

New Technologies and the Labor Market

Enghin Atalay Phai Phongthiengtham Sebastian Sotelo Daniel Tannenbaum*

November 5, 2017

Abstract

We examine how the introduction of information and communication technologies (ICTs) affects the tasks that workers perform in their jobs, workers' occupational choices, and the wages that workers of different skill levels earn. Using the text from help wanted ads published between 1960 and 2000, we construct a data set that measures the adoption of 40 ICTs. We find that new technologies are associated with an increase in nonroutine analytic tasks, and a decrease in nonroutine interactive, routine cognitive, and manual tasks. We embed these interactions in a quantitative model of worker sorting across occupations and technology adoption, and evaluate the impact of the arrival of ICTs on the aggregate demand for worker-performed tasks and on earnings inequality. Through the lens of the model, the arrival of ICTs generates a large shift away from routine tasks, and, consequently, an increase in inequality since (i) high wage workers tend to adopt ICTs and (ii) relative to high wage workers, low wage workers have a comparative advantage in performing routine tasks. **JEL Codes: J24, M51, O33**

*Atalay and Phongthiengtham: Department of Economics, University of Wisconsin-Madison. Sotelo: Department of Economics, University of Michigan-Ann Arbor. Tannenbaum: Department of Economics, University of Nebraska-Lincoln. We acknowledge financial support from the Washington Center for Equitable Growth.

1 Introduction

Enabled by increasingly powerful computers and the proliferation of new, ever more capable software, the fraction of workers’ time spent using information and communication technologies (ICTs) has increased considerably over the last half century.¹ In this project, we quantify the impact of 40 individual and recognizable ICTs on the aggregate demand for routine and nonroutine tasks, on the allocation of workers across occupations, and on earnings inequality.

We start by constructing a data set tracking the adoption rates of 40 ICTs across occupations and years. We assemble this data set through a text analysis of 6.6 million job vacancy ads appearing in newspapers between 1960 and 2000 in the *Boston Globe*, *New York Times*, and *Wall Street Journal*.² We extract information about jobs’ ICT use and task content, as measured by their appearance in the text of job postings. In addition, we use the job titles posted in the ads to recover SOC codes, allowing us to link our processed data to economy-wide occupation data in the U.S. Census.

The technologies we study constitute a wide set, ranging from office software (including Lotus 123, Word Perfect, Microsoft Word, Excel, PowerPoint), enterprise programming languages (Electronic Data Processing, Sybase), general-purpose programming languages (COBOL, Fortran, Java), and hardware (UNIVAC, IBM 360, IBM 370), among others. With this data set, we document rich interactions of individual ICTs and the task content of individual occupations. One of the strengths of the data is being able to measure ICT adoption separately by technology type, and indeed we find substantial heterogeneity in the impact of individual ICTs. We show that, for the most part, job ads that mention a new technology tend to also mention nonroutine analytic tasks more frequently, while mentioning other tasks less frequently — this provides preliminary evidence that new technologies are complementary with particular tasks.³ An important exception is office software, which is more likely to appear alongside words associated with nonroutine interactive tasks.

Since our data set includes a wide range of occupations and technologies, we can speak

¹Nordhaus (2007) estimates that, between 1960 and 1999, the total cost of a standardized set of computations fell by between 30 and 75 percent annually, a rapid rate of change that far outpaced earlier historical periods.

²We introduced part of this data set in our earlier paper, namely the task measurement and the job title-to-SOC mapping (Atalay, Phongthientham, Sotelo, and Tannenbaum, 2017). Here, we build these data further by extracting information about job-specific technology adoption. In Atalay, Phongthientham, Sotelo, and Tannenbaum (2017), we use the text of job vacancy ads to explore trends in occupations’ task content, showing that within-occupation changes in the tasks workers perform are at least as large as the changes that happen between occupations.

³Building on a mapping between survey question titles and task categories introduced by Spitz-Oener (2006), we have identified words that represent nonroutine (analytic, interactive, and manual) and routine (cognitive and manual) tasks.

directly to the macroeconomic implications of changes in ICT prices while maintaining a detailed analysis of individual occupations. Informed by our micro estimates on the relationship between the tasks that workers perform and the technologies they use on the job, we build a quantitative model of occupational sorting and technology adoption. In the model, workers sort into occupations based on their comparative advantage. They also choose which ICT to adopt, if any, based on the price of each piece of technology and the technology’s complementarity with the tasks involved in their occupation. Within the model, the availability of a new technology — which we model as a decline in the technology’s price — alters the types of tasks workers perform in their occupation.

To explore the implications of new technologies on the labor market, we consider three sets of counterfactual exercises. These exercises investigate the effects of three groups of technologies: i) Unix, ii) the Microsoft Office suite: Microsoft Excel, Microsoft PowerPoint, and Microsoft Word, and iii) all 40 of the technologies in our sample. In each of the counterfactual exercises, we quantify the impact of the new technologies on occupations’ overall task content, workers’ sorting across occupations, and economy-wide income inequality.

One of our main findings is that new technologies result in an increase in occupations’ nonroutine analytic task content (relative to other tasks). As we have documented elsewhere (Atalay, Phongthientham, Sotelo, and Tannenbaum, 2017) and confirm again here, workers with observable characteristics indicating high skill levels (experienced and highly educated workers) have a comparative advantage in producing nonroutine analytic tasks. Because new technologies increase the demand for worker-performed tasks in which high-skilled workers have a comparative advantage, the introduction of ICTs has (for the most part) led to an increase in income inequality. Overall, in a counterfactual economy in which our ICT technologies were never introduced, earnings would have been 11.5 percent lower for the median worker; the College-High School skill premium would have been 3.2 percentage points lower.⁴ Unlike the other technologies in our data, Microsoft Office technologies are only weakly correlated with nonroutine analytic tasks. Concomitantly, the impact of Microsoft Office software has been to decrease the skill premium, and income inequality. However, the effects of these technologies are small. Individual technologies whose use is concentrated in a few high-earning occupations, such as Unix, tend to modestly increase inequality.

This paper relates to a rich literature in labor economics exploring the implications of technological change for skill prices and the wage distribution (Katz and Murphy, 1992; Juhn, Murphy, and Pierce, 1993; Berman, Bound, and Machin, 1998; Krusell, Ohanian, Rios-Rull, and Violante, 2000). More recent work has argued that information technology complements high-skilled workers performing abstract tasks and substitutes for middle-skilled workers

⁴Between 1960 and 2000, the College-High School skill premium increased by 23 percentage points.

performing routine tasks (Autor, Levy, and Murnane, 2003; Goos and Manning, 2007; Autor, Katz, and Kearney, 2005; Acemoglu and Autor, 2011).

Our paper adopts the task approach as well, and examines how new technologies complement (or substitute for) the types of tasks that workers of different skill groups perform, finding that ICTs tend to substitute for routine tasks (especially routine manual tasks) which are disproportionately performed by low skill workers. In turn ICTs allow high skill workers to focus on the activities in which they are the most productive, which in our model is the essence of the complementarity. One of our contributions to this literature is to measure both technological change and the task content of occupations directly, over a period of immense technological change.

Our paper relates to a second literature that measures directly the adoption of specific technologies and its effect on wages and the demand for skills. These include studies of the effect of computer adoption (e.g., Krueger, 1993; Entorf and Kramarz, 1998; Haisken-DeNew and Schmidt, 1999; Autor, Katz, and Krueger, 1998) or the introduction of broadband internet (e.g., Brynjolfsson and Hitt, 2003; Akerman, Gaarder, and Mogstad, 2015) on worker productivity and wages.⁵ Also exploiting text descriptions of occupations, Michaels, Rauch, and Redding (2016) provide evidence that, since 1880, new technologies that enhance human interaction have reshaped the spatial distribution of economic activity. Focusing on a more recent technological revolution, Burstein, Morales, and Vogel (2015) document how the diffusion of computing technologies has contributed to the rise of inequality in the U.S. Our paper builds on this literature by introducing a rich data set measuring the adoption of ICTs at the job vacancy level.

The rest of the paper is organized as follows. Section 2 of the paper introduces our new data set. Section 3 provides direct evidence on the interaction between individual ICT adoption and task contents. Section 4 takes our micro estimates and uses a quantitative model to study the aggregate impact of ICTs. Section 5 concludes.

2 A New Data Set Measuring ICT Adoption

The construction of this new data set builds on our previous work with newspaper help wanted ads (Atalay, Phongthientham, Sotelo, and Tannenbaum, 2017). In that paper, we showed how to transform the text of help wanted ads into time-varying measures of the task content of occupations. In this paper, we turn to previously unexamined content of the ad:

⁵Additional investigations of technology-driven reorganizations within specific firms or industries include Levy and Murnane (1996)'s study of a U.S. bank and Bartel, Ichniowski, and Shaw (2007)'s study of the steel valve industry.

mentions of ICTs.

Our main data set is built from the universe of job vacancies published in three major metropolitan newspapers — the *Boston Globe*, *New York Times*, and *Wall Street Journal* — which we purchased from ProQuest. We use the text contained in each vacancy to measure the tasks that will be performed on the job and the computer and information technologies that will be used on the job. Our sample period spans 1960 to 2000.

The original newspapers were digitized by ProQuest using an Optical Character Recognition (OCR) technology. We briefly describe the steps we take to transform this digitized text into a structured database. First, the raw text does not distinguish between job ads and other types of advertisements. Hence, we apply a machine learning algorithm to determine which pages of advertisements are job ads. Figure 1 presents a portion of a page of job ads. This snippet of text refers to three job ads, first for a Software Engineer position, then a Senior Systems Engineer position, and finally for a Software Engineer position. Within this page of ads, we first determine the boundaries of each individual advertisement (where, e.g., the Software Engineer ad ends and the Senior Systems Engineer ad begins) and the job’s title. We then extract, from each advertisement, words that refer to skill requirements, tasks the new hire is expected to perform, and technologies that will be used in the job. So that we may link our text-based data to occupation-level variables in the Decennial Census, including wages, education, and demographics, our procedure also finds the SOC code corresponding to each job title (for example 151132 for the “Software Engineers” job title.)⁶

We adopt the mapping of words to task categories based on [Spitz-Oener \(2006\)](#). The five tasks are *nonroutine analytic*, *nonroutine interactive*, *nonroutine manual*, *routine cognitive*, and *routine analytic*.⁷ Because we do not want our analysis to be sensitive to trends in language — either word usage or meaning — we adopt a machine-learning algorithm called the *continuous bag of words* to define a set of synonyms for each of our task-related words. The idea is that words that commonly share surrounding words in the text are likely to

⁶For additional details on the steps mentioned here, see [Atalay, Phongthientham, Sotelo, and Tannenbaum \(2017\)](#). In that paper we also address issues regarding the representativeness of newspaper ads, and the validity of task measures extracted from the text. Our data set, including information on occupations’ task and technology mentions are available at http://ssc.wisc.edu/~eatalay/occupation_data. In addition, on that website we list the full list of words and phrases we associate with each task and technology.

⁷We use the mapping of words to tasks as described in Footnote 15 of [Atalay, Phongthientham, Sotelo, and Tannenbaum \(2017\)](#) and for convenience listed again here: 1) nonroutine analytic: analyze, analyzing, design, designing, devising rule, evaluate, evaluating, interpreting rule, plan, planning, research, researching, sketch, sketching; 2) nonroutine interactive: advertise, advertising, advise, advising, buying, coordinate, coordinating, entertain, entertaining, lobby, lobbying, managing, negotiate, negotiating, organize, organizing, presentation, presentations, presenting, purchase, sell, selling, teaching; 3) nonroutine manual: accommodate, accommodating, accommodation, renovate, renovating, repair, repairing, restore, restoring, service, serving; 4) routine cognitive: bookkeeping, calculate, calculating, correcting, corrections, measurement, measuring; 5) routine manual: control, controlling, equip, equipment, equipping, operate, operating.

Figure 1: Unprocessed Text from the *New York Times*, January 12, 1997, Display Ad #87

SOFTWARE ENGINEERS - Modal Software Develop air-to-surface modal software, including design, code, unit test, integration and test, and documentation. Requires 5+ years software engineering experience with a BSEE/CS or Computer Engineering. Software development for real-time, multi-tasking/multi-processor, embedded systems experience a must. 3+ years C programming experience in a Unix environment and familiarity with modern software design methodologies essential. Knowledge of radar design principles a plus. Joint STARS The premiere ground surveillance system for the U.S. and allied forces. The DoD has authorized the full production of Joint STARS. In addition, significant activity on Joint STARS upgrades is underway. SENIOR SYSTEMS ENGINEERS Design and develop advanced, high-resolution radar imaging systems, including ultra-high resolution SAR and Moving Target Imaging Systems in real-time or near real-time environments. Represent the engineering organization to senior technical management, potential partners and customers in industry and government; plan/coordinate R&D program activities; lead a team of hardware/software/systems engineers; develop and test complex signal processing modes and algorithms in a workstation environment; support development with analyses, reports, documentation and technical guidance. Requires an MS or PhD in Engineering, Physics or Mathematics with experience in specification, Imaging and testing of Advanced Coherent Radar High-Resolution Must have strong math, physics and signal processing skills, C/C++ and FORTRAN programming expertise, plus familiarity with workstations and analytical tools such as The following require knowledge of emulators, debuggers, and logic analyzers. Knowledge of Ada, Unix, VxWorks, DigitalAlpha Processor and assembly language desirable. Radar systems experience plus. SOFTWARE ENGINEERS Define requirements and develop software for RCU or Intel microprocessor-based RSEs. Help define software requirements for LRU ECPs and the Contractor Logistics software program, including design, code, integration and test, and documentation. BSCS/EE preferred with 3-5 years real-time software development experience using Ada and/or FORTRAN programming languages. U IS- * SOFTWARE

Notes: The figure presents text from three vacancy postings in a page of display ads in the *New York Times*.

share the same meaning. For example, one of the words corresponding to the nonroutine analytic task is *researching*. The continuous bag of words method uses the text itself to find synonyms of *researching*; these synonyms include *interpreting*, *investigating*, *reviewing*, etc. In our analysis, we include the union of these synonyms as words mapping to the nonroutine analytic task, which limits the sensitivity of our analysis to variations in diction over time. In addition to tasks, we extract 40 different pieces of technology based on word appearances in the text.

Figure 2 presents the output of our text processing algorithm. This algorithm has been able to correctly identify the boundaries between the three job ads, as well as the positions of each of the three job titles. However, since the initial text contained, “Sofware,” a misspelled version of “Software,” we have incorrectly identified the first job ad as referring to an engineering position. Our algorithm identifies nine mentions of nonroutine analytic tasks: “design” and “plan” were words in Spitz-Oener (2006)’s definitions of nonroutine task related words. In addition, our continuous bag of words model identifies “develop” and “define” as referring to nonroutine analytic tasks. We also identify one mention of a nonroutine interactive task — based on the word “coordinate” — and three mentions of software: two mentions of Unix and one of Fortran. While our data set contains some measurement error in identifying each job ad’s title and task and technology content, there is considerable information within the text that can be usefully extracted.

Table 1 lists the technologies in our sample together with information on their timing of adoption, as measured by the number of mentions in job ads, and the year the technology was introduced.⁸ The columns titled “First Year” and “Last Year” list the first and last years within the 1960 to 2000 period in which the frequency of technology mentions is at least one-third of the mentions in the year when the technology is mentioned most frequently. Using this one-third cutoff, the lag between technology introduction and technology adoption (i.e. the difference between the “Introduction” and the “First Year” column) is 8 years on average. The final column lists the overall frequency of mentions, across the 6.6 million job ads in our data set, of each piece of technology.

Figure 3 plots the trends in technology mentions in our data set. Over the sample period, there is a broad increase in the frequency with which employers mention technologies, from less than 0.02 mentions per ad in the beginning of the sample to 0.20 mentions by 2000. While there is a broad increase in technology adoption rates throughout the sample, certain technologies have faded from use over time. The right panel of Figure 4 documents adoption rates for each of the 40 technologies in our sample, with seven of these highlighted. Certain technologies which were prevalent in the 1960s and 1970s — including Electronic

⁸We obtained the year of introduction from the Wikipedia page of each technology.

Figure 2: Processed Text from the *New York Times*, January 12, 1997, Display Ad #87

engineers|- modal software **develop** air-to-surface modal software , including **design** , code , unit test , integration and test , and documentation . requires 5+ years software engineering experience with a b see cs or computer engineering . software development for real-time , multitasking multiprocessor , embedded systems experience a must . 3+ years c programming experience in a **UNIX** environment and familiarity with modern software design methodologies essential . knowledge of radar design principles a plus . joint stars the premiere ground surveillance system far the u . s . and allied forces . the DOD has authorized the full production of joint stars . in addition , significant activity on joint stars upgrades is underway .

senior system engineer| **design** and **develop** advanced , high-resolution radar imaging systems , including ultra-high resolution sear and moving target imaging systems in real-time or near real-time environments . represent the engineering organization ta senior technical management , potential partners and customers in industry and government ; **plan** **coordinate** ; d program activities ; lead a team of hardware soared systems engineers ; **develop** and test complex signal processing modes and algorithms in a workstation environment ; support development with analysis , reports , documentation and technical guidance . requires an ms or PhD in engineering , physics or mathematics with experience in specification , imaging ans and testing of advanced coherent radar high-resolution must have strong math , physics and signal processing skills , c c and , an programming expertise , plus familiarity with workstations and analytical tools such as the following require knowledge of emulators , debuggers , and logic Ana . knowledge of Ada , **UNIX** , vxworks , digital alpha processor and assembly language desirable . radar systems experience plus.

software engineers|define requirements and **develop** software far r cu or Intel microprocessor-based rs es . help **define** software requirements far lr u e cps and the contractor logistics software program , including **design** , code , integration and test , and documentation . bscs ee preferred with 3-5 years real-time software development experience using Ada and or **FORTRAN** programming languages . u is- software

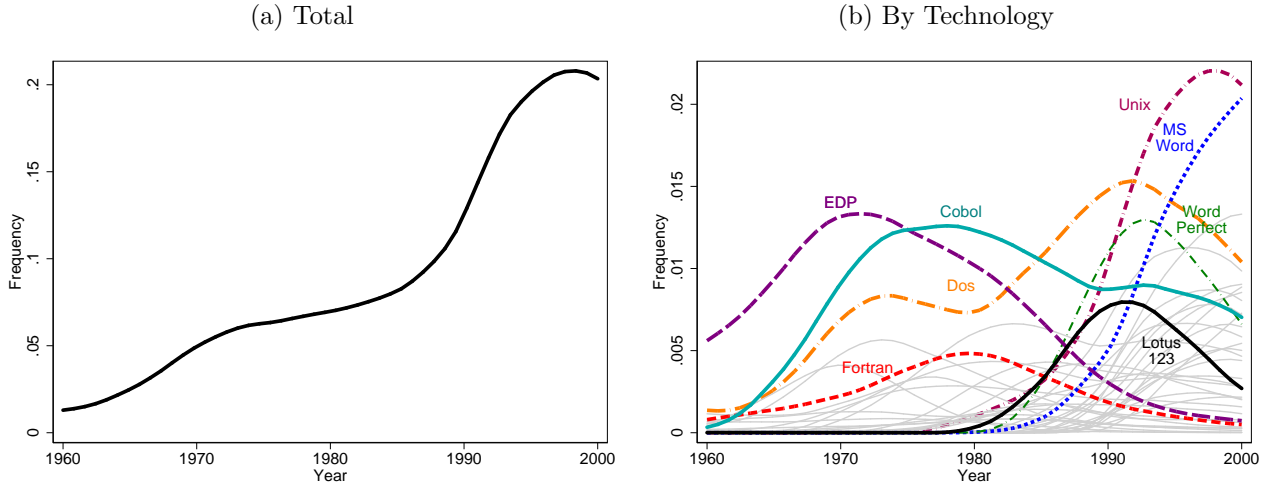
Notes: The figure presents text from three vacancy postings in a page of display ads in the *New York Times*. Highlighted text, within a rectangle, refers to a mention of a nonroutine analytic task. Highlighted text, within an oval, refers to a mention of a nonroutine interactive task. Text within a rectangle refers to a technology mention. Within these three ads, there are zero mentions of nonroutine manual, routine cognitive, and routine manual tasks.

Table 1: Technologies

Technology	Introduction	First Year	Last Year	Frequency (%)
APL	1957	1961	1998	0.05
BAL	1964	1968	1983	0.30
CAD	1966	1981	1985	0.04
CICS	1968	1974	1998	0.30
COBOL	1959	1968	1998	0.83
C++	1983	1993	1999	0.02
DB2	1983	1989	1998	0.08
DOS	1966	1969	1999	0.72
EDP	1960	1963	1986	0.91
Fortran	1957	1965	1992	0.27
Foxpro	1989	1992	1998	0.02
HTML	1993	1996	≥2000	0.04
IBM 360	1964	1965	1974	0.18
IBM 370	1970	1972	1982	0.13
IBM RPG	1959	1970	1992	0.04
IMS	1966	1960	≥2000	0.26
Java	1995	1996	≥2000	0.08
JCL	1964	1969	1998	0.17
LAN	1970	1990	1998	0.19
Lotus 123	1983	1987	1995	0.12
Lotus Notes	1989	1994	1998	0.03
MS Excel	1987	1993	≥2000	0.04
MS PowerPoint	1990	1995	≥2000	0.05
MS Word	1983	1993	1999	0.16
MVS	1974	1979	1998	0.15
Novell	1983	1994	1998	0.07
Oracle	1977	1995	1999	0.10
Pascal	1970	1982	1990	0.05
Quark	1987	1992	1999	0.07
SQL	1986	1993	1999	0.08
Sybase	1984	1995	1997	0.05
TCP	1974	1994	1999	0.03
TSO	1971	1977	1998	0.06
Univac	1951	1960	1984	0.06
Unix	1971	1992	1999	0.22
Vax	1977	1982	1998	0.11
VisualBasic	1991	1995	1998	0.04
VMS	1977	1985	1996	0.07
VSAM	1970	1982	1998	0.05
Word Perfect	1979	1988	1998	0.15

Notes: This table lists the 40 technologies in our sample. The “First Year” and “Last Year” columns report the first year and last year at which the frequency of technology mentions was at least one-third of the frequency of the year with the maximum mention frequency (number of technology mentions per job ad). The ≥ 2000 symbol indicates that the technology was still in broad use at the end of the sample period.

Figure 3: Mentions of Technologies



Notes: This plot gives the smoothed frequency with which job ads mention our set of technologies. The left panel depicts the sum frequency of all 40 technologies. The right panel depicts the frequencies of each of the 40 technologies separately, eight which are highlighted in thick dark lines and thirty-two which are depicted by thin, light gray lines.

Data Processing (EDP) and COBOL — have declined in usage. Other technologies — Word Perfect and Lotus 123 — quickly increased and then decreased in newspaper mentions.

In Figure 4, we examine the heterogeneity across occupations in their adoption rates. Here, we plot the frequency of job ads which mention each technology, across 4-digit SOC groups of four different technologies: Fortran, Unix, Word Perfect, and Microsoft Word. Each plot indicates with a vertical line the year of release of the technology to the public. These plots suggest several new facts. First, technological adoption is uneven across occupations, occurring at different times and to different degrees. For instance Fortran is quickly adopted by Computer Programmers, while the adoption by Engineers lags behind and is more limited. Second, for technologies that perform the same function, such as Word Perfect and Microsoft Word, the figures suggest dramatic substitution between technologies. Lastly, we see that office software is adopted widely across diverse occupations, whereas other types of software, such as CAD, are adopted more narrowly. Finally, between the time of release to the public and the peak of adoption, adoption rates increase first quickly and then slowly. This pattern is consistent with the S-shaped documented in the diffusion of many technologies (e.g., [Griliches, 1957](#); [Gort and Klepper, 1982](#)). While we do not offer a theory of the pattern of adoption of new technologies for each occupation, we will exploit the time variation in adoption rates to gauge their impact on the macroeconomy.

While our data set is new in its measurement of the adoption of a large number ICTs across time and occupations, there are existing data sets — O*NET and the October CPS

— which measure ICT usage across occupations. O*NET contains information on multiple ICTs, while the October CPS tracks computer usage rates across a number of years. In Appendix A, we document that our technology measures align with those in these two existing data sets.

3 Task and Technology Complementarity

This section documents empirically how new technologies interact with occupational task content. We investigate the relationship between mentions of the technologies that employees use on the job and the tasks that these employees are expected to perform. This estimated relationship will be a critical input into the equilibrium model in the following section.

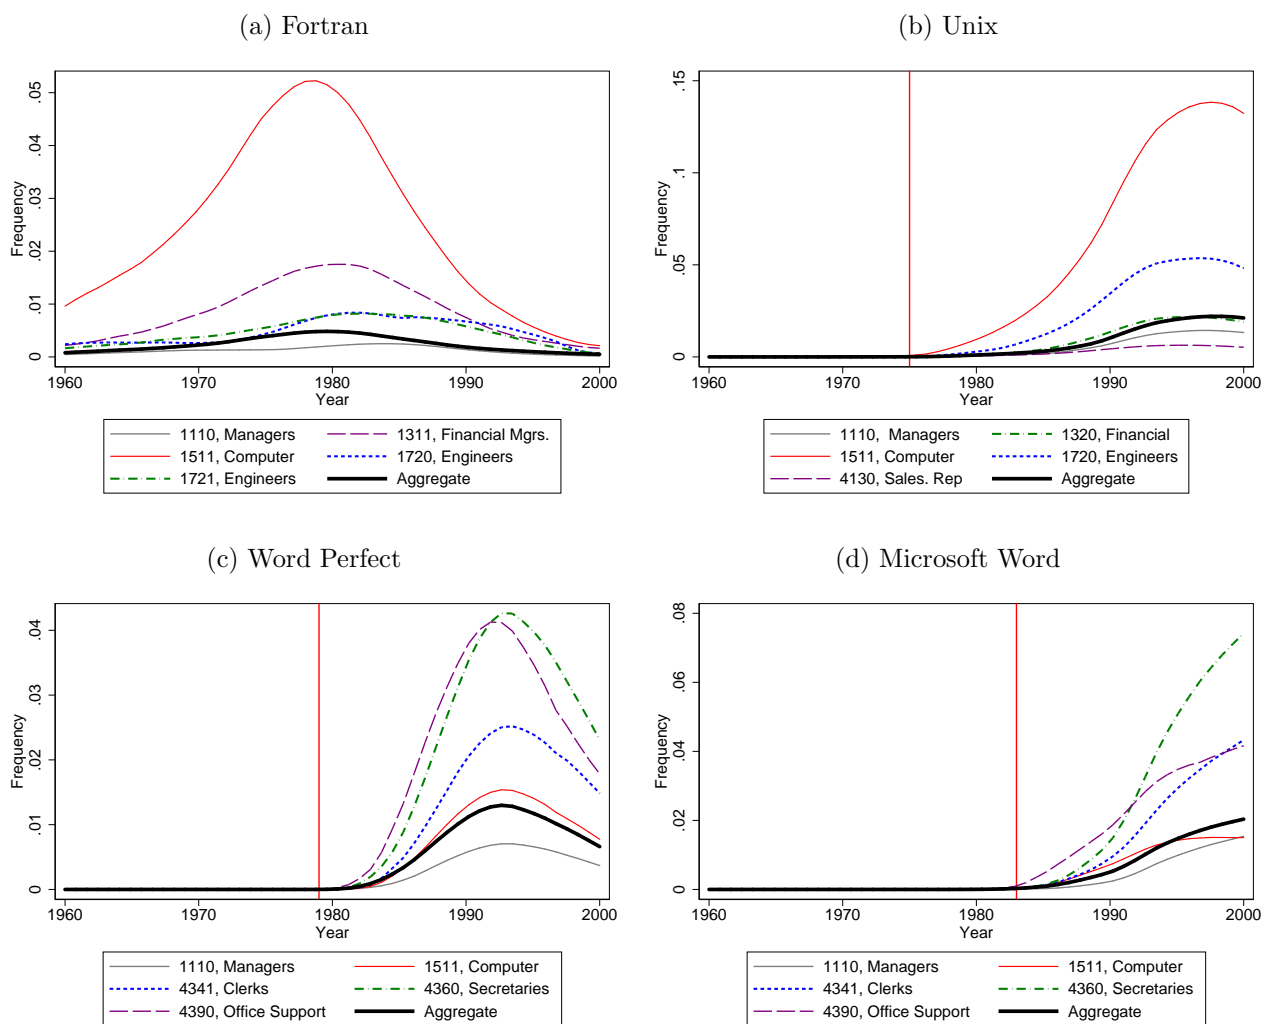
As new technologies are introduced and developed, the implicit price of technology adoption falls. As the price falls, in certain jobs employers will find it profitable to have their employees adopt the new technology. Based on the applicability of the new technology, jobs will differ in the extent to which adoption occurs, even if the price of adopting the technology is the same across occupations. Exploiting this temporal and occupational variation in the extent to which workers adopt technologies, we estimate the following equation:

$$\text{task}_{ajt}^h = \beta_{hk} \cdot \text{tech}_{ajkt} + f_h(\text{words}_{ajt}) + \iota_{jh} + \iota_{th} + \epsilon_{ahjkt} \quad (1)$$

In Equation 1, h refers to one of five potential routine and nonroutine task categories; tech_{ajkt} gives the number of mentions of a particular technology k in individual job ad a , published in year t for an occupation j ; ι_{jh} and ι_{th} refer to occupation and year fixed effects, respectively; and $f_h(\text{words}_{ajt})$ is a quartic polynomial controlling for the number of words in the ad, since the word count varies across ads. We run the regressions characterized by Equation 1 separately for each technology k and task h . The occupation fixed effects and year-fixed effects respectively control for occupation-specific differences in the frequency of task mentions and economy-wide trends in the tasks that workers perform unrelated to technology adoption.

In interpreting the regression coefficient, β_{hk} , a key challenge is that technology adoption may be correlated with unobserved attributes of the job (Athey and Stern, 1998). For instance, within a particular 4-digit SOC (e.g., SOC 1721–Engineers) certain jobs (e.g., Mechanical Engineers relative to Industrial Engineers) potentially could be both more likely to adopt a new technology and more intensive in nonroutine analytic tasks. In other words, instead of concluding that ICT adoption and nonroutine analytic tasks are complements,

Figure 4: Mentions of Technologies across Occupations



Notes: This plot gives the smoothed frequency with which job ads in different occupations mention technologies. Each plot depicts the frequencies of technology mentions for five of the top (largest and most-intensively adopting of new technologies) Standard Occupation Classification (SOC) occupations along with the economy-wide average frequency of technology mentions. The red vertical lines depict the date the technology was introduced. (Fortran was introduced in 1957, right before the beginning of our sample.)

one may conclude that jobs that are high in nonroutine analytic tasks tend to adopt the technology. This distinction is important for the interpretation of the empirical results, and we explore it in Appendix B. There, we re-estimate the regressions specified by Equation 1 with increasingly detailed job-level fixed effects, showing that the relationship between ICT adoption and task content does not change with these more detailed controls.⁹

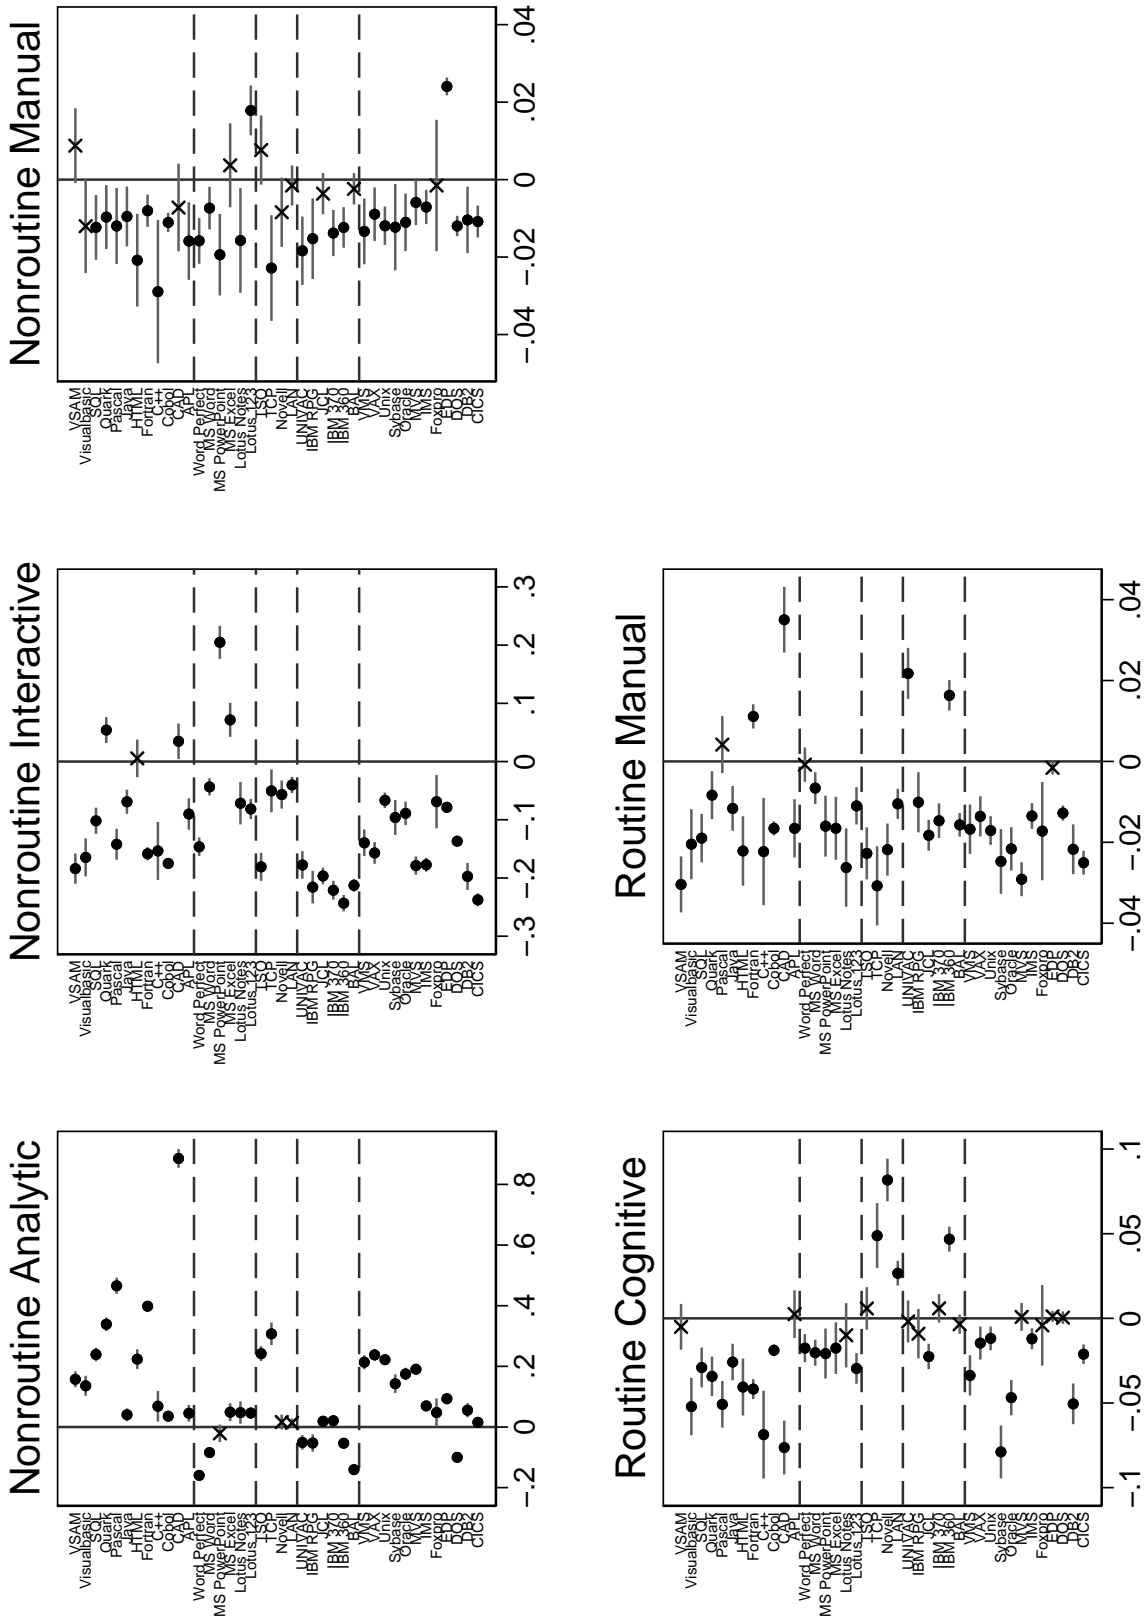
Figure 5 presents the estimates of β_{hk} for each task-technology pair. Within each panel, technologies are grouped according to their type, with database management systems first, then office software, networking software/hardware third, other hardware fourth, and general purpose software fifth. According to the top-left panel, the relationship between nonroutine analytic task mentions and technology mentions is increasing for database management systems, networking software/hardware, and general purpose software. Among the 40 technologies in our sample, the median effect of an additional technology-related mention is an additional 0.05 nonroutine analytic task mentions per job ad. On the other hand, technology mentions and task mentions are broadly inversely related for the other four task categories: An additional mention of a technology is associated (again, according to the median of the 40 coefficient estimates) with 0.137 fewer mentions of nonroutine interactive tasks, 0.018 fewer mentions of nonroutine manual tasks, 0.011 fewer mentions of routine cognitive tasks, and 0.017 fewer mentions of routine manual tasks.¹⁰ But there are important exceptions to these interactions: Quark XPress, CAD, Microsoft Excel, and PowerPoint are the four technologies which are associated with an increasing frequency of nonroutine interactive task-related words. Three of the networking technologies — LAN, Novell, and TCP — are associated with increased mentions of routine cognitive task mentions.

To sum up, our job ads data set allows us to investigate the degree of complementarity between tasks and technologies for the adopting occupations. In our data, new technologies tend to be mentioned jointly with analytic tasks, not with nonroutine interactive, nonroutine manual, routine cognitive, or routine manual tasks. There are important exceptions, however, such as the widely adopted office software and interactive tasks.

⁹If job titles with the highest nonroutine analytic task content were more likely to adopt ICTs, controlling for job title fixed effects would diminish our main estimates, as they would be partially driven by the composition of job titles across occupations. As Appendix B shows, this does not appear to happen.

¹⁰The frequencies with which employers mention tasks — and with which our text-processing algorithm detects task-related words — differ across the five task categories. Stating our coefficients in a comparable scale, the median effect of an individual technology mention is associated with a 0.07 standard deviation increase in nonroutine analytic task mentions, and a decline in nonroutine interactive, nonroutine manual, routine cognitive, and routine manual task mentions of (respectively) 0.20, 0.06, 0.05, and 0.11 standard deviations.

Figure 5: Relationship between Task and Technology Mentions



Notes: Each panel presents the 40 coefficient estimates and corresponding 2-standard errors confidence intervals, one for each technology, of β_{hk} from Equation 1. An “x” indicates that the coefficient estimate differs significantly from zero, while an “•” indicates that the coefficient estimate does not. Horizontal, dashed lines separate technologies into the following groups: general software, office software, networking software/hardware, other hardware, and database management systems.

4 The Macroeconomic Implications of ICTs

In this section, we develop a general equilibrium model, based on the model of [Autor, Levy, and Murnane \(2003\)](#), [Michaels, Rauch, and Redding \(2016\)](#), [Burstein, Morales, and Vogel \(2015\)](#), and most directly [Atalay, Phongthientham, Sotelo, and Tannenbaum \(2017\)](#). In our framework, new technologies directly alter the task content of occupations and, through changes in the value of occupations' output, indirectly reduce the demand for workers who were originally producing tasks now substituted by the new technologies. We use our model to study how new technologies alter the types of tasks that workers perform, and as a result, reshape their occupational choices and the wages which they earn. We first describe the model (Section 4.1), explain how we estimate workers' skills in producing tasks (Section 4.2), delineate our procedure for computing counterfactual changes in equilibrium allocations and prices in response to changes in the price of ICT capital (Section 4.3), provide details of our calibration (Section 4.4), and finally present the results from our counterfactual exercises (Section 4.5).

4.1 An Equilibrium Model of Occupation and Technology Choice

Workers belong to one of many groups $g = 1, \dots, G$, and sort across occupations $j = 1, \dots, J$. There are $k = 1, \dots, K$ ICT technologies which workers can use to perform their occupations. Workers' observable characteristics, captured by their group g , shape their ability to perform tasks. In addition, workers have an unobservable comparative advantage across occupation-ICT pairs. Workers supply one unit of labor inelastically to their jobs.¹¹

Preferences The representative consumer has constant elasticity of substitution preferences across outputs of each of the J occupations, given by the following utility function: $U = \left(\sum_j a_j^{1/\rho} Y_j^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}}$. In this function, Y_j equals the sum of the production of individual workers who work in occupation j , ρ equals the elasticity of substitution, while a_j controls the importance of each occupation in the economy.

Production The focus of our analysis is the technology for producing output in each occupation. We model an occupation as a combination of labor and capital. Labor is used to produce tasks $h = 1, \dots, H$. We model occupations as a bundle of tasks that workers

¹¹Our model does not capture the decision to leave the labor market. An extension to examine the employment margin — but one we do not pursue here — would be to include household production as an additional occupation.

need to perform. Occupations are different in the intensity with which they require tasks, as well as their complementarity with each ICT.

After choosing an occupation, each worker allocates her labor optimally across these H tasks. Moreover, workers can adopt an ICT technology $k = 1, \dots, K$ or not adopt a technology at all, $k = 0$, according to the returns of doing so. We adopt, in particular, the following formulation for occupation output of a worker from group g , if working in occupation j and using κ units of technology k :

$$\tilde{V}_{gjk}(\epsilon) = \epsilon^{\bar{\alpha}_k} \prod_{h=1}^H \left[\frac{q_{hgjk}(\epsilon)}{\alpha_{hjk}} \right]^{\alpha_{hjk}} \times \left(\frac{\kappa_{gjk}}{1 - \bar{\alpha}_k} \right)^{1 - \bar{\alpha}_k},$$

where ϵ is the worker's idiosyncratic efficiency term, which varies across occupations and ICTs; q_{hgjk} equals the units of task h produced by the worker; and κ_{gjk} equals the units of ICT k used in production. We impose that $\bar{\alpha}_k \equiv \sum_h \alpha_{hjk}$ equals 1 if $k = 0$ (where no technology is adopted), and $\bar{\alpha}_k < 1$ for technologies $k \in 1, \dots, K$. This formulation allows for flexible cost shares α_{hjk} , to reflect that at the occupation level some tasks are complementary with ICT k , while others are substitutable. We assume that ϵ is drawn i.i.d. from a Fréchet distribution, such that $\Pr[\epsilon < x] = \exp(-x^{-\theta})$.

A worker decides how to allocate her unit endowment of time to perform the H tasks that the occupation requires. The worker's skill to perform each task is determined by the group g to which she belongs, according to

$$q_{hgjk} = S_{hg} l_{hgjk},$$

where l_{hgjk} is the time allocated to task h by the worker.

ICT $k = 1, \dots, K$ is produced with a constant returns to scale technology that employs only the final good as input, with productivity $1/c_k$.

Equilibrium Payments per efficiency unit of labor for group g workers in occupation j using ICT k is

$$w_{gjk} = p_j^{\frac{1}{\bar{\alpha}_k}} (c_k)^{-\frac{1 - \bar{\alpha}_k}{\bar{\alpha}_k}} \prod_{h=1}^H S_{gh}^{\frac{\alpha_{hjk}}{\bar{\alpha}_k}}, \quad (2)$$

where p_k is the price of ICT k .¹² These payments reflect that workers allocate their time to each task h according to their comparative advantage, that ICTs are used as to maximize profits in an occupation, and that workers appropriate all of the residual value of their job,

¹²Appendix C contains the proofs to all the analytical results we obtain from the model.

net of payments to ICTs.¹³ The fraction of workers in group g that sorts into occupation j and technology k is then

$$\lambda_{gjk} = \frac{w_{gjk}^\theta}{\sum_{k'} \sum_{j'} w_{gj'k'}^\theta}. \quad (3)$$

Note that our distributional assumptions imply that the average total payment to workers in group g , which is the same as the average total payments to workers in that group who select into occupation j using ICT k , is equal to

$$\bar{W}_g = \Gamma(1 - 1/\theta) \cdot \left(\sum_j \sum_k w_{gjk}^\theta \right)^{1/\theta}, \quad (4)$$

where $\Gamma(\cdot)$ is the Gamma function.

We let the final good be the numeraire, so we set $P = 1$. Given c_k , the price of ICTs, an equilibrium is given by prices of occupational output $\{p_j\}$ and ICT uses $\{\kappa_{gjk}\}$ such that: (i) occupational-output markets clear,

$$\underbrace{a_j \left(\frac{p_j}{P} \right)^{1-\rho} E}_{\text{total spending on } j\text{-output}} = \underbrace{\sum_{g=1}^G \sum_{k=1}^K \bar{W}_g \lambda_{gjk} L_g}_{\text{wage bill in } j} + \underbrace{\sum_{g=1}^G \sum_k c_k \kappa_{gjk} \lambda_{gjk} L_g}_{\text{payments to all ICTs in occ. } j} \quad \forall j, \quad (5)$$

and (ii) ICT markets clear,¹⁴

$$c_k \kappa_{gjk} \lambda_{gjk} L_g = \underbrace{(1 - \bar{\alpha}_k)}_{\text{fraction of factor payments going to } k} \times \underbrace{\frac{\bar{W}_g \lambda_{gjk} L_g}{\bar{\alpha}_k}}_{\text{total factor payments in } g, j} \quad \forall g, j, k, \quad (6)$$

In Equation 5, total expenditure E is given by the sum of payments to all factors of production:

$$E = \sum_{g=1}^G \left(\bar{W}_g L_g + \sum_{j=1}^J \sum_{k=1}^K c_k \kappa_{gjk} \right);$$

the employment shares λ_{gjk} are consistent with sorting, as in 3; efficiency wages are consistent

¹³A way to rationalize this result, as in [Burstein, Morales, and Vogel \(2015\)](#), is to assume that each occupation's output is produced by single-worker firms that enter freely into the market, ensuring zero profits are earned.

¹⁴This market clearing condition is equivalent to a condition in terms of ICT use per worker

$$c_k \kappa_{gjk} = \frac{(1 - \bar{\alpha}_k)}{\bar{\alpha}_k} \bar{W}_g \quad \forall g, j, k.$$

with the worker's optimal time allocation and with free entry, as in 2, and our normalization relates occupational prices according to

$$1 = \left(\sum_{j=1}^J a_j \cdot p_j^{1-\rho} \right)^{\frac{1}{1-\rho}}.$$

This system of equations contains $J + G \cdot J \cdot K \cdot 3 + 1$ equations and the same number of unknowns: $\{p_j\}, \{\kappa_{gjk}, w_{gjk}, \lambda_{gjk}\}$, and E (together with a normalization).

4.2 Estimating Groups' Skills

A key input into the calibration of our model and our counterfactual exercises are measures of comparative advantage of worker groups across occupations and for using ICTs. We parameterize the skill of worker group g in producing task h , S_{gh} , as in our earlier paper:

$$\log S_{gh} = a_{h,gender} \cdot D_{gender,g} + a_{h,edu} \cdot D_{edu,g} + a_{h,exp} \cdot D_{exp,g}. \quad (7)$$

In this equation, $D_{gender,g}$, $D_{edu,g}$, and $D_{exp,g}$ are dummies for gender, education and experience, which define demographic groups, g . In our parameterization, we have two genders, five education groups, and four experience groups. As a result, there are $40 = (2 - 1) \cdot (5 - 1) \cdot (4 - 1) \cdot 5 a_h$ parameters which we need to estimate.

Our model delivers three aggregate moments that we take to data using a method of moments estimator. Let Θ denote the vector of parameters we estimate. Let \tilde{x} denote the value of variable x observed in the data and $x(\Theta)$ denote the model-implied dependence of variable x on the set of parameters. Our moments are, first, the fraction of workers of group g who work in occupation j :

$$\tilde{\lambda}_{gj} = \sum_{k=1}^K \left[\frac{w_{gjk}^\theta(\Theta)}{\sum_j w_{gjk'}^\theta(\Theta)} \right] \quad \forall g, j, \quad (8)$$

where $\lambda_{gj} \equiv \sum_{k=1}^K \lambda_{gjk}$; second, the fraction of workers in occupation j which adopt ICT k :

$$\tilde{\pi}_{jk} = \sum_g \frac{\lambda_{gjk}(\Theta) \tilde{L}_{gj}}{\sum_{g'} \tilde{L}_{g'j}} \quad \forall j, k, \quad (9)$$

and, third, the average earnings per group:

$$\tilde{W}_g = \Gamma(1 - 1/\theta) \cdot \left(\sum_j \sum_k w_{gjk}^\theta(\Theta) \right)^{1/\theta} \quad \forall g. \quad (10)$$

This system contains $G \cdot J + K \cdot J + G$ moments each decade, which we use to estimate $40 + 3 \times (J + K)$ moments: 40 a parameters, and, as fixed effects, J occupational prices, and K ICT prices. We estimate the latter two for the decades of 1960, 1980, and 2000, while the a parameters are obtained by using only data from 2000. ¹⁵

To compute the fraction of group g workers who sort into occupation j (the left hand-side of Equation 8) and the average earnings of group g workers (Equation 10), we draw on the public use sample of the decennial censuses (Ruggles, Genadek, Goeken, Grover, and Sobek, 2015).¹⁶ We use our new data set to compute the share of workers who adopt various ICT technologies (the left-hand side of Equation 9): We set this adoption rate equal to the fraction of ads corresponding to SOC code j which mention ICT technology k .

These data moments allow us to estimate the patterns of comparative advantage of worker groups across tasks, which Table 2 contains. An additional outcome of our estimation are the ICT prices, c_k , that rationalize the patterns of technology adoption we observe in the data.

4.3 Computing Counterfactual Equilibria

In this section we use our estimated model to compute the effect of changes to exogenous variables, $\{c_k\}$, and $\{L_g\}$, exploiting the “exact hat algebra” approach popularized by Dekle, Eaton, and Kortum (2008) and used in a similar context to ours by Burstein, Morales, and Vogel (2015). The advantage of this approach is that it does not require us to fully parameterize the model, and instead incorporates information about the parameters contained in employment shares and technology adoption rates observed directly in the data.

Throughout, for any variable x , we use x' to refer to the counterfactual value of that variable in response to changes in either labor supply or ICT prices, and \hat{x} to refer to its relative change, x'/x . We start by rewriting all of our equations in terms of changes. We obtain the following system of equilibrium conditions which depends on the observed shares of payments to labor and ICT and on exogenous shocks, which act as forcing variables:

¹⁵We do not estimate the model on all five decades’ worth of data because it is computationally infeasible.

¹⁶We restrict our sample to full time workers — workers who were are between the age of 16 and 65, who worked at least 40 weeks in the preceding year, who work for wages, and have non-imputed gender, age, occupation, and education data.

Table 2: Estimates of Skills

	Nonroutine Analytic	Nonroutine Interactive	Nonroutine Manual	Routine Cognitive	Routine Manual
Gender					
Female	-1.249	0.416	-2.012	3.254	-9.919
Education					
<HS	-2.272	-1.089	1.792	-1.210	3.597
High School	-1.100	-0.678	1.289	-0.187	2.736
College	1.513	0.549	-0.803	-1.212	-9.616
Post Graduate	2.275	0.773	-1.162	-3.262	-15.639
Experience					
0-9 Years	-0.553	-0.705	0.273	-0.339	-1.920
10-19 Years	-0.048	-0.291	0.432	-0.174	-1.086
30+ Years	-0.044	-0.027	0.439	0.070	-1.678

Notes: The table presents the estimates of $a_{h,gender}$, $a_{h,edu}$, and $a_{h,exp}$ for the five tasks h in our main classification of tasks. The omitted demographic groups are males, workers with Some College, and workers with 20-29 years of potential experience.

(i) occupational-output markets

$$\sum_{j=1}^J \hat{p}_j^{1-\sigma} \hat{E} \Psi_j = \Xi \sum_{g=1}^G \sum_{k=1}^K \widehat{W}_g \hat{\lambda}_{gjk} \hat{L}_g \chi_{gjk} + (1 - \Xi) \sum_{g=1}^G \sum_{k=1}^K \xi_{gjk} \hat{c}_k \hat{\kappa}_{gjk} \hat{\lambda}_{gjk}, \quad (11)$$

where Ψ_j is the share of payments to occupation j in total expenditure, Ξ is the share of labor in aggregate payments, χ_{gjk} is the share of group g , occupation j using ICT k in total labor payments, and ξ_{gjk} is the share of ICT k used by group g in occupation j in total payments to ICT;

(ii) ICT market clearing

$$\hat{\kappa}_{gjk} = \widehat{W}_g / \hat{c}_k; \quad (12)$$

(iii) Changes in aggregate income

$$\hat{E} = \Xi \sum_{g=1}^G \widehat{W}_g \hat{L}_g \zeta_g + (1 - \Xi) \sum_{g=1}^G \sum_{j=1}^J \sum_{k=1}^K \xi_{gjk} \hat{c}_k \hat{\kappa}_{gjk} \hat{\lambda}_{gjk}, \quad (13)$$

where ζ_g is group g 's share of total payments to labor (i.e., $\zeta_g \equiv \sum_{j,k} \chi_{gjk}$);

(iv) changes in employment shares

$$\hat{\lambda}_{gjk} = \frac{\hat{w}_{gjk}^\theta}{\sum_{j'} \sum_{k'} \hat{w}_{gj'k'}^\theta \lambda_{gj'k'}}; \quad (14)$$

(v) Changes in wages per efficiency unit of labor

$$\hat{w}_{gjk} = (\hat{p}_j)^{\frac{1}{\bar{\alpha}_k}} (\hat{c}_k)^{-\frac{1-\bar{\alpha}_k}{\bar{\alpha}_k}}; \text{ and} \quad (15)$$

(vi) Changes in average wages per group¹⁷

$$\widehat{W}_g = \left(\sum_{jk} \lambda_{gjk} \hat{w}_{gjk}^\theta \right)^{1/\theta}. \quad (16)$$

We use this system to study the effect of the availability of ICTs on task content, wages, and inequality, driven in our model by changes in the price of individual ICT pieces, \hat{c}_k . Since we are also interested in changes in aggregate task content for task h produced in occupation j , we also compute the changes in the aggregate content of task h ,¹⁸

$$\widehat{T}_{hj} = \frac{\sum_{g,k} \frac{\alpha_{hjk}}{\bar{\alpha}_k} \cdot L_g \pi_{gjk} \hat{\pi}_{gjk}}{\sum_{g,k} \frac{\alpha_{hjk}}{\bar{\alpha}_k} \cdot L_g \pi_{gjk}}, \quad (17)$$

where $\pi_{gjk} \equiv \lambda_{gjk} / (\sum_{k'} \lambda_{gjk'})$ equals the fraction of group g , occupation j workers who adopt capital k .

4.4 Calibration

In this section, we explain how to calibrate the shares required for computing our counterfactuals. The primitive data for our calibration are: (i) the frequency of task mentions in each occupation, (ii) our task-technology regression coefficients from Section 3, (iii) average wages per group \bar{W}_g , (iv) employment shares by group and occupation, $\lambda_{gj} = \sum_k \lambda_{gjk}$, and (v) the fraction of adopters in occupation j , π_{gjk} .

First, our calibrated α_{hjk} emerge from the coefficient estimates from our Section 3 regressions. To compute α_{hj0} — the parameter which governs the importance of task h in

¹⁷Our normalization of prices becomes

$$1 = \left(\sum_{j=1}^J \Psi_j \hat{P}_j^{1-\rho} \right)^{\frac{1}{1-\rho}}.$$

¹⁸We define the aggregate content of task h as

$$T_{hj} = \sum_{g,k} (\alpha_{hjk} / \bar{\alpha}_k) L_g \pi_{gjk}.$$

occupation j when no ICT technology is being used — we take the predicted value for each occupation-task pair (plugging in the occupation fixed effect, the average of the year fixed effects, and the average ad length) when no technologies are mentioned. Since the sum of the task shares equals 1, we normalize these predicted values to sum to 1. We repeat the same procedure when k is mentioned once: To calibrate $\alpha_{hjk}/\sum_{h'}\alpha_{h'jk}$ for $k \neq 0$, we take the predicted number of task h mentions when the k technology is mentioned once.

In addition, in Appendix C.5 we explain how to construct each of the shares we list below. We start by constructing aggregates, such as the payments to ICT pieces across groups and occupations, as well as total expenditures in the economy. We then calibrate shares related to occupations, groups, and ICT use. We calibrate the share of labor in total payments, Ξ , as:

$$\Xi = \frac{\sum_g \bar{W}_g L_g}{E}.$$

To match this moment, we use information from the Bureau of Economic Analysis.¹⁹ Next we compute the share of group g , occupation j , using k in total labor payments

$$\chi_{gjk} = \frac{\bar{W}_g L_g \lambda_{gj} \pi_{gjk}}{\Xi E}.$$

Finally we compute the share of ICT k used by group g in occupation j in total payments to ICT

$$\xi_{gjk} = \frac{(1 - \bar{\alpha}_k) \bar{W}_g \pi_{gjk} L_g \lambda_{gj}}{\bar{\alpha}_k (1 - \Xi) E}.$$

Importantly, we do not observe variation across groups of adoption rates of ICT k , so we use the estimates of group skills, S , together with our estimates of task contents, α , to impute them. Appendix C.5 explains this imputation in detail.

4.5 Results

We now explore a set of counterfactual scenarios, aimed at understanding how ICTs have transformed the US labor market. More specifically, we analyze the impact of increasing

¹⁹We compute payments to labor using the data series on wage and salary disbursements in private industries. To compute, payments to ICT capital, we begin by taking the stock of ICT capital — Information Processing Equipment and Software. From these capital stocks, we compute the value of capital services by multiplying each of the stocks with the sum of the real interest rate and depreciation rate. We set the real interest rate at 0.04, the depreciation rate on Information Processing Equipment at 0.18, and the depreciation rate on Software at 0.40. The average ratio, over the 1960 to 2000 sample, of payments to ICT capital to payments to labor equals 0.053. While we use the sample average when calibrating $\bar{\alpha}$, note that the ratio of payments to ICT capital to payments to labor increases from 0.020 in 1960 to 0.088 in 2000. Our model will be able to match, at least qualitatively, the increased share of payments to ICT capital through increased ICT adoption rates (which occur, in the model, as a result of declines in the various c_k).

the price of different sets of ICTs on inequality, adoption rates, and aggregate task content, taking as a baseline the economy in the year 2000. Our choice of taking the end of the sample as the baseline reflects the fact that, in that year, the ICTs we study were already available and widely adopted, which allows us to exploit the method described in Section 4.3 and thus rely on observed adoption shares.²⁰ In all of our counterfactuals, we simulate a situation where ICTs are less available, by increasing their price (i.e., setting $\hat{c}_k > 1$).

We study three sets of shocks. First, exploiting the granularity of our ICT data, we study the impact of Unix, which was disproportionately adopted in computer programming and engineering occupations. Second, we study the impact of the Microsoft Office suite (consisting of Excel, Word, and PowerPoint), a set of office technologies widely adopted across occupations. Finally, we study the impact of all 40 of the ICTs in our data set.

A common theme in our applications is a tension of two forces that shape the effect of ICTs on inequality. On the one hand, adoption of ICTs is not homogeneous across groups of workers, which we estimate to have different skills for performing tasks. Consider, for example, a worker who has relatively high productivity in nonroutine tasks. When an ICT arrives that changes the task composition of her occupation towards more nonroutine tasks, the worker benefits because the ICT frees up her time to be allocated to more productive activities.

On the other hand, the arrival of an ICT acts as a supply shock to the occupations that adopt the technology most intensively, decreasing the price of this occupation's output, and thus lowering the wage of the workers who specialize disproportionately in this occupation. The following example with (i) two occupations (j, j'), (ii) two ICTs, and (iii) two types of workers (with $L_g = L/2$ for each group) clarifies the intuition. Workers sort according to

$$\lambda_{gjk} = \frac{\left(p_j^{1/\bar{\alpha}} c_k^{(\bar{\alpha}-1)/\bar{\alpha}} \prod_{h=1}^H S_{gh}^{(\alpha_{hjk}/\bar{\alpha})} \right)^\theta}{\bar{W}_g^\theta}.$$

In a symmetric equilibrium, where $\bar{W}_g = \bar{W}$, the relative price $p_j/p_{j'}$ reflects the relative

²⁰The opposite exercise, namely, starting the economy in the year 1960, is difficult since most technologies had not yet been introduced, and thus their impact through the lens of the model would be negligible. Studying the removal of specific technologies that were widely used in 2000 – as we do – is analogous to the exercise in the international trade literature of comparing the current, observable situation with the counterfactual scenario with autarky.

supplies of both occupations' outputs:

$$\frac{p_j}{p_{j'}} = \left[\frac{\sum_g \sum_k \left(c_k^{(\bar{\alpha}-1)/\bar{\alpha}} \prod_{h=1}^H S_{gh}^{(\alpha_{hj k}/\bar{\alpha})} \right)^\theta}{\sum_g \sum_{k'} \left(c_{k'}^{(\bar{\alpha}-1)/\bar{\alpha}} \prod_{h=1}^H S_{gh}^{(\alpha_{hj' k'}/\bar{\alpha})} \right)^\theta} \right]^{\frac{\bar{\alpha}}{\bar{\alpha}(1-\sigma)-\theta}}.$$

The exponent is negative for $\theta > \bar{\alpha}$ (which we have assumed throughout), meaning that a relative increase in output reduces relative prices unambiguously. Furthermore, this elasticity will be larger the more complementary are the occupations, attaining its maximum at $\sigma = 0$. Thus, when occupations are substitutable in consumption, there will be larger equilibrium movements of workers across occupations, which limits the effect on relative prices.

A decrease in the price of one of the ICTs, c_k , will have a disproportionate effect on the occupation-group pair which uses the ICT more intensively, as measured by $\prod_{h=1}^H S_{gh}^{(\alpha_{hj k}/\bar{\alpha})}$. In turn, the effect of this decrease in the relative price will disproportionately affect workers which specialize in that occupation, as shown in Equation 16. The availability of the new ICT increases inequality if workers in occupations whose relative prices decrease had a low wage before the shock.


4.5.1 The impact of Unix

In this counterfactual, we increase the price of Unix, c_{Unix} , as to decrease the adoption rates to essentially zero. Again, the spirit of the exercise is to get close to what the economy would look like if this ICT were not available. Although this is a large shock, the aggregate effect is somewhat muted, as it is concentrated on a small fraction of the population.

We first plot in Figure 6 the counterfactual changes in occupations' task content which would have prevailed in an environment without Unix. Across all occupations, the counterfactual nonroutine analytic task content would have been lower by 0.6 percent — and the corresponding routine manual task content would have been 1.2 percent higher — had Unix not been present. Moreover, the occupations with the largest counterfactual task changes are those which were originally adopting Unix most intensively.

Turning to the implications for the earnings distribution, the top left panel of Figure 7 shows that making Unix unavailable tends to reduce inequality, which we interpret as saying that the arrival of Unix increased inequality. The smallest losers in this counterfactual are workers with less than high school education who have a 0.5 percent decrease in their earnings, while the biggest losers (male workers with less than 10 years of experience and college education) lose about 1.6 percent of their baseline real earnings.

Figure 6: The Impact of Decreasing ICT Availability on Tasks



Notes: In each panel, the vertical axis presents the percent change in the task content of occupations (measured as the number of task mentions per ad) in a counterfactual environment without Unix. The horizontal axis in each panel plots the frequency of mentions of Unix per ad, as observed in our newspaper data.


4.5.2 The impact of the Microsoft Office Suite

In this counterfactual, we increase the price of three technologies – Excel, Word, and PowerPoint– as to decrease their adoption rates to zero. The impact of increasing their price is larger and contrary to that of Unix. To begin, these ICTs are used by many occupations and groups, and thus are more widespread than Unix (or other specialty ICTs). Also unlike in the previous Unix exercise, a counterfactual drastic increase in the price of Microsoft Office software would lead to an increase in the economy-wide nonroutine analytic task content and a reduction in nonroutine interactive task content, by 1.1 percent and 1.1 percent, respectively.


The top right panel of Figure 7 shows that reducing the availability of the Microsoft Office Suite decreases average wages. The decrease is least severe for workers with moderate levels of education: workers without a high school degree would have their earnings decline by 2.51 percent, while the earnings of high school graduates, workers with some college education, college graduates, and post-graduates would decline by 2.28, 2.37, 2.50, and 2.55 percent, respectively. We interpret these patterns as suggesting that the arrival of these set of Microsoft Office has increased wages and had no systematic impact on inequality. The reason for the disparate impact across demographic groups is that the Microsoft Office products tend to increase aggregate (manual and cognitive) routine and nonroutine manual content, benefiting low education workers relative to middle and high education workers.

Figure 7: The Impact of Decreasing ICT Availability on Earnings


(a) Unix



(b) Microsoft Office Suite



(c) All Observed ICT



Notes: Within each panel, each point gives the growth in earnings for one of the 40 g groups. The first character — “M” or “F” — describes the gender; the second set of characters — “<HS,” “HS,” “Some C,” “C,” or “>C” — the educational attainment; and the third set of characters the number of years of potential experience for the demographic group. The correlation is weighted by the number of people in each demographic group.

4.5.3 The impact of all observed ICTs

In this counterfactual, we increase the price of all ICT technologies, as to reduce average adoption rates to essentially zero. Such a large shock has important macroeconomic implications.

The most important effect of this shock is to reduce earnings across the board. The bottom panel of Figure 7 shows that earnings drop by 13 percent, on average, in a counterfactual without ICTs. However, the reduction is unevenly distributed across workers of different demographic groups. In the counterfactual equilibrium, the ratio of nonroutine analytic to routine manual aggregate task content is approximately 19 log points lower. As a result of these economy-wide task changes, counterfactual earnings declines are concentrated on workers at the top and very bottom of the initial earnings distribution. Moreover, the removal of ICTs is associated with a 3.2 percentage point decline in the earnings of College graduates, relative to High School graduates. This counterfactual reduction in the college premium is 3.8 percentage points for males, and 2.5 percentage points for females. In this way, the introduction of ICTs account for approximately 14 percent of the 23 percentage point increase in the College-to-High School premium observed from 1960 to 2000.²¹

This 14 percent figure is substantially smaller than in [Burstein, Morales, and Vogel \(2015\)](#). There, the authors report that computerization accounts for 60 percent of the increase in the skill premium that occurred from 1984 to 2003. There are two key differences between their setup and ours. First, while we study the effect of a particular set of ICTs, [Burstein, Morales, and Vogel \(2015\)](#) consider the effect of computer use as a whole. Second, while our model features comparative advantage of worker groups based on how ICTs change occupational tasks, in [Burstein, Morales, and Vogel \(2015\)](#), worker groups' comparative advantage in using computers is based on idiosyncratic shocks. But regardless of these differences, in applying the hat algebra approach, we both condition on observed shares of workers across occupations and technologies. Therefore, our different modeling approaches only yield different results because of the larger share of all computing in payments, compared to that of ICTs, as well as how we use the present model to impute the baseline observed shares of workers.

Also responsible for the relatively low figure in this section's counterfactual exercise is measurement error in ads' reporting of technologies, which will tend to attenuate the coefficient estimates presented in Section 3. Attenuated coefficient estimates in our ad-level regressions lead to calibrated α_{hjk} coefficients which vary less across k , within h, j pairs,

²¹To compute this 23 percentage point figure, we draw on our sample of full time workers in the public use sample of the decennial census. We compute the College-High School premium by regressing log earnings against education, potential experience, and gender dummies, then comparing the coefficient estimates on the College and High School category dummies.

and in turn a smaller role that lower capital prices can play in shaping occupations’ task content and workers’ earnings.

5 Conclusion

This paper contributes to the literature on the labor market effects of the computer revolution of the second half of the 20th century, a transformative period of technological change. In particular, we study the effect of ICT adoption on the task content of occupations, the sorting of workers across occupations, and earnings inequality.

Our first contribution is to measure technological adoption at the job ad level. We extract these data from the job descriptions of 6.6 million ads appearing between 1960 and 2000 in the *Boston Globe*, *New York Times*, and *Wall Street Journal*. This data set, as far as we are aware, is the most comprehensive available that includes time-varying information on tasks and technologies at the job level. We use the job title as recorded in the text, and associate it with an SOC code, to aggregate and produce a publicly available occupation-year data set.

With this new and rich source of data, we have several main findings. First, we show that technology adoption is associated with an increase in nonroutine analytic tasks. This represents an important piece of evidence that the development of computer technologies has reshaped occupational tasks (Acemoglu and Autor, 2011). Second, through the lens of the model estimation and counterfactual analysis, we are able to show that the introduction of ICTs has increased welfare but also earnings inequality, although the overall magnitude of the effects are somewhat small. Overall, our paper provides evidence that the introduction of new computer technologies has played a key role in the occupational changes of the 20th century.

References

- ACEMOGLU, D., AND D. AUTOR (2011): “Skills, Tasks and Technologies: Implications for Employment and Earnings,” *Handbook of Labor Economics*, 4, 1043–1171.
- AKERMAN, A., I. GAARDER, AND M. MOGSTAD (2015): “The Skill Complementarity of Broadband Internet,” *Quarterly Journal of Economics*, 130(4), 1781–1824.
- ATALAY, E., P. PHONGTHIENGTHAM, S. SOTELO, AND D. TANNENBAUM (2017): “The Evolving U.S. Occupational Structure,” Discussion paper.

- ATHEY, S., AND S. STERN (1998): “An Empirical Framework for Testing Theories about Complementarity in Organizational Design,” NBER Working Paper No. 6600.
- AUTOR, D. H., L. KATZ, AND M. KEARNEY (2005): “Rising Wage Inequality: The Role of Composition and Prices,” Discussion paper.
- AUTOR, D. H., L. KATZ, AND A. B. KRUEGER (1998): “Computing Inequality: Have Computers Changed the Labor Market?,” *Quarterly Journal of Economics*, 113(4), 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(4), 1279–1333.
- BARTEL, A., C. ICHNIOWSKI, AND K. SHAW (2007): “How Does Information Technology Affect Productivity? Plant-Level Comparisons of Product Innovation, Process Improvement, and Worker Skills,” *Quarterly Journal of Economics*, 122(4), 1721–1758.
- BERMAN, E., J. BOUND, AND S. MACHIN (1998): “Implications of Skill-Biased Technological Change: International Evidence,” *Quarterly Journal of Economics*, 113(4), 1245–1280.
- BRYNJOLFSSON, E., AND L. M. HITT (2003): “Computing Productivity: Firm-Level Evidence,” *Review of Economics and Statistics*, 85(4), 793–808.
- BURSTEIN, A., E. MORALES, AND J. VOGEL (2015): “Accounting for Changes in Between-Group Inequality,” Working Paper 20855, National Bureau of Economic Research.
- DEKLE, R., J. EATON, AND S. KORTUM (2008): “Global Rebalancing with Gravity: Measuring the Burden of Adjustment,” *IMF Staff Papers*, 55(3), 511–540.
- ENTORF, H., AND F. KRAMARZ (1998): “The Impact Of New Technologies On Wages: Lessons From Matching Panels On Employees And On Their Firms,” *Economics of Innovation and New Technology*, 5(2-4), 165–198.
- GOOS, M., AND A. MANNING (2007): “Lousy and Lovely Jobs: The Rising Polarization of Work in Britain,” *Review of Economics and Statistics*, 89(1), 118–133.
- GORT, M., AND S. KLEPPER (1982): “Time Paths in the Diffusion of Product Innovations,” *Economic Journal*, 92(367), 630–653.
- GRILICHES, Z. (1957): “Hybrid Corn: An Exploration in the Economics of Technological Change,” *Econometrica*, 25(4), 501–522.

- HAIKEN-DENEW, J., AND C. SCHMIDT (1999): “Money for Nothing and Your Chips for Free? The Anatomy of the PC Wage Differential,” IZA Discussion Paper No. 86.
- JUHN, C., K. M. MURPHY, AND B. PIERCE (1993): “Wage Inequality and the Rise in Returns to Skill,” *Journal of Political Economy*, 101(3), 410–442.
- KATZ, L. F., AND K. M. MURPHY (1992): “Changes in Relative Wages, 1963-1987: Supply and Demand Factors,” *Quarterly Journal of Economics*, 107(1), 35–78.
- KRUEGER, A. B. (1993): “How Computers Have Changed the Wage Structure: Evidence From Microdata, 1984-1989,” *Quarterly Journal of Economics*, 108(1), 33–61.
- KRUSELL, P., L. E. OHANIAN, J.-V. RIOS-RULL, AND G. L. VIOLANTE (2000): “Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis,” *Econometrica*, 68(5), 1029–1053.
- LEVY, F., AND R. J. MURNANE (1996): “With What Skills Are Computers a Complement?,” *American Economic Review*, 86(2), 258–262.
- MICHAELS, G., F. RAUCH, AND S. J. REDDING (2016): “Tasks and Technology in the United States 1880-2000,” Discussion paper, National Bureau of Economic Research.
- NORDHAUS, W. D. (2007): “Two Centuries of Productivity Growth in Computing,” *Journal of Economic History*, 67(1), 128–159.
- RUGGLES, S., K. GENADEK, R. GOEKEN, J. GROVER, AND M. SOBEK (2015): “Integrated Public Use Microdata Series: Version 6.0,” Minneapolis, MN: Historical Census Projects, University of Minnesota.
- SPITZ-OENER, A. (2006): “Technical Change, Job Tasks, and Rising Educational Demands: Looking Outside the Wage Structure,” *Journal of Labor Economics*, 24(2), 235–270.

A Comparison of Technology Adoption in Our Data Set to Adoption in Existing Data Sources

In this appendix, we compare our technology measures with those in existing data sets. Data from O*NET permit the measurement of technology adoption for multiple types of ICTs, but not long-run measurement of technology adoption. On the other hand, the October CPS

permits the measurement of technology adoption over a portion of our sample period, but not across technologies.

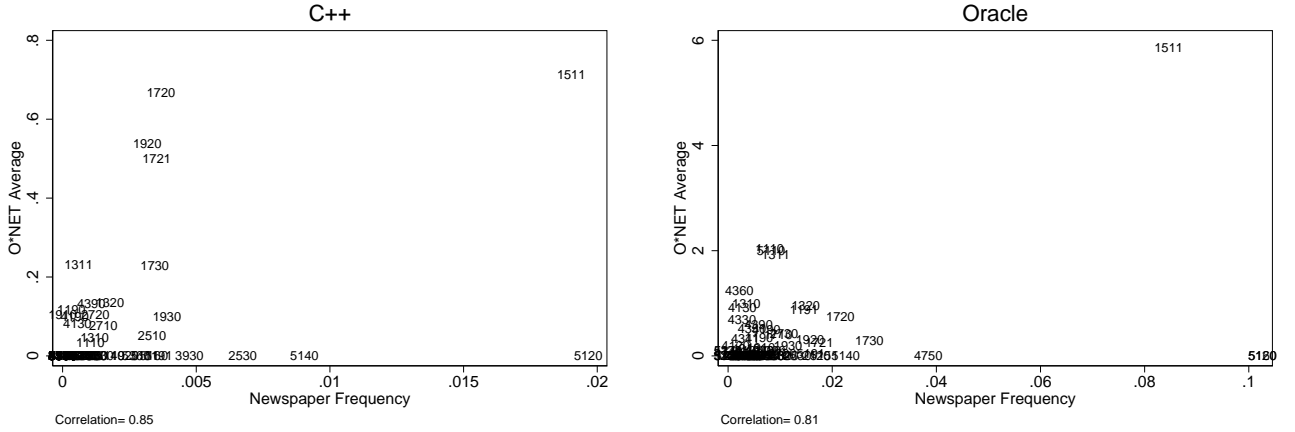
As a first comparison, from O*NET’s Tools and Technologies file, we compute the average number of mentions of eighteen ICTs (per 8-digit SOC) in each 4-digit SOC code.²² In Figure 8 we compare the average mentions per ad in our newspaper data set (as of 2000, at the end of the sample period) to the number of mentions in the O*NET Tools and Technology data set. According to the left panel of this figure, there are 0.02 mentions per ad of C++ for computer programmers (SOC=1511) in our data set, and 0.71 mentions per 8-digit SOC code according to O*NET. Weighted by the number of vacancy postings in our newspaper data, the correlation between the two data sets’ measures of C++ adoption is 0.46; the unweighted correlation is 0.26. In the right panel, we display the same comparison for a second ICT, Oracle software. Here the analogous weighted and unweighted correlations are 0.81 and 0.35. Overall, averaging across the eighteen ICTs for which we can compare occupations’ technology adoption rates, the median weighted and unweighted correlations are 0.80 and 0.31. In sum, technology adoption rates measured in our new data set broadly correlate with the rates measured in O*NET. O*NET, however, only permits measuring ICT adoption in the cross-section.

As a second check, we compare computer usage across occupations in the October CPS and the sum of the 40 technology mentions in our newspaper data. Here, we apply three editions of the October CPS — 1989, 1993, and 1997.²³ In the left panel of Figure 9 we plot the fraction of occupations’ full time workers who directly use computers at work, according to the CPS, to the sum of the technology mentions in our newspaper data. The correlation (across years and occupations), again weighting by the number of job ads in our newspaper data, is 0.40. According to this figure, nearly all of the CPS survey respondents report using computers on the job. In the newspaper data, the number of ICT mentions per ad increases from 0.66 in 1989, to 0.96 in 1993, then 1.04 in 1997. Exploiting the time variation in the October CPS and in our newspaper data, we next de-mean (within occupations) of the technology measures in the two data sets. In the right panel we plot the result of this

²²These are the eighteen (among the 40 in our original data set) which are measured in O*NET. They are C++, CAD, Foxpro, HTML, Java, LAN Lotus Notes, Microsoft Excel, Microsoft PowerPoint, Microsoft Word, Novell, Oracle, Quark, SQL, Sybase, Unix, VisualBasic, and WordPerfect.

²³For these years, the October CPS measures computer adoption rates by use (e.g., using computers for analysis, using computers for bookkeeping; using computers for calendar/scheduling; using computers for communications; using computers for databases, etc.) While computer use at work (according to the question: “Do you directly use computers at work?”) is broadly increasing from 1989 to 1993 to 1997, the average of computer adoption rates by use are decreasing from 1993 to 1997. Moreover, many of the individual questions regarding computer adoption rates by use are missing for substantially more survey respondents than for the question about overall computer use at work. For this reason, we restrict our comparison to only the overall measure of computer use.

Figure 8: Relationship between Technology Mentions in Newspaper Data and O*NET



Notes: Each panel plots the relationship of ICT adoption according to O*NET (on the vertical axis) and our newspaper data (on the horizontal axis). For each 4-digit SOC, the O*NET average is constructed by taking the number of mentions of the ICT across all 8-digit SOCs in our data and dividing by the number of 8-digit SOCs within the 4-digit SOC. We calculate the newspaper frequency by first computing the locally weighted number of ICT mentions per ad across years, within occupations, throughout the sample period, and then taking the predicted value for the year 2000. The weighted correlations for the two plotted panels are 0.85 and 0.81, respectively. Among the other ICTs, the same correlations range from 0.10, 0.38, and 0.43 (for Lotus Notes, Novell, and Microsoft PowerPoint) to 0.93, 0.95, and 0.96 (for Unix, Java and SQL).

exercise: Again, focusing on the computer programmer SOC (1511), the leftmost point in the right panel indicates that in 1989 our newspaper frequency of ICT adoption was below the average within the $t \in \{1989, 1993, 1997\}$ sample period. The positive correlation indicates that variation in computer usage is correlated across time, within occupations.


To summarize, our new data set containing measurement of the adoption of 40 ICTs within occupations over time concurs with existing data sets' measurements of ICT adoption across occupations and ICTs (according to O*NET) and across occupations and time (according to the October CPS).

B Robustness Checks Related to Section 3

In this appendix, we consider two additional exercises related to our Section 3 investigation of the relationship between ads' task and technology mentions. In Section 3, we interpret our β_{hk} coefficients as evidence for complementarity between tasks and technologies. The main concern for this interpretation is the endogeneity of technology adoption at the ad-level.

In this of exercises, we adopt specifications which include increasingly detailed occupation-level fixed effects: first, at the 6-digit SOC level (Figure 10) and second at the job title level

Figure 9: Relationship between Technology Mentions in Newspaper Data and October CPS



Notes: Each point represents a combination of a 4-digit SOC and year. In the left panel, we plot the raw ICT measures. In the right panel, we plot the measures relative to their SOC averages. The stated correlations refer to correlations weighted by the number of newspaper ads in the corresponding SOC-year combination.

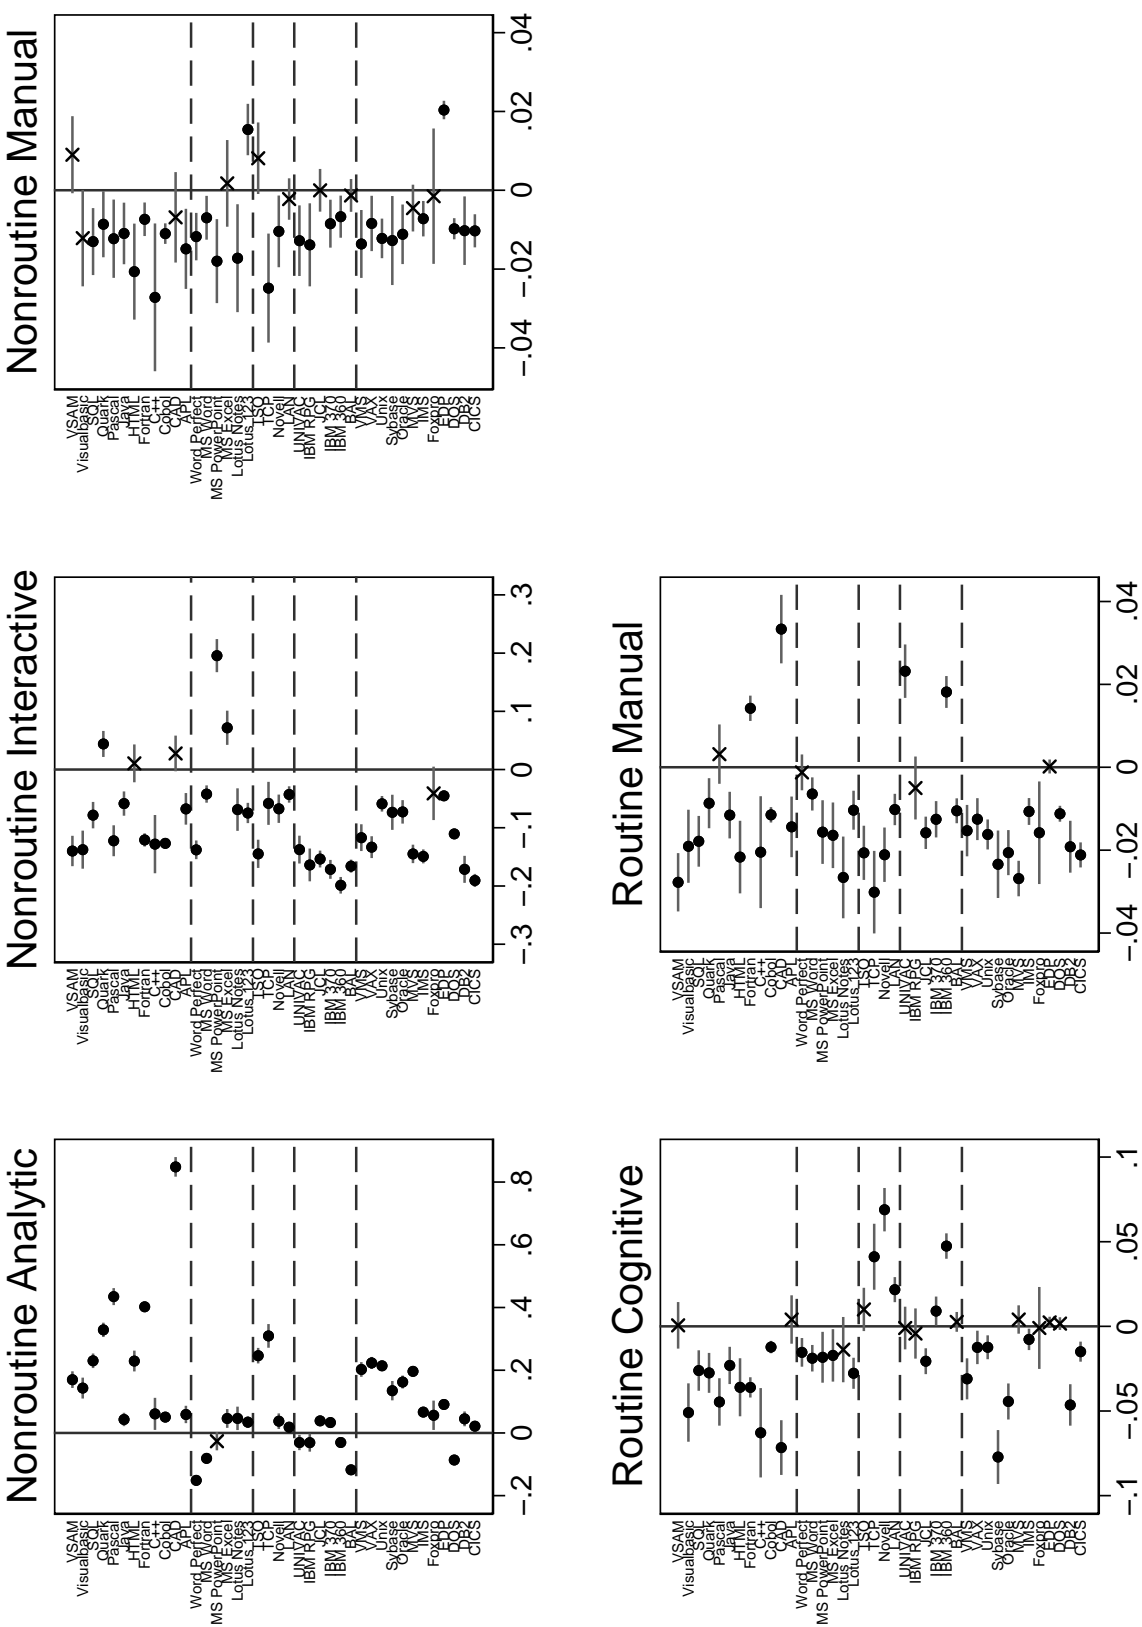
Table 3: Technologies and Tasks: Sensitivity Analysis

	(1)	(2)	(3)
Nonroutine Analytic	0.052	0.058	0.073
Nonroutine Interactive	-0.137	-0.110	-0.078
Nonroutine Manual	-0.018	-0.014	-0.012
Routine Cognitive	-0.011	-0.010	-0.007
Routine Manual	-0.017	-0.014	-0.013
Occupation Fixed Effects	4-Digit SOC	6-Digit SOC	Job Title

Notes: This table summarizes the coefficient estimates given in Figures 5, 10, and 11. Each cell gives the median coefficient estimate, across the 40 technologies.

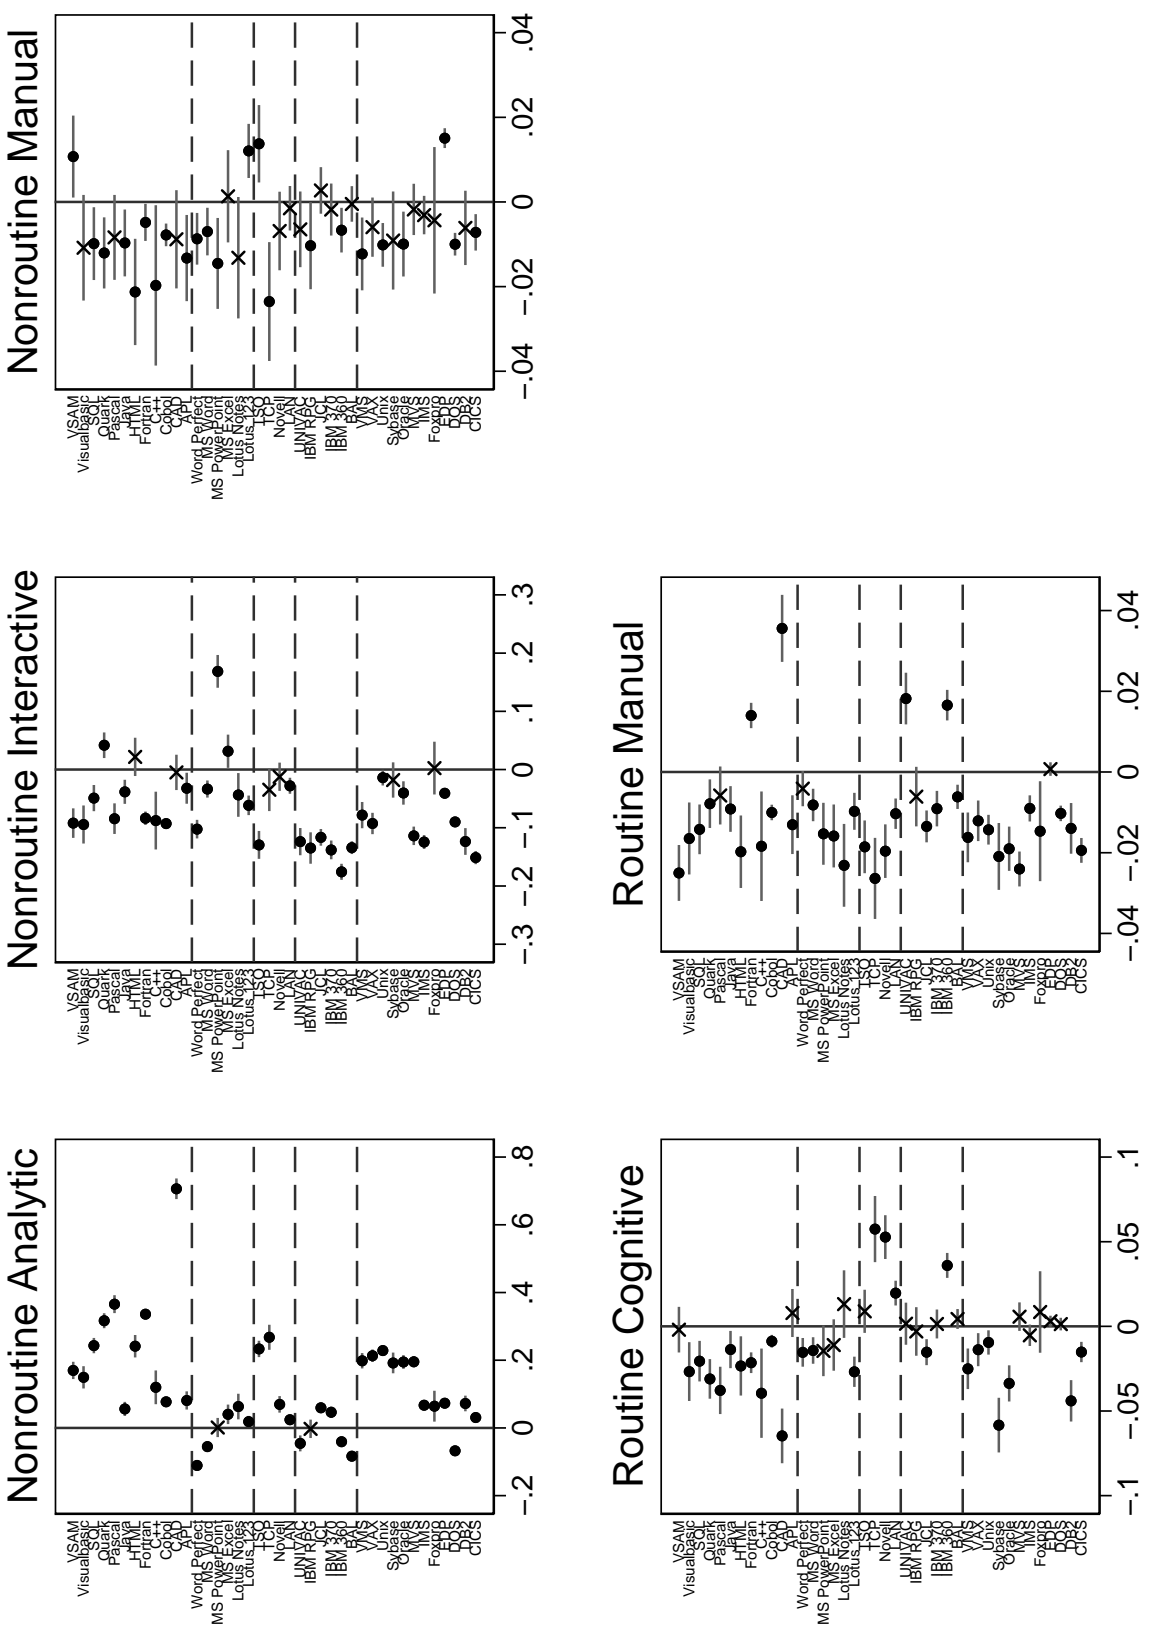
(Figure 11). The coefficient estimates given in these two figures are similar to those given in Figure 5. Whereas the median estimate (across the 40 technologies) of the relationship between technology mentions and nonroutine analytic task mentions is 0.052 when using 4-digit SOC fixed effects, the analogous coefficient is 0.058 when using 6-digit SOC fixed effects and 0.073 when using fixed effects for each job title. (See Table 3 for comparisons for the other four task measures). That the estimates are not diminished by adding job title fixed effects suggests that the estimates are not driven by endogenous adoption.

Figure 10: Relationship between Task and Technology Mentions



Notes: See the notes for Figure (5). Compared to this figure, here we apply fixed effects at the 6-digit SOC code, as opposed to the 4-digit level. Horizontal, dashed lines separate technologies into the following groups: general software, office software, networking software/hardware, other hardware, and database management systems.

Figure 11: Relationship between Task and Technology Mentions



Notes: See the notes for Figure (5). Compared to this figure, here we apply fixed effects at the job title level, as opposed to the 4-digit level. Horizontal, dashed lines separate technologies into the following groups: general software, office software, networking software/hardware, other hardware, and database management systems.

C Model Derivations

C.1 Payments to workers

We adopt the following formulation for occupation output of a worker from group g , if working in occupation j and using κ units of technology k :

$$\tilde{V}_{gjk}(\epsilon) = \epsilon^{\bar{\alpha}_k} \prod_{h=1}^H \left[\frac{q_{hgjk}(\epsilon)}{\alpha_{hjk}} \right]^{\alpha_{hjk}} \times \left(\frac{\kappa_{gjk}}{1 - \bar{\alpha}_{jk}} \right)^{1 - \bar{\alpha}_k},$$

where ϵ is an efficiency which allows for flexible cost shares, as well as productivity augmenting effects, and $\bar{\alpha}_k \equiv \sum_h \alpha_{hjk}$.

We solve the problem in stages. First, the firm takes p_j as given and chooses the amount of capital optimally. That is, κ_{gjk} solves the following first order conditions

$$p_j (1 - \bar{\alpha}_k) \tilde{V}_{gjk}(\epsilon) = c_k \kappa_{gjk}.$$

Plugging this back in the expression above, we obtain the optimized value function $V_{gjk}(\epsilon)$ that only depends on the worker's time allocations:

$$\begin{aligned} V_{gjk}(\epsilon) &= \epsilon^{\bar{\alpha}_k} \prod_{h=1}^H \left[\frac{q_{hgjk}(\epsilon)}{\alpha_{hjk}} \right]^{\alpha_{hjk}} \left(\frac{p_j V_{gjk}(\epsilon)}{c_k} \right)^{1 - \bar{\alpha}_k} \\ &\Rightarrow \\ V_{gjk}(\epsilon) &= \left[\epsilon^{\bar{\alpha}_k} \prod_{h=1}^H \left[\frac{q_{hgjk}(\epsilon)}{\alpha_{hjk}} \right]^{\alpha_{hjk}} \left(\frac{p_j}{c_k} \right)^{1 - \bar{\alpha}_{jk}} \right]^{\frac{1}{\bar{\alpha}_k}} \\ &= \epsilon \prod_{h=1}^H \left[\frac{q_{hgjk}(\epsilon)}{\alpha_{hjk}} \right]^{\frac{\alpha_{hjk}}{\bar{\alpha}_k}} \left(\frac{p_j}{c_k} \right)^{\frac{1 - \bar{\alpha}_k}{\bar{\alpha}_k}} \end{aligned}$$

Taking the function V_{gjk} as given, the worker chooses his time allocation as to maximize his payoff:

$$\max_{l_{hgjk}} \bar{\alpha}_k p_j V_{gjk}(\epsilon)$$

subject to his unit time endowment

$$\sum_h l_{hgjk} = 1.$$

This means that, in equilibrium, the worker allocates her time according to

$$l_{hjdk} = \frac{\alpha_{hjk}}{\bar{\alpha}_k}.$$

Plugging this back, we get that the worker's payment per efficiency unit of labor, conditional on working in occupation j , is

$$\begin{aligned} w_{gjk} &= \bar{\alpha}_k p_j \epsilon \prod_{h=1}^H \left[\frac{S_{gh}}{\bar{\alpha}_k} \right]^{\frac{\alpha_{hjk}}{\bar{\alpha}_k}} \left(\frac{p_j}{c_k} \right)^{\frac{1-\bar{\alpha}_k}{\bar{\alpha}_k}} \\ &= p_j^{\frac{1}{\bar{\alpha}_k}} (c_k)^{-\frac{1-\bar{\alpha}_k}{\bar{\alpha}_k}} \prod_{h=1}^H S_{gh}^{\frac{\alpha_{hjk}}{\bar{\alpha}_k}}. \end{aligned}$$

C.2 Labor supply

Using the assumption that idiosyncratic shocks are drawn from a Fréchet distribution, i.i.d across occupations and ICTs, the fraction of workers in group g that work in occupation j using ICT k is

$$\lambda_{gjk} = \frac{w_{gjk}^\theta}{\sum_{k'} \sum_{j'} w_{gj'k'}^\theta}.$$

We aggregate this labor supply at different levels, as to match what we observe in the data. The fraction of g workers who work in occupation j is given by the aggregation of such workers across all ICT uses:

$$\begin{aligned} \lambda_{gj} &= \sum_k \lambda_{gjk} \\ &= \sum_k \frac{w_{gjk}^\theta}{\sum_{k'} \sum_{j'} w_{gj'k'}^\theta} \\ &= \frac{\sum_k p_j^{\frac{1}{\bar{\alpha}_k} \theta} (c_k)^{-\theta \frac{1-\bar{\alpha}_k}{\bar{\alpha}_k}} \prod_{h=1}^H S_{gh}^{\theta \frac{\alpha_{hjk}}{\bar{\alpha}_k}}}{\sum_{k'} \sum_{j'} p_{j'}^{\theta \frac{1}{\bar{\alpha}_{k'}}} (c_{k'})^{-\theta \frac{1-\bar{\alpha}_{k'}}{\bar{\alpha}_{k'}}} \prod_{h=1}^H S_{gh}^{\theta \frac{\alpha_{hj'k'}}{\bar{\alpha}_{k'}}}}. \end{aligned}$$

C.3 ICT market clearing

The use of a worker from group g , in occupation j using ICT k is κ_{gjk} . We want to calculate aggregate ICT use over the fraction of workers who select into j , from g , which we denote Ω_{gjk} . Since all workers in g, j use the same amount of ICT, we can just multiply κ_{gjk} by the

amount of workers, $\kappa_{gjk}\lambda_{gjk}L_g$. With that, ICT markets clearing states

$$\begin{aligned}
c_k\Omega_{gjk} &\equiv c_k\kappa_{gjk}\lambda_{gjk}L_g \\
&= (1 - \bar{\alpha}_k) \frac{\bar{W}_g\lambda_{gjk}L_g}{\bar{\alpha}_k} \\
&\Leftrightarrow \\
c_k\kappa_{gjk} &= (1 - \bar{\alpha}_k) \frac{\bar{W}_g}{\bar{\alpha}_k}.
\end{aligned}$$

where the second line follows from the fact that $\bar{\alpha}_k$ is the fraction of total payments to factors that goes to workers.

C.4 Derivations of hat algebra

1. Occupational-output markets clear

$$\begin{aligned}
\sum_{s=1}^S \left(\frac{\hat{p}_j}{\hat{P}_s}\right)^{1-\sigma} \hat{P}_s^{1-\rho} \hat{E} b_{sj} \left(\frac{p_j}{P_s}\right)^{1-\sigma} a_s P_s^{1-\rho} E &= \sum_{g=1}^G \widehat{W}_g \bar{W}_g \sum_{k=1}^K \lambda_{gjk} \hat{\lambda}_{gjk} \hat{L}_g L_g + \sum_{g=1}^G \sum_k \hat{c}_k \hat{\Omega}_{jk} c_k \Omega_{gjk} \\
\sum_{s=1}^S \left(\frac{\hat{p}_j}{\hat{P}_s}\right)^{1-\sigma} \hat{P}_s^{1-\rho} \hat{E} \Psi_{sj} \Gamma_s &= \frac{1}{E} \sum_{g=1}^G \sum_{k=1}^K \widehat{W}_g \hat{\lambda}_{gjk} \hat{L}_g \bar{W}_g \lambda_{gjk} L_g \\
&\quad + \frac{1}{E} \sum_{g=1}^G \sum_k \hat{c}_k \hat{\kappa}_{jk} \hat{\lambda}_{gjk} c_k \Omega_{gjk} \\
\sum_{s=1}^S \left(\frac{\hat{p}_j}{\hat{P}_s}\right)^{1-\sigma} \hat{P}_s^{1-\rho} \hat{E} \Psi_{sj} \Gamma_s &= \Xi \sum_{g=1}^G \sum_{k=1}^K \widehat{W}_g \hat{\lambda}_{gjk} \hat{L}_g \chi_{gjk} \\
&\quad + (1 - \Xi) \sum_{g=1}^G \sum_k \xi_{gjk} \hat{c}_k \hat{\kappa}_{gjk} \hat{\lambda}_{gjk}
\end{aligned}$$

where Ψ_{sj} is the share of occupation j in sector s expenditure, Γ_s is sector s share in total spending, Ξ is the share of labor in aggregate payments, χ_{gjk} is the share of group g , occupation j using ICT k in total labor payments, and ξ_{gjk} is the share of ICT k used by group g in occupation j in total payments to ICT. The first line uses the definition $\hat{x} \equiv x'/x$ where x' is the counterfactual value of variable x . The second line forms expenditure shares, and the third line collects shares.

2. ICT markets clear

$$c_k \kappa_{gjk} = (1 - \bar{\alpha}_k) \frac{\bar{W}_g}{\bar{\alpha}_k}$$

$$\hat{c}_k \hat{\kappa}_{gjk} = \widehat{W}_g$$

which implies

$$\hat{\kappa}_{gjk} = \hat{\kappa}_{gk} = \frac{\widehat{W}_g}{\hat{c}_k}$$

3. Income

$$E = \sum_{g=1}^G \left(\bar{W}_g L_g + \sum_{j=1}^J \sum_{k=1}^K c_k \Omega_{gjk} \right)$$

$$E \hat{E} = \sum_{g=1}^G \left(\widehat{W}_g \hat{L}_g \bar{W}_g L_g + \sum_{j=1}^J \sum_{k=1}^K \hat{c}_k \hat{\kappa}_{gjk} c_k \Omega_{gjk} \right)$$

$$\hat{E} = \Xi \sum_{g=1}^G \widehat{W}_g \hat{L}_g \zeta_g + (1 - \Xi) \sum_{g=1}^G \sum_{j=1}^J \sum_{k=1}^K \hat{c}_k \hat{\kappa}_{gjk} \hat{\lambda}_{gjk} \xi_{gjk}$$

where ζ_g is the share of group g in total payments to labor (i.e., $\zeta_g \equiv \sum_{j,k} \chi_{gjk}$). That is, changes in income reflect changes in all factor payments.

4. Employment shares

$$\hat{\lambda}_{gjk} \lambda_{gjk} = \frac{\hat{w}_{gjk}^\theta w_{gjk}^\theta}{\sum_{j'} \sum_{k'} \hat{w}_{gj'k'}^\theta \hat{w}_{gj'k'}^\theta} \Rightarrow$$

$$\hat{\lambda}_{gjk} = \frac{\hat{w}_{gjk}^\theta}{\sum_{j'} \sum_{k'} \hat{w}_{gj'k'}^\theta \lambda_{gj'k'}}$$

5. Wages per efficiency unit of labor

$$w_{gjk} = p_j^{\frac{1}{\bar{\alpha}_k}} (c_k)^{-\frac{1-\bar{\alpha}_k}{\bar{\alpha}_k}} \prod_{h=1}^H S_{gh}^{\frac{\alpha_{hjk}}{\bar{\alpha}_k}}$$

$$\hat{w}_{gjk} = (\hat{p}_j)^{\frac{1}{\hat{\alpha}_k}} (\hat{c}_k)^{-\frac{1-\hat{\alpha}_k}{\hat{\alpha}_k}}$$

6. Sectoral prices

$$P_s = \left(\sum_{j=1}^J b_{sj} p_j^{1-\sigma} \right)^{\frac{1}{1-\sigma}}$$

$$\hat{P}_s = \left(\sum_{j=1}^J \Psi_{sj} \hat{p}_j^{1-\sigma} \right)^{\frac{1}{1-\sigma}}$$

7. Normalization

$$1 = \left(\sum_{s=1}^S \Gamma_s \hat{P}_s^{1-\rho} \right)^{\frac{1}{1-\rho}}$$

8. Changes in aggregate task content

$$T_{hj} \equiv \sum_{g,k} \frac{\alpha_{jhk}}{\bar{\alpha}_k} \cdot L_g \pi_{gjk}$$

$$\hat{T}_{hj} = \frac{\sum_{g,k} \frac{\alpha_{jhk}}{\bar{\alpha}_k} \cdot L_g \pi_{gjk} \hat{\pi}_{gjk}}{\sum_{g,k} \frac{\alpha_{jhk}}{\bar{\alpha}_k} \cdot L_g \pi_{gjk}}$$

C.5 Calibration of shares according to the model

The primitive data for our calibration are: (i) average wages per group \bar{W}_g , (ii) employment shares by group and occupation, $\lambda_{gj} = \sum_k \lambda_{gjk}$, (iii) the fraction of adopters in occupation j , π_{jk} , and (iv) the estimated cost shares α_{hjk} . We observe (i) and (ii) from the decennial census for various decades; we observe (iii) in our newspaper data, measured as the number of ads for occupation j that mention ICT k , relative to the total number of ads for occupation j (both in a given year); finally, (iv) we estimate using the newspaper data, as explained in Section 3.

In this appendix, our notation allows for heterogeneity across sectors, which we index by $s = 1, \dots, S$. In this extension, sectors differ according to their weight in the representative consumer's utility function. The output of each sector is a constant elasticity of substitution composite of the production in different occupations of employees working in the sector. While the main analysis in the paper considers only a single-sector economy, in future drafts we plan on analyzing how technological change affects workers who do not themselves adopt a new technology but are exposed through sectoral links. Towards this goal, it will be necessary to analyze a multi-sector economy.

ICT use by group of worker. We start by producing figures for adoption rates that depend on the worker group. Since we do not observe these directly in the data, we rely on the model to fill in the gaps. Consider the fraction of group g , occupation j workers who adopt capital k

$$\left(\frac{\lambda_{g j k t}}{\lambda_{g j 0 t}}\right)^{1/\theta} = \left(\frac{c_{k,t}}{p_{j,t}}\right)^{1-\frac{1}{\alpha_k}} \prod_h (S_{gh})^{\frac{\alpha_{hjk}}{\alpha_k} - \alpha_{hj0}}.$$

And consider the ratio of this fraction for two different demographic groups, g and g' , which will depend exclusively on groups characteristics and task shares:

$$\begin{aligned} \frac{\left(\frac{\lambda_{g j k t}}{\lambda_{g j 0 t}}\right)^{1/\theta}}{\left(\frac{\lambda_{g' j k t}}{\lambda_{g' j 0 t}}\right)^{1/\theta}} &= \prod_h \left(\frac{S_{gh}}{S_{g'h}}\right)^{\frac{\alpha_{hjk}}{\alpha_k} - \alpha_{hj0}} \\ \frac{\left(\frac{\lambda_{g j k t}}{\lambda_{g j 0 t}}\right)}{\left(\frac{\lambda_{g' j k t}}{\lambda_{g' j 0 t}}\right)} &= \prod_h \left(\frac{S_{gh}}{S_{g'h}}\right)^{\theta \frac{\alpha_{hjk}}{\alpha_k} - \theta \alpha_{hj0}}. \end{aligned}$$

Because that $\lambda_{g j k t} = \Pr(j, k|g, t) = \Pr(j|g, t) \cdot \Pr(k|j, g, t) = \lambda_{g j t} \cdot \pi_{g j k t}$, we can take logs and re-arrange to write an expression for $\log\left(\frac{\pi_{g j k t}}{\pi_{\bullet j k t}}\right)$:

$$\begin{aligned} \log\left(\frac{\pi_{g j k t}}{\pi_{g j 0 t}}\right) - \left(\frac{\pi_{g' j k t}}{\pi_{g' j 0 t}}\right) &= \theta \sum_h \left[\frac{\alpha_{hjk}}{\alpha_k} - \alpha_{hj0} \right] [\log S_{gh} - \log S_{g'h}] \\ \log\left(\frac{\pi_{g j k t}}{\pi_{g j 0 t}}\right) - \left(\frac{\pi_{j k t}}{\pi_{j 0 t}}\right) &= \theta \sum_h \left[\frac{\alpha_{hjk}}{\alpha_k} - \alpha_{hj0} \right] \left[\log S_{gh} - \sum_{g'} \frac{L_{g'} \lambda_{g' j t}}{\sum_{g''} L_{g''} \lambda_{g'' j t}} \log S_{g'h} \right] \\ \log\left(\frac{\pi_{g j k t}}{\pi_{g j 0 t}}\right) &= \log\left(\frac{\pi_{j k t}}{\pi_{j 0 t}}\right) + \theta \sum_h \left[\frac{\alpha_{hjk}}{\alpha_k} - \alpha_{hj0} \right] \left[\log S_{gh} - \sum_{g'} \frac{L_{g'} \lambda_{g' j t}}{\sum_{g''} L_{g''} \lambda_{g'' j t}} \log S_{g'h} \right] \\ \frac{\pi_{g j k t}}{\pi_{g j 0 t}} &= \frac{\pi_{j k t}}{\pi_{j 0 t}} \cdot \exp \left[\theta \sum_h \left[\frac{\alpha_{hjk}}{\alpha_k} - \alpha_{hj0} \right] \left[\log S_{gh} - \sum_{g'} \frac{L_{g'} \lambda_{g' j t}}{\sum_{g''} L_{g''} \lambda_{g'' j t}} \log S_{g'h} \right] \right] \end{aligned}$$

The terms on the right hand side are directly observable or estimated. The $\frac{L_{g'} \lambda_{g' j t}}{\sum_{g''} L_{g''} \lambda_{g'' j t}}$ come from the decennial census, the $\frac{\alpha_{hjk}}{\alpha_k}$ from our micro regressions, and the $\log S_{gh}$ come from our model estimation. We use these expressions to impute $\pi_{g j k}$, on the basis of $\pi_{j k}$, which we actually observe.

Expenditure in ICT k . Next we build from these data total expenditure in ICT k , using the market clearing equation:

$$c_k \Omega_{gjk} = (1 - \bar{\alpha}_{jk}) \frac{\bar{W}_g \lambda_{gjk} L_g}{\bar{\alpha}_{jk}}.$$

Manipulating the right-hand side, we get

$$\begin{aligned} c_k \Omega_{gjk} &= (1 - \bar{\alpha}_k) \frac{\bar{W}_g}{\bar{\alpha}_{jk}} \times \lambda_{gjk} L_g \\ &= (1 - \bar{\alpha}_k) \frac{\bar{W}_g}{\bar{\alpha}_k} \times \underbrace{\frac{\lambda_{gjk}}{\sum_{k'} \lambda_{gjk'}}}_{=\pi_{gjk}} \left(\sum_k \lambda_{gjk} \right) L_g \\ &= (1 - \bar{\alpha}_k) \frac{\bar{W}_g}{\bar{\alpha}_k} \pi_{gjk} L_g \lambda_{gj}, \end{aligned}$$

where we remove λ_{gjk} and instead we use π_{gjk} , which we observe.

Aggregate expenditure. We now compute aggregate expenditure in the economy, in a manner consistent with our framework. Our definition states that expenditure comes from the income of worker and ICTs:

$$\begin{aligned} E &= \sum_g \left\{ \bar{W}_g L_g + \sum_j \sum_k c_k \Omega_{gjk} \right\} \\ &= \sum_g \bar{W}_g L_g + \sum_j \sum_k \sum_g c_k \Omega_{gjk} \\ &= \sum_g \bar{W}_g L_g + \sum_j \sum_k \frac{(1 - \bar{\alpha}_k)}{\bar{\alpha}_k} \sum_g \pi_{gjk} \bar{W}_g L_{gj}, \end{aligned}$$

where the last expression is observable.

The share of labor in total payments, which we denote Ξ , is:

$$\Xi = \frac{\sum_g \bar{W}_g L_g}{E},$$

which implies the value of $1 - \Xi$.

Group g 's share in labor payments. Next we need to compute χ_{gjk} , the share of group g , occupation j , using k in total labor payments

$$\begin{aligned}\chi_{gjk} &= \frac{\bar{W}_g L_g \lambda_{gjk}}{\sum_g \bar{W}_g L_g} = \frac{1}{\Xi E} \bar{W}_g L_g \lambda_{gjk} \times \frac{\sum_l \lambda_{gjl}}{\sum_{k'} \lambda_{gjk'}} \\ &= \frac{1}{\Xi E} \bar{W}_g L_g \left(\sum_l \lambda_{gjl} \right) \pi_{gjk} \\ &= \frac{1}{\Xi E} \bar{W}_g L_g \lambda_{gj} \pi_{gjk}.\end{aligned}$$

Finally we compute the share of ICT k used by group g in occupation j in total payments to ICT

$$\begin{aligned}\xi_{gjk} &= \frac{c_k \Omega_{gjk}}{(1 - \Xi) E} \\ &= \frac{(1 - \bar{\alpha}_k) \frac{\bar{W}_g \lambda_{gjk} L_g}{\bar{\alpha}_k}}{(1 - \Xi) E} \\ &= \frac{(1 - \bar{\alpha}_k) \frac{\bar{W}_g}{\bar{\alpha}_k} \pi_{gjk} L_g \lambda_{gj}}{(1 - \Xi) E}.\end{aligned}$$

Sectoral shares Now we compute shares related to the importance of each sector. The only additional information we need is the total payments to all workers who work in sector s , occupation j .

We start by computing Ψ_{sj} is the share of occupation j in sector s expenditure. Recall that the total payment to occupation j firms that employ group g workers is

$$\sum_k \bar{W}_g \lambda_{gjk} L_g + \sum_k c_k \kappa_{gjk} \lambda_{gjk} L_g.$$

The average payment per firm (since the number of workers equals the number of firms) is

$$\begin{aligned}\frac{\sum_k \bar{W}_g \lambda_{gjk} L_g}{\sum_k \lambda_{gjk} L_g} + \frac{\sum_k c_k \kappa_{gjk} \lambda_{gjk} L_g}{\sum_k \lambda_{gjk} L_g} &= \bar{W}_g + \frac{\sum_k (1 - \bar{\alpha}_k) \frac{\bar{W}_g}{\bar{\alpha}_k} \pi_{gjk} L_g \lambda_{gj}}{\sum_k \lambda_{gjk} L_g} \\ &= \bar{W}_g + \bar{W}_g \sum_k \frac{(1 - \bar{\alpha}_k)}{\bar{\alpha}_k} \pi_{gjk}.\end{aligned}$$

Since there is no selection of workers across sectors, total payments to occupation j (both workers and ICT) in sector s is given by the following expression, where we denote by Λ_{gjs}

the number of workers from group g , who work in sector s and occupation j ²⁴

$$\begin{aligned}
\psi_{sj} &= \sum_g \Lambda_{gjs} \times \text{average payment to occ } j, \text{ group } g \\
&= \sum_g \Lambda_{gjs} \left\{ \bar{W}_g + \bar{W}_g \sum_k \frac{(1 - \bar{\alpha}_k)}{\bar{\alpha}_k} \pi_{gjk} \right\} \\
&= \sum_g \Lambda_{gjs} \bar{W}_g \left\{ 1 + \sum_k \frac{(1 - \bar{\alpha}_k)}{\bar{\alpha}_k} \pi_{gjk} \right\} \\
&= \left\{ 1 + \sum_k \frac{(1 - \bar{\alpha}_k)}{\bar{\alpha}_k} \pi_{gjk} \right\} \underbrace{\sum_g \Lambda_{gjs} \bar{W}_g}_{=\text{total payments to all workers in } s, j} .
\end{aligned}$$

The share we are looking for is

$$\Psi_{sj} = \psi_{sj} / \sum_{j'} \psi_{sj'}.$$

Finally, we compute Γ_s , sector s share in total spending,

$$\Gamma_s = \frac{\sum_j \psi_{sj}}{\sum_{s',j'} \psi_{s',j'}}.$$

²⁴The model does not make a prediction for these quantities, but note that we will never need them separately for the calibration; we just need them insofar as we need data on total payments to all workers in sector s , occupation j . Note that $\sum_s \Lambda_{gjs} = L_{gj}$.