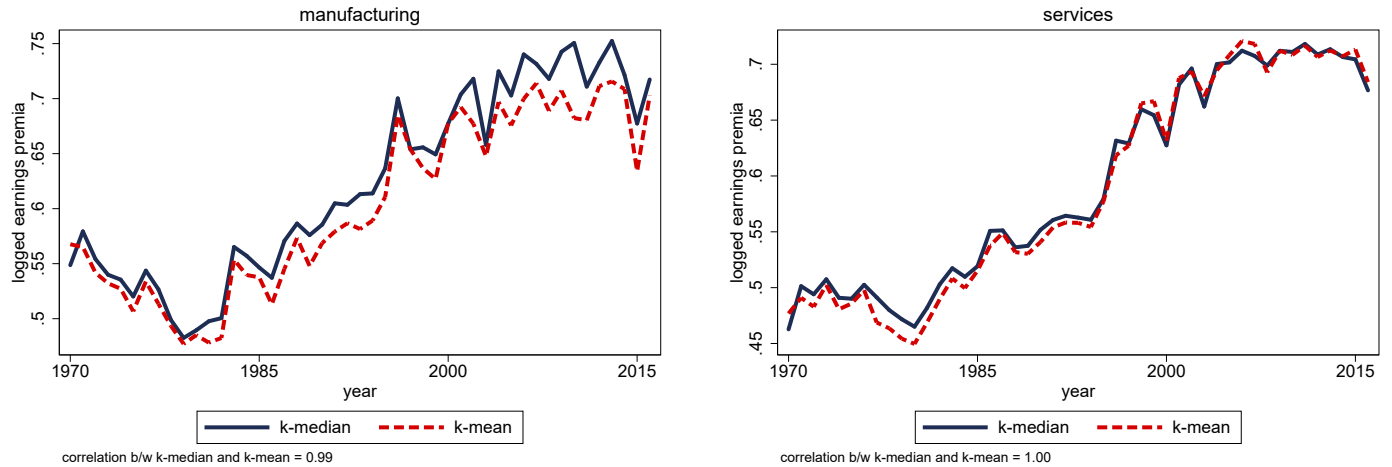


Figure 22: Comparing CPS Earnings Premia with K -medians and K -means

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 23: 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 24 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 24: 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.

A.3.4 The Decline for C/NR Skills 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 C/NR skilled 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 25, 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 25, 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 25: 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 26. 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 26: 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.

A.4 Supplement to Quantitative Model

A.4.1 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):

$$\Delta Y_{g,r,t}(\pi) = \Delta \omega_{g,r,t}^{IT} l_{g,r,t_0}^{IT}(\pi) + \Delta \omega_{g,r,t}^{NIT} l_{g,r,t_0}^{NIT}(\pi) + \Delta \nu_{g,r,t}(\pi) \quad (13)$$

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:

$$\Delta Y_{g,r,t}(\pi) = \Delta \omega_{g,r,t}^{IT,man} l_{g,r,t_0}^{IT,man}(\pi) + \Delta \bar{\omega}_{g,r,t}^{IT,man} (1 - l_{g,r,t_0}^{IT,man}(\pi)) + \Delta \nu_{g,r,t}(\pi) \quad (14)$$

$$\Delta Y_{g,r,t}(\pi) = \Delta \omega_{g,r,t}^{IT, serv} l_{g,r,t_0}^{IT, serv}(\pi) + \Delta \bar{\omega}_{g,r,t}^{IT, serv} (1 - l_{g,r,t_0}^{IT, serv}(\pi)) + \Delta \nu_{g,r,t}(\pi) \quad (15)$$

$$\Delta Y_{g,r,t}(\pi) = \Delta \omega_{g,r,t}^{NIT, man} l_{g,r,t_0}^{NIT, man}(\pi) + \Delta \bar{\omega}_{g,r,t}^{NIT, man} (1 - l_{g,r,t_0}^{NIT, man}(\pi)) + \Delta \nu_{g,r,t}(\pi) \quad (16)$$

$$\Delta Y_{g,r,t}(\pi) = \Delta \omega_{g,r,t}^{NIT, serv} l_{g,r,t_0}^{NIT, serv}(\pi) + \Delta \bar{\omega}_{g,r,t}^{NIT, serv} (1 - l_{g,r,t_0}^{NIT, serv}(\pi)) + \Delta \nu_{g,r,t}(\pi) \quad (17)$$

where $\bar{\omega}_{g,r,t}^{IT, man}$, 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:

$$\Delta Y_{g,r,t}(\pi) = \Delta \bar{\omega}_{g,r,t}^{IT, man} + (\Delta \omega_{g,r,t}^{IT, man} - \Delta \bar{\omega}_{g,r,t}^{IT, man}) l_{g,r,t_0}^{IT, man}(\pi) + \Delta \nu_{g,r,t}(\pi) \quad (18)$$

$$\Delta Y_{g,r,t}(\pi) = \Delta \bar{\omega}_{g,r,t}^{IT, serv} + (\Delta \omega_{g,r,t}^{IT, serv} - \Delta \bar{\omega}_{g,r,t}^{IT, serv}) l_{g,r,t_0}^{IT, serv}(\pi) + \Delta \nu_{g,r,t}(\pi) \quad (19)$$

$$\Delta Y_{g,r,t}(\pi) = \Delta \bar{\omega}_{g,r,t}^{NIT, man} + (\Delta \omega_{g,r,t}^{NIT, man} - \Delta \bar{\omega}_{g,r,t}^{NIT, man}) l_{g,r,t_0}^{NIT, man}(\pi) + \Delta \nu_{g,r,t}(\pi) \quad (20)$$

$$\Delta Y_{g,r,t}(\pi) = \Delta \bar{\omega}_{g,r,t}^{NIT, serv} + (\Delta \omega_{g,r,t}^{NIT, serv} - \Delta \bar{\omega}_{g,r,t}^{NIT, serv}) l_{g,r,t_0}^{NIT, serv}(\pi) + \Delta \nu_{g,r,t}(\pi) \quad (21)$$

A.4.2 Additional Characterizations of Prices & Quantities

We now present several additional statistics implied by our model. Figure 27 plots the price growth premium for IT to non-IT jobs separately by decade. Like before, high IT-intensity jobs have greater price growth in every decade. Price growth is especially large for manufacturing jobs between 1990 and 2000 with a mean of 0.20 for IT jobs and 0.04 for non-IT jobs, whereas price growth is much weaker between 1980 and 1990 with a mean of -0.07 for IT jobs and -0.12 for non-IT jobs. We observe a similar pattern in the services sector. For example, price growth is relatively strong between 1990 and 2000 with 0.17 mean growth among IT jobs and 0.13 for non-IT jobs, whereas it is as low as -0.14 for IT jobs and -0.18 for non-IT jobs between 2000 and 2013. In this sense, the boom in IT price growth was strongest between 1990 and 2000, which coincides with the technology boom (and partial bubble).

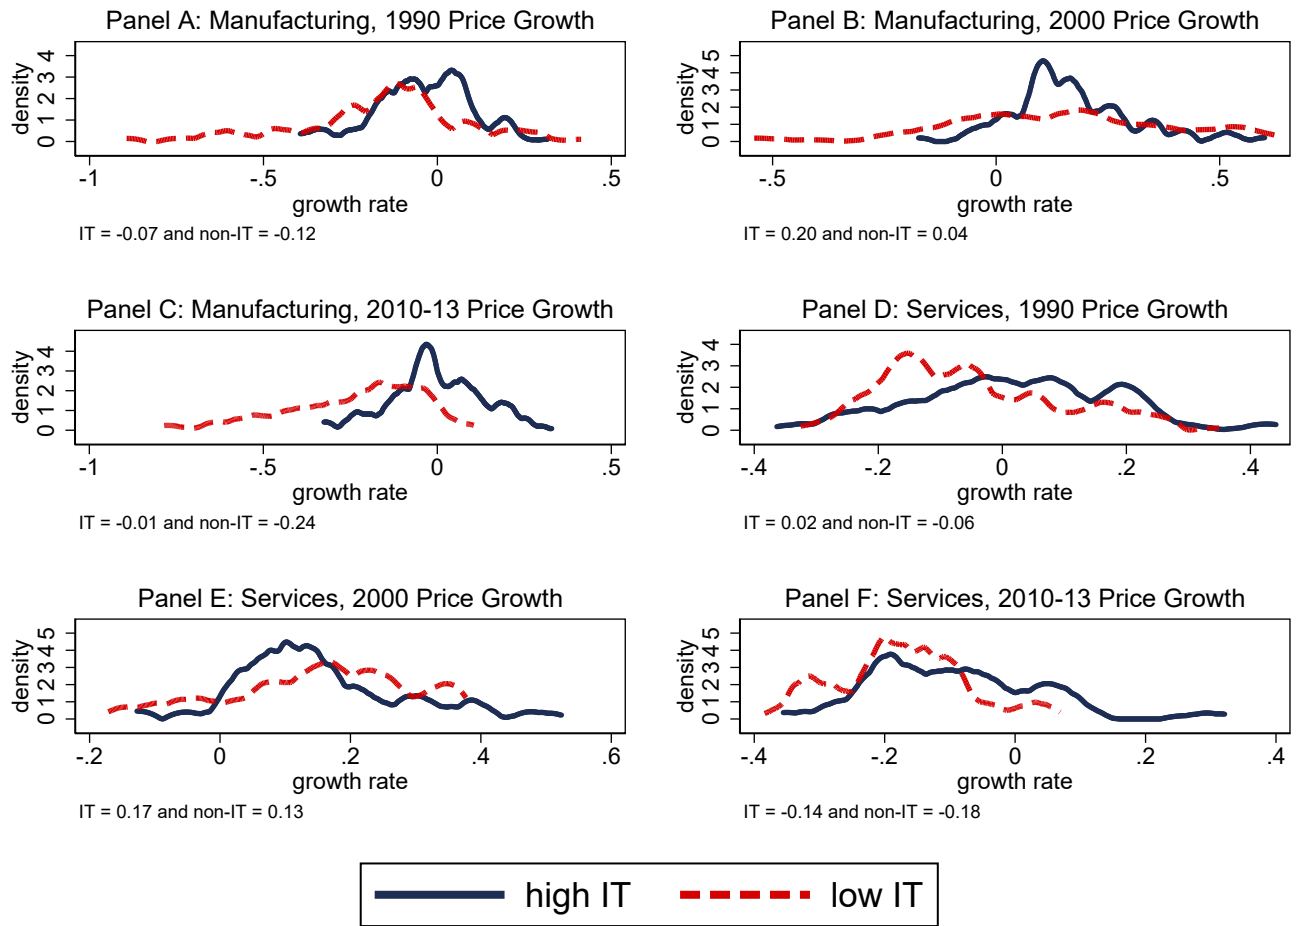


Figure 27: Distribution of Labor Price Growth in IT and non-IT Jobs by Decade

Notes.—Sources: Census Bureau (1990, 2000, and 2013-2015 one-year waves). The figure plots the distributions of the price effects separately by year pooling across group, and region separately for manufacturing / services sectors. Price effects are estimated from the following regression $\Delta Y_{g,r,t}(\pi) = \Delta \omega_{g,r,t}^{IT} + [\Delta \omega_{g,r,t}^{NIT} - \Delta \omega_{g,r,t}^{IT}] t_{g,r,t_0}^{NIT}(\pi) + \mu_{g,r,t} X_{g,r,t}(\pi) + \Delta \nu_{g,r,t}(\pi)$ for $t_0 = 1980$. Controls include: the age distribution (fraction of people between ages 20-29, 30-39, 40-49, 50-65), marital status, race, gender, and dummies on the bottom, middle, and top of the wage distribution. Composition effects are obtained by taking the ratio of the gross wage growth and the price effects.

A.5 Supplement to Revisiting the Solow Paradox

We partition industries into high and low IT based on their median IT intensity collapsed to a three-digit industry level. We subsequently use the BEA data to generate a measure of real value added deflating using price indices normalized to 2009 as the base year, which we use to compare the heterogeneous productivity growth over time in different sectors and in high versus low IT industries. Figure 28 documents these results. We see a stark difference in productivity between not only the services and manufacturing sectors, but also, and more heavily, the high versus low IT sectors. In other words, our alternative measurement strategy suggests that productivity

growth in industries with high concentrations of IT jobs is higher than productivity growth in their counterpart industries.




Figure 28: Real Value Added in Services and Manufacturing with High / Low IT

Notes.—Sources: Current Population Survey and Bureau of Economic Analysis, 1970-2016. The figure plots the smoothed real value added (deflating using 2009 price indices) in high versus low IT in manufacturing and services sectors. We use the CPS to produce IT employment shares, which we use to classify an industry as high versus low IT based on whether its standardized z -score is above the median or not.

Some of the literature on the Solow paradox and structural transformation highlights the role of multi-factor productivity (or total factor productivity, TFP). Our model focuses on task-specific productivity changes in IT and non-IT jobs, hence it is not suited to examine the properties of aggregate multi-factor productivity. Nonetheless, in what follows we compare the task-specific productivity changes implied by our model to measures of TFP based on calculations from national accounts data. Before discussing these comparisons, we describe our exercise. Using the Bureau of Economic Analysis regional data, we produce employment-weighted state \times industry \times year data for employment and real output between 1982 and 2015. Since there is no publicly available capital expenditure data, we estimate TFP in two alternative ways. First, we draw on the national time series of fixed private capital assets and equipment. Second, we draw on measures of state \times industry \times year capital income.⁵⁵ Using both these measures of capital, we produce two different

⁵⁵Although capital income is an imperfect proxy for actual capital assets, it is a flow that is driven (at least in

TFP proxies by regressing logged real output on logged employment, logged capital measures, obtaining sector-specific TFP processes.

We produce a model-based state-year measure of IT task productive growth (relative to non-IT tasks) by first regressing the growth rate of our IT price premium on the relative IT to non-IT task input changes, and then taking the residual. 29 plots the latter measure against the two measures of TFP from the state-level BEA data, weighting observations by the employment share of each state relative to the total U.S. population in 1980. As displayed in each panel, growth in relative IT-task productivity is positively associated to higher TFP at the state-year level. Estimates for the services sector are more compressed than those in the manufacturing sector, which contains more extreme observations. However, the fact that our model-based estimate of relative IT productivity growth lines up with higher TFP suggests that part of residual TFP may be related to the evolution of task-specific productivity growth.

part) by local ownership of capital assets.

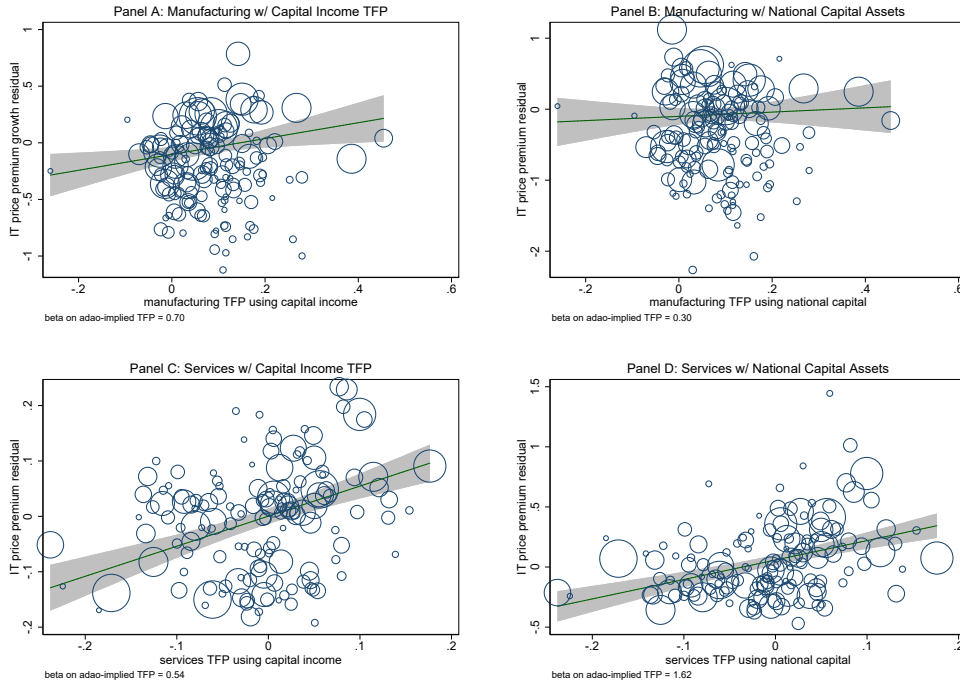


Figure 29: Comparison of Model Approximation of TFP and Actual TFP

Notes.—Sources: Census Bureau, O*NET, Bureau of Economic Analysis, 2000-2014. Each plot displays IT-task relative productivity growth against a measure of TFP, by industry. After collapsing the data to employment-weighted manufacturing and services sectors, we measure TFP by regressing logged real output on logged employment, and a logged proxy for capital; we take the residual as an estimate of total factor productivity (TFP) measure. Because of the lack of state-by-industry data on productive capital, we measure capital using either (i) state \times industry capital income or (ii) using national private capital assets on equipment. In both cases, we estimate sectoral-specific TFP processes. Moreover, we measure state-year growth in IT-tasks productivity (relative to non-IT tasks) by regressing the IT price premium growth on the IT tasks input premium growth, in a specification with no intercept. Then, we take the residual as our measure of yearly task-productivity growth. That is, we estimate the following industry-specific regression:

$$\Delta_t (\omega_i^{IT} - \omega_i^{NIT}) = (\nu_i - 1) \Delta_t \ln \left(L_i^{IT} / L_i^{NIT} \right) + \underbrace{\Delta_t \ln (\alpha_i / (1 - \alpha_i)) + \nu_i \Delta_t \left(\varepsilon_i^{IT} / \varepsilon_i^{NIT} \right)}_{\text{IT task productivity growth (relative to non-IT tasks)}}.$$

Observations are weighted by the state employment share of 1980.