University of Nebraska - Lincoln

DigitalCommons@University of Nebraska - Lincoln

Economics Department Faculty Publications

Economics Department

2018

New Technologies and the Labor Market

Enghin Atalay

Phai Phongthientham

Sebastian Sotelo

Daniel Tannenbaum

Follow this and additional works at: <https://digitalcommons.unl.edu/econfacpub>

This Article is brought to you for free and open access by the Economics Department at DigitalCommons@University of Nebraska - Lincoln. It has been accepted for inclusion in Economics Department Faculty Publications by an authorized administrator of DigitalCommons@University of Nebraska - Lincoln.

New Technologies and the Labor Market

Enghin Atalay, Phai Phongthientham, Sebastian Sotelo, Daniel Tannenbaum*

March, 2018

Abstract

Using newspaper job ad text from 1960 to 2000, we measure job tasks and the adoption of individual information and communication technologies (ICTs). Most new technologies are associated with an increase in nonroutine analytic tasks, and a decrease in nonroutine interactive, routine cognitive, and routine manual tasks. We embed these interactions in a quantitative model of worker sorting across occupations and technology adoption. Through the lens of the model, the arrival of ICTs broadly shifts workers away from routine tasks, which increases the college premium. A notable exception is the Microsoft Office suite, which has the opposite set of effects. **JEL Codes: E24, J20, O33**

*Atalay and Phongthientham: 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 thank the participants of the Carnegie-Rochester-NYU Conference for helpful comments, Jennifer Bernard for outstanding research assistance, and Erin Robertson for excellent editorial assistance. We especially thank Brad Hershbein for his thoughtful and constructive discussion of our paper. 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 48 distinct and widely-adopted 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 48 ICTs across occupations and over time. We assemble this data set through a text analysis of 4.2 million job vacancy ads appearing 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.³

We study a wide set of technologies, ranging from office software (including Lotus 123, Word Perfect, Microsoft Word, Excel, and PowerPoint), enterprise programming languages (Electronic Data Processing and Sybase), general-purpose programming languages (COBOL, FORTRAN, and Java), to hardware (UNIVAC, IBM 360, and IBM 370), among others. With this data set, we document rich interactions between individual ICTs and the task content of individual occupations. One of the strengths of the data is that we observe 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. An important exception is office software, which, compared to other technologies, is relatively less likely to appear alongside words associated with nonroutine analytic tasks.

Since our data set includes a wide range of occupations and technologies, we can speak directly to the macroeconomic implications of changes in the availability of ICTs 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

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 periods.

²We introduce part of this data set in an earlier paper; in particular, the measurement of job tasks and the mapping between job titles and SOC (Atalay, Phongthientham, Sotelo, and Tannenbaum, 2017). In it, we use the text of job vacancy ads to explore trends in the task content of occupations over the second half of the 20th century, showing that within-occupation changes in the tasks workers perform are at least as large as the changes that happen between occupations. In the current paper, we build on this earlier data set to include information about job-specific technology adoption.

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.

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 48 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, highly educated workers have a comparative advantage in producing nonroutine analytic tasks. Because new technologies increase the demand for nonroutine analytic tasks, 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 15 log points lower for the average worker, and the college-high school skill premium would have been 4.0 log points lower.⁴ Unlike most other technologies in our data, Microsoft Office technologies are only weakly correlated with nonroutine analytic tasks, and are positively correlated with nonroutine interactive tasks. As a result, we find that the introduction of Microsoft Office software has decreased the skill premium, the gender gap, and income inequality, although the magnitude of these effects is 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 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 performing routine tasks (Autor, Levy, and Murnane, 2003; Goos and Manning, 2007; Autor, Katz, and Kearney, 2005; Acemoglu and Autor, 2011). Researchers have also studied the implications of changes

⁴Between 1960 and 2000, the college-high school skill premium increased by 23 log points.

in the demand for tasks on the male-female wage gap and the female share of employment in high-wage occupations (Black and Spitz-Oener, 2010; Cortes, Jaimovich, and Siu, 2018). Our paper contributes to this literature by studying how new technologies complement (or substitute for) the types of tasks that workers of different skill groups perform. We find that ICTs tend to substitute for routine tasks (especially routine manual tasks) which are disproportionately performed by low skill workers. ICTs also allow high skill workers to focus on the activities in which they are most productive, which in our model is the essence of the complementarity between tasks and technologies. A key contribution of this paper is that we measure both technological adoption and the task content of occupations directly, over a period of immense technological change.

Our paper relates to a second literature that directly measures the adoption of specific technologies and its effect on wages and the demand for skills. These include studies of the effect of computer adoption (Krueger, 1993; Entorf and Kramarz, 1998; Autor, Katz, and Krueger, 1998; Haisken-DeNew and Schmidt, 1999) or the introduction of broadband internet (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 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 content. Section 4 takes our micro estimates and uses a quantitative model to study the aggregate impact of ICTs, while Section 5 assesses three extensions of the model. Section 6 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 show how to transform the text of help wanted ads into time-varying measures of the

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.

task content of occupations. In this paper, we turn to previously unexamined ad content: 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, as well as to examine 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. To begin, the raw text does not distinguish between job ads and other types of advertisements. Hence, in a first step, we apply a machine learning algorithm to determine which pages of advertisements are job ads. The top panel of Figure 1 presents a portion of a page that, according to our algorithm, contains 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 second Software Engineer position. Within this page of ads, we then determine the boundaries of each individual advertisement (for instance, where the first Software Engineer ad ends and the Senior Systems Engineer ad begins) and the job title. In the second step we extract, from each advertisement, words that refer to tasks the new hire is expected to perform and technologies that will be used on the job. So that we may link our text-based data to occupation-level variables in the decennial census, including wages, education, and demographic variables, our procedure also finds the Standard Occupation Classification (SOC) code corresponding to each job title (for example, 151132 for the “Software Engineers” job title).⁶

We extract job tasks from the text using a mapping between words and task categories based on Spitz-Oener (2006). The five tasks are *nonroutine analytic*, *nonroutine interactive*, *nonroutine manual*, *routine cognitive*, and *routine analytic*.⁷ To retrieve a more complete

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 is available at http://ssc.wisc.edu/~eatalay/occupation_data. There, we also provide 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 Atalay, Phongthientham, Sotelo, and Tannenbaum (2017). For convenience, we list the taxonomy 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, serving;

Figure 1: 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 analysis and testing of Advanced Coherent Radar High-Resolution Must have strong math, physics and signal processing skills, C/C++ and LAN programming expertise, plus familiarity with workstations and analytical tools such as The following require knowledge of emulators, debuggers, and logic analyzer. 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

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 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 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 to 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 for r cu or Intel microprocessor-based rs es. help **define** software requirements for 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 top panel presents text from three vacancy postings in a page of display ads in the *New York Times*. The bottom panel presents the results from our text processing algorithm. Highlighted text, within a rectangle, represents a mention of a nonroutine analytic task. Highlighted text, within an oval, represents a mention of a nonroutine interactive task. Text within an open rectangle represents a technology mention. Within these three ads, there are zero mentions of nonroutine manual, routine cognitive, or routine manual tasks.

measure of these task groups, and because we do not want our analysis to be sensitive to trends in 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 of the algorithm is that two words that share surrounding words in the text are likely to be synonyms. For example, one of the words corresponding to the nonroutine analytic task is *researching*. The continuous bag of words method uses our corpus of job ad text to find synonyms of *researching*; these synonyms include *interpreting*, *investigating*, *reviewing*, etc. In our analysis, we take the union of our original task-related words (from footnote 7) with the synonyms identified from our continuous bag of words (CBOW) model. In addition to the five task measures, we extract 48 different pieces of technology based on word appearances in the text.

The bottom panel of Figure 1 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 ten 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,” “define,” and “engineering” 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. Overall, while our data set contains some measurement error in identifying each job ad’s title and task and technology content, there is still considerable information in the text.

Table 1 lists the technologies in our sample together with information on their timing of introduction and adoption, and on their overall usage.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 next column lists the overall frequency of mentions of each piece of technology, across the 4.2 million job ads in our data set.

4) routine cognitive: bookkeeping, calculate, calculating, correcting, corrections, measurement, measuring;
5) routine manual: control, controlling, equip, equipment, equipping, operate, operating.

⁸Some of the introduction dates are ambiguous. We assign the introduction date for CAD to 1968, the date at which UNISURF (one of the original CAD/CAM systems) was introduced. Regarding point of sales technologies, Charles Kettering invented the electric motor cash register in 1906. Computerized point of sales systems were introduced in the early 1970s.

Table 1: Technologies

Technology	Year			Frequency		Year			Frequency	
	Introduced	First	Last	(%)	Technology	Introduced	First	Last	(%)	
APL	1964 ^a	1965	1998	0.05%	MS Word	1983 ^o	1993	≥2000	0.15%	
BAL	≤1964 ^b	1968	1983	0.28%	MVS	1974 ^h	1979	1998	0.14%	
CAD	1968 ^c	1981	1985	0.03%	Novell	1983 ^q	1994	1998	0.06%	
CICS	1969 ^d	1974	1998	0.28%	Oracle	1977 ^r	1995	1999	0.08%	
CNC	Late 1950s ^e	1979	≥2000	0.01%	PASCAL	1970 ^s	1982	1991	0.05%	
COBOL	1959 ^d	1968	1998	0.81%	Point of Sale	1906 ^t /1970s ^u	1963	1998	0.03%	
C++	1985 ^f	1993	1999	0.01%	PowerBuilder	1990 ^v	1995	1997	0.01%	
DB2	1983 ^g	1989	1998	0.06%	Quark	1987 ^w	1992	1999	0.07%	
DOS	1964 ^h	1969	1999	0.68%	Sabre	1960 ^x	1982	1999	0.08%	
EDP	≤1960 ⁱ	1963	1986	0.88%	SQL	1974 ^y	1993	1999	0.07%	
FORTRAN	1957 ^d	1965	1987	0.27%	Sybase	1984 ^z	1995	1997	0.04%	
FoxPro	1989 ^j	1992	1999	0.02%	TCP	1974 ^{aa}	1994	1999	0.03%	
HTML	1993 ^k	1996	≥2000	0.03%	TSO	1974 ^h	1977	1997	0.06%	
IBM 360	1964 ^b	1965	1975	0.17%	UNIVAC	1951 ^d	1960	1984	0.06%	
IBM 370	1970 ^b	1972	1982	0.13%	Unix	1969 ^d	1992	1999	0.19%	
IBM 5520	1979 ^l	1983	1987	0.02%	VAX	1977 ^d	1982	1998	0.10%	
IBM RPG	1959 ^m	1968	1993	0.04%	Visual Basic	1991 ^{ab}	1995	1998	0.03%	
Java	1995 ^d	1996	≥2000	0.07%	VMS	1977 ^d	1985	1996	0.06%	
JCL	≤1964 ^h	1970	1998	0.16%	VSAM	Early 1970s ^h	1982	1997	0.05%	
LAN	Early 1970s ⁿ	1990	1998	0.17%	Vydec	Early 1970s ^{ac}	1977	1985	0.05%	
Lotus 123	1983 ^o	1987	1997	0.11%	WordPerfect	1980 ^o	1988	1998	0.13%	
Lotus Notes	1989 ^p	1994	1998	0.03%	Xerox 630	1982 ^{ad}	1984	1988	0.01%	
MS Excel	1985 ^o	1993	≥2000	0.04%	Xerox 800	1974 ^{ad}	1977	1985	0.01%	
MS PowerPoint	1987 ^o	1995	≥2000	0.04%	Xerox 860	1979 ^{ad}	1982	1987	0.03%	

Notes: This table lists the 48 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. We define some of the less standard acronyms here: BAL (IBM Basic Assembly Language); JCL (Job Control Language); MVS (Multiple Virtual Storage); TSO (TCP Segment Offloading); and VMS (OpenVMS). Sources: *a* : Falkoff and Iverson (1978); *b* : Pugh, Johnson, and Palmer (1991); *c* : Bezier (1974); *d* : Ceruzzi (2003); *e* : Ross (1978); *f* : Stroustrup (1996); *g* : Haderle and Saracco (2013); *h* : Auslander, Larkin, and Scherr (1981); *i* : Mann and Williams (1960); *j* : Stark and Satonin (1991); *k* : Berners-Lee and Connolly (1993); *l* : May (1981); *m* : Baer (2003); *n* : Clark, Porgan, and Reed (1978); *o* : Evans, Nichols, and Reddy (1999); *p* : Rangaswamy and Lilien (1997); *q* : Major, Minshall, and Powell (1994); *r* : Preger (2012); *s* : Wirth (1971); *t* : Kettering (1906); *u* : Brobeck, Givins, Meads, and Thomas (1976); *v* : Goodall (1992); *w* : Srinivasan, Lilien, and Rangaswamy (2004); *x* : Kirkman, Rosen, Gibson, Tesluk, and McPherson (2002); *y* : Chamberlin and Boyce (1974); *z* : Epstein (2013); *aa* : Cerf and Kahn (1974); *ab* : Bronson and Rosenthal (2005); *ac* : Haigh (2006); *ad* : Xerox (1987).

The top left panel of Figure 2 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 0.01 mentions per ad in the beginning of the sample to 0.19 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 top right panel of Figure 2 documents adoption rates for each of the 48 technologies in our sample, with eight of these highlighted. Certain technologies which were prevalent in the 1960s and 1970s — including Electronic 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 the next four panels of Figure 2, 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.⁹ In each plot, a vertical line indicates 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. Third, we see that office software is adopted widely across diverse occupations, whereas other types of software, such as FORTRAN and Unix, 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-shape documented in the diffusion of many technologies (Griliches, 1957; Gort and Klepper, 1982). Here, we do not offer a theory of the pattern of adoption of new technologies for each occupation, but we do 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

⁹A foundational assumption in our work is that the words within job titles in the body of each job ad have fixed semantic meaning. Individual words (including the words within job titles) may change their semantic meaning. For instance, in 1900, the word “wanting” usually represented “lacking” or “insufficient.” In 1990, the primary meaning of “wanting” was closer to that of “wishing;” see Table 5 of Hamilton, Leskovec, and Jurafsky (2016). Another example, one which requires careful attention: In the beginning of the sample, “server” almost always represented someone in a food service occupation. Near the end of the sample, “server” appeared in job titles both for food service occupations and for computer / systems engineering occupations. For the most part, though, modifiers within job titles help distinguish between the two cases: “server - diner” and “sql server” exemplify job titles within the two occupations.

Throughout the paper, we assume that occupation titles describe bundles of tasks that are stable enough to warrant a comparison over time — e.g., it is valid to compare computer programmers in 1980 to computer programmers in 2000. Without a stable relation between job titles and occupations, there is no hope of studying trends in employment, task intensities, and ICT use across occupations.

Figure 2: Mentions of Technologies

Notes: These plots give the smoothed frequency with which job ads mention our set of technologies. The top left panel depicts the sum frequency — the number of technology mentions per job ad — of all 48 technologies. The top right panel depicts the frequencies of each of the 48 technologies separately, eight of which are highlighted in thick, dark lines and 40 of which are depicted by thin, light gray lines. Each of the bottom four panels depicts the frequencies of technology mentions for five of the top (those with the most mentions) Standard Occupation Classification (SOC) occupations, along with the economy-wide average frequency of technology mentions. The vertical lines depict the date the technology was introduced. FORTRAN was introduced in 1957, shortly before the beginning of our sample.

across time and occupations, there are existing data sets — O*NET and the October CPS — that measure ICT usage across occupations. O*NET contains information on multiple ICTs over a relatively short horizon, 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 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 is 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 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.¹⁰

¹⁰Since our job vacancy data originate from two metropolitan areas — New York and Boston — there is a potential external validity concern that the consequences of ICT adoption for occupational change may not generalize beyond these regions. We explore the extent to which the task content of occupations in Boston and New York differs substantially from the rest of the U.S. over a more recent period (2012-2017) in Appendix D.3 of [Atalay, Phongthientham, Sotelo, and Tannenbaum \(2017\)](#) and find relatively minor differences. With the same data, we perform a similar exercise in Appendix B of this paper, comparing the task-technology relationships in Boston and New York to those in the country more generally. We find that the relationship between technologies and routine manual tasks is stronger in the New York and Boston

Figure 3 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, followed by office software, networking software and hardware, other hardware, and general purpose software. According to the left panel, the relationship between nonroutine analytic task mentions and technology mentions is increasing for database management systems, networking software and hardware, and general purpose software. Among the 48 technologies in our sample, the median effect of an additional technology-related mention is an additional 0.061 nonroutine analytic task mentions per job ad. On the other hand, technology mentions and task mentions are broadly inversely related for three of the other task categories: An additional mention of a technology is associated (again, according to the median of the 48 coefficient estimates) with 0.125 fewer mentions of nonroutine interactive tasks, 0.017 fewer mentions of routine cognitive tasks, and 0.011 fewer mentions of routine manual tasks.¹¹ But there are important exceptions to these interactions: Quark XPress, Point of Sale systems, Microsoft Excel, and PowerPoint are the four technologies associated with an increasing frequency of nonroutine interactive task-related words. For most office technologies, nonroutine analytic task mentions are negatively related to technology mentions. Finally, for the fifth task category — nonroutine manual tasks — the task-technology relationships are increasing for all four of the networking technologies (LAN, Novell NetWare, TCP, and TSO) and all six of the hardware technologies (BAL, IBM 360, IBM 370, IBM RPG, JCL, and UNIVAC), with no clear relationship for the other technologies.¹²

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

metro areas than in the rest of the U.S., while the relationship between technologies and the other four task measures is broadly similar.

¹¹The 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 0.09 standard deviation increase in nonroutine analytic task mentions, a 0.02 standard deviation increase in nonroutine manual tasks, and a decline in nonroutine interactive, routine cognitive, and routine manual task mentions of 0.18, 0.07, and 0.07 standard deviations.

¹²The relationships that we estimate between point of sale technologies and nonroutine interactive tasks and between computer numerical control production technologies and routine manual tasks are exceptionally strong. These estimated relationships represent, in part, an unfortunate consequence of the way in which our text processing algorithm identifies tasks and technologies. For these two technologies, the words that refer to tasks are to some extent the same words that refer to technologies: “sale” is one word that refers to nonroutine interactive tasks; “machining” is a word that both refers to routine manual tasks and also regularly appears next to CNC in our job ad text. However, since these two technologies represent such a small share of overall technology mentions in our newspaper text, these two spuriously estimated task-technology relationships will not alter the aggregate impact of ICTs that we discuss in the following section.

Figure 3: Relationship between Task and Technology Mentions

Notes: Each panel presents the 48 coefficient estimates and corresponding 2-standard error confidence intervals, one for each technology, of β_{hk} from Equation 1. An “•” indicates that the coefficient estimate significantly differs from zero, while an “×” indicates that the coefficient estimate does not. Horizontal, dashed lines separate technologies into the following groups: general software and other technologies, office software, networking software/hardware, other hardware, and database management systems.

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, 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 C. 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.¹³ Within this appendix, we also estimate Equation 1 using occupation-year fixed effects. This specification identifies β_{hk} from comparisons of adopting jobs to non-adopting jobs within the same occupation-year cell. Here, too, the estimates of β_{hk} are close to those presented in Figure 3.¹⁴ Finally, in Appendix C, we also demonstrate that the task-technology relationships that we document within this section are, for the most part, highly correlated across ICT-task pairs over time.

In sum, 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 complementarity between the widely adopted Microsoft Office suite and interactive tasks.

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 tasks that workers perform, and as a result, reshape their occupational choices and the wages they earn. We first describe the model (Section

¹³If 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 C shows, this does not appear to happen.

¹⁴The specification with occupation-year fixed effects lessens the danger of spuriously attributing the impact of new technologies on occupations' task content to unobserved variables with coincident timing with these new technologies. Nevertheless, we prefer the specification with occupation fixed effects and year fixed effects separately. The occupation-year fixed effects remove variation which we believe to be the primary channel through which occupational change is occurring: the declining price of technologies over time.

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 that workers can use to perform their occupations, and we reserve $k = 0$ for no ICT adoption. 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.¹⁵

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

Production The focus of our analysis is on the technology used to produce output in each occupation. We model an occupation as a combination of tasks and ICTs. Labor is used to produce a bundle of tasks $h = 1, \dots, H$ that workers need to perform. Occupation-ICT combinations are different in the intensity with which they require tasks.

Workers jointly choose their occupation and whether to adopt one of the ICTs. Conditional on their ICT-occupation choice, workers choose how to allocate their time among the H tasks. We adopt, in particular, the following formulation for occupation output of a worker from group g , if working in occupation j and using technology k :

$$\tilde{V}_{gjk}(\epsilon) = \epsilon^{\bar{\alpha}_k} \cdot \left(\prod_{h=1}^H \left[\frac{q_{hgjk}(\epsilon)}{\alpha_{hjk}} \right]^{\alpha_{hjk}} \right) \cdot \left(\frac{\kappa_{gjk}}{1 - \bar{\alpha}_k} \right)^{1 - \bar{\alpha}_k}, \quad (2)$$

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

¹⁵Our benchmark model does not capture the decision to leave the labor market. An extension in Section 5 relaxes this assumption of inelastic labor supply.

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 $\bar{\alpha} < 1$ for technologies $k = 1, \dots, K$. This formulation allows for flexible cost shares α_{hjk} , reflecting 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 $\Pr[\epsilon < x] = \exp(-x^{-\theta})$.

Workers decide how to allocate their unit endowment of time to perform the H tasks that the occupation requires. Each worker's skill, S_{gh} , to perform each task h is determined by the group g to which she belongs. The number of units of task h that the worker produces is a function of the worker's skill and the time she allocates, l_{hgjk} , to task h :

$$q_{hgjk} = S_{gh} \cdot l_{hgjk} \ .$$

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

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

$$w_{gjk} = p_j^{\frac{1}{\bar{\alpha}_k}} (c_k)^{-\frac{1-\bar{\alpha}_k}{\bar{\alpha}_k}} \prod_{h=1}^H S_{gh}^{\frac{\alpha_{hjk}}{\bar{\alpha}_k}} \ , \quad (3)$$

where c_k is the price of ICT k in terms of the final good, and p_j is the price of occupation j output.¹⁶ 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, net of payments to ICTs.¹⁷ The fraction of workers in group g that sort into occupation j and technology k is then

$$\lambda_{gjk} = \frac{w_{gjk}^\theta}{\sum_{j'=1}^J \sum_{k'=0}^K w_{gj'k'}^\theta} \ . \quad (4)$$

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

¹⁶Appendix D contains the proofs to all the analytic results we obtain from the model.

¹⁷A 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.

select into occupation j using ICT k , is equal to

$$\bar{W}_g = \Gamma(1 - 1/\theta) \cdot \left(\sum_{j=1}^J \sum_{k=0}^K w_{gjk}^\theta \right)^{1/\theta}, \quad (5)$$

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

Given the price of ICTs $\{c_k\}$, an equilibrium is given by prices of occupational output $\{p_j\}$ and ICT uses $\{\kappa_{gjk}\}$ such that: (i) occupational-output markets clear,

$$\underbrace{a_j \left(\frac{p_j}{P} \right)^{1-\sigma} E}_{\text{total spending on occupation } j \text{ output}} = \underbrace{\sum_{g=1}^G \sum_{k=0}^K \bar{W}_g \lambda_{gjk} L_g}_{\text{wage bill in } j} + \underbrace{\sum_{g=1}^G \sum_{k=1}^K c_k \kappa_{gjk} \lambda_{gjk} L_g}_{\text{payments to all ICTs in occupation } j} \quad \forall j, \quad (6)$$

and (ii) ICT markets clear,¹⁸

$$c_k \kappa_{gjk} \lambda_{gjk} L_g = \underbrace{(1 - \bar{\alpha}_k)}_{\text{fraction of factor payments going to } k} \times \underbrace{\frac{\bar{W}_g \lambda_{gjk} L_g}{\bar{\alpha}_k}}_{\text{total factor payments in } g,j} \quad \forall g, j, k, \quad (7)$$

In Equation 6, 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 Equation 4; efficiency wages are consistent with the worker's optimal time allocation and with free entry, as in Equation 3; and our price index relates to occupational prices according to

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

This system of equations contains $J + G \cdot J \cdot K \cdot 3 + 2$ equations and the same number of unknowns: $\{p_j\}$, $\{\kappa_{gjk}, w_{gjk}, \lambda_{gjk}\}$, P , and E (together with a normalization).¹⁹

¹⁸This market clearing condition is equivalent to a condition in terms of ICT use per worker

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

¹⁹To aid in mapping the model to data, going forward we set $\bar{W}_g$ for a particular group g as the numeraire. The choice of numeraire does not alter our results.

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:

$$\log S_{gh} = a_{h,gender} \cdot D_{gender,g} + a_{h,edu} \cdot D_{edu,g} + a_{h,exp} \cdot D_{exp,g}. \quad (8)$$

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 $40 = [1 + 4 + 3] \cdot 5$ parameters a_h that we need to estimate.

Our model delivers three aggregate moments that we take to the 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 :

$$\tilde{\lambda}_{gj} = \sum_{k=0}^K \left[\frac{w_{gjk}^{\theta}(\Theta)}{\sum_{j'=1}^J \sum_{k'=0}^K w_{gj'k'}^{\theta}(\Theta)} \right] \quad \forall g, j, \quad (9)$$

where $\lambda_{gj} \equiv \sum_{k=0}^K \lambda_{gjk}$; second, the fraction of workers in occupation j that adopt ICT k :

$$\tilde{\pi}_{jk} = \sum_{g=1}^G \frac{\lambda_{gjk}(\Theta) \tilde{L}_{gj}}{\sum_{g'=1}^G \tilde{L}_{g'j}} \quad \forall j, k, \quad (10)$$

and, third, the average earnings per group:

$$\tilde{W}_g = \Gamma(1 - 1/\theta) \cdot \left(\sum_{j=1}^J \sum_{k=0}^K w_{gjk}^{\theta}(\Theta) \right)^{1/\theta} \quad \forall g. \quad (11)$$

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_h parameters, and, as fixed effects, J occupational prices and K ICT prices. We estimate the a_h parameters using only data from 2000. To limit the number of parameters we need to estimate, we use the values of $\theta = \sigma = 1.78$ from [Burstein, Morales, and Vogel \(2015\)](#).²⁰

²⁰We do not estimate the model on all five decades' worth of data because it is computationally infeasible. Estimating the model using data for the year 1980 yields a smaller effect for the impact of the Microsoft Office suite on the male-female earnings differential; and it somewhat dampens the effect of overall ICT

Table 2: Estimates of Skills

	Nonroutine Analytic	Nonroutine Interactive	Nonroutine Manual	Routine Cognitive	Routine Manual
Gender					
Female	-0.520	-0.241	-1.434	2.168	-5.964
Education					
< HS	-2.351	0.015	1.534	-1.413	3.189
High School	-1.210	-0.051	1.098	-0.376	2.231
College	1.878	0.515	-2.444	-0.589	-9.018
Post-Graduate	2.466	0.698	0.588	-1.915	-20.127
Experience					
0-9 Years	-0.941	-0.232	-0.451	-0.229	-1.469
10-19 Years	-0.190	-0.102	0.076	0.005	-0.410
30+ Years	-0.135	0.135	0.080	0.300	-0.380

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 education, and workers with 20-29 years of potential experience.

To compute the fraction of group g workers who sort into occupation j (the left hand-side of Equation 9) and the average earnings of group g workers (Equation 11), we draw on the public use sample of the decennial censuses (Ruggles, Genadek, Goeken, Grover, and Sobek, 2015).²¹ We use our new data set to compute the share of workers who adopt various ICT technologies (the left-hand side of Equation 10). 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

adoption on the college premium.

²¹We restrict our sample to workers who were between the age of 16 and 65, who worked at least 40 weeks in the preceding year, who work for wages, and who have non-imputed gender, age, occupation, and education data.

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 that depends on the observed shares of payments to labor and ICT and on exogenous shocks, which act as forcing variables:

(i) occupational-output markets

$$\left(\hat{p}_j/\hat{P}\right)^{1-\sigma} \hat{E}\Psi_j = \Xi \sum_{g=1}^G \sum_{k=0}^K \widehat{W}_g \hat{\lambda}_{gjk} \hat{L}_g \chi_{gjk} + (1 - \Xi) \sum_{g=1}^G \sum_{k=1}^K \xi_{gjk} \hat{c}_k \hat{\kappa}_{gjk} \hat{\lambda}_{gjk} \hat{L}_g, \quad (12)$$

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 ICTs;

(ii) ICT market clearing

$$\hat{\kappa}_{gjk} = \frac{\widehat{W}_g}{\hat{c}_k}; \quad (13)$$

(iii) changes in aggregate income

$$\hat{E} = \Xi \sum_{g=1}^G \widehat{W}_g \hat{L}_g \zeta_g + (1 - \Xi) \sum_{g=1}^G \sum_{j=1}^J \sum_{k=1}^K \xi_{gjk} \hat{c}_k \hat{\kappa}_{gjk} \hat{\lambda}_{gjk} \hat{L}_g, \quad (14)$$

where ζ_g is group g 's share of total payments to labor (i.e., $\zeta_g \equiv \sum_{j=1}^J \sum_{k=0}^K \chi_{gjk}$);

(iv) changes in employment shares

$$\hat{\lambda}_{gjk} = \frac{\hat{w}_{gjk}^\theta}{\sum_{j'=1}^J \sum_{k'=0}^K \hat{w}_{gj'k'}^\theta \lambda_{gj'k'}}; \quad (15)$$

(v) changes in wages per efficiency unit of labor

$$\hat{w}_{gjk} = (\hat{p}_j)^{\frac{1}{\bar{\alpha}_k}} (\hat{c}_k)^{-\frac{1-\bar{\alpha}_k}{\bar{\alpha}_k}}; \text{ and} \quad (16)$$

(vi) changes in average wages per group²²

$$\widehat{W}_g = \left(\sum_{j=1}^J \sum_{k=0}^K \lambda_{gjk} \hat{w}_{gjk}^\theta \right)^{1/\theta}. \quad (17)$$

We use this system to study the effect of the availability of ICTs — driven in our model by changes in the price of individual ICT pieces, $\hat{c}_k$ — on task content, wages, and inequality. 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 ,²³

$$\hat{T}_{hj} = \frac{\sum_{g=1}^G \sum_{k=0}^K \frac{\alpha_{jhk}}{\bar{\alpha}_k} \cdot L_g \lambda_{gjk} \hat{\lambda}_{gjk} \hat{L}_g}{\sum_{g=1}^G \sum_{k=0}^K \frac{\alpha_{jhk}}{\bar{\alpha}_k} \cdot L_g \lambda_{gjk}}. \quad (18)$$

4.4 Calibration

In this section, we explain how to calibrate the shares required for computing our counterfactual exercises. 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, $\lambda_{gj} = \sum_{k=0}^K \lambda_{gjk}$, and (v) the fraction of adopters in occupation j , π_{jk} .

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 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. To calibrate

²²The change for the price index is given by

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

while the change in the prices of ICTs is given by

$$\hat{c}_k = \hat{P} \hat{c}_k.$$

²³We define the aggregate content of task h as

$$T_{hj} = \sum_{g=1}^G \sum_{k=0}^K (\alpha_{hjk} / \bar{\alpha}_k) L_g \lambda_{gjk}.$$

$\alpha_{hjk} / \sum_{h'=1}^H \alpha_{h'jk}$ for $k \neq 0$, we take the predicted number of task h mentions when the k technology is mentioned once.

In addition, in Appendix D.6 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:

$$\Xi = \frac{\sum_{g=1}^G \bar{W}_g L_g}{E}.$$

To match this moment, we use information from the Bureau of Economic Analysis.²⁴ Next, we compute the share of group g , occupation j , using ICT 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

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

Importantly, we do not observe variation across groups in adoption rates of ICT k , so we use the estimates of group skills, S , together with our estimates of task content, α , to impute π_{gjk} . Appendix D.6 explains this imputation in detail.

4.5 Results

We now explore a set of counterfactual scenarios, aimed at understanding how ICTs have transformed the U.S. labor market. More specifically, we analyze the impact of increasing the price of different sets of ICTs on inequality and aggregate task content, taking the economy in the year 2000 as a baseline. 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

²⁴We 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 as 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. This procedure yields a value of 0.57 for $\bar{\alpha}$. 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).

adopted, which allows us to exploit the method described in Section 4.3 and thus rely on observed adoption shares.²⁵ In all of our counterfactual exercises, we simulate a situation where ICTs are less available by increasing their price (i.e., setting $\hat{c}_k > 1$).²⁶

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 48 of the ICTs in our data set. We choose these counterfactual exercises to study the effects of ICTs that affect particular groups more than others, and also to compare micro and macro shocks.²⁷

A common theme in our applications is a tension between two forces that shape the effect of ICTs on inequality. On the one hand, adoption of ICTs differs across groups of workers, who we estimate to have different skills for performing tasks. Consider, for example, a worker who has relatively high productivity in nonroutine tasks. The introduction of an ICT which is complementary to nonroutine tasks benefits the worker, since it shifts the allocation of her time to tasks in which she has a comparative advantage. 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.²⁸

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.

²⁵The 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 a counterfactual autarky scenario.

²⁶Note that while in our model we allow for many margins of adjustment in general equilibrium, we keep other choices fixed. For instance, human capital accumulation decisions — which would manifest as changes in the relative size of L_g — are fixed.

²⁷As we have argued above, Unix is mostly adopted by programmers and engineers, and tends to complement analytic tasks (as do the large majority of ICTs), while adoption of the Microsoft Office suite has been widespread and tends to complement interactive skills.

²⁸Appendix D.4 shows that, when occupations are substitutable in consumption, there will be larger equilibrium movements of workers across occupations in response to shocks, which limits the effect on relative prices, and thus decreases the strength of the second force.

We first plot in Figure 4 the counterfactual changes in occupations’ task content which would have prevailed in an environment without Unix. Had Unix not been present, across all occupations the counterfactual nonroutine analytic task content would have been lower by 0.6 percent, while the corresponding routine cognitive task content would have been 0.6 percent higher. Moreover, the occupations with the largest counterfactual task changes are those which adopted Unix most intensely.

Turning to the implications for the earnings distribution, the bottom right panel of right panel of Figure 4 shows that making Unix unavailable tends to reduce inequality, which we interpret as saying that the arrival of Unix increased inequality. Workers with less than high school education are least affected; their earnings are 0.3 percent lower in a counterfactual environment without Unix. On the other hand, workers with a post-graduate degree lose about 1.2 percent of their baseline real earnings.

4.5.2 The Impact of the Microsoft Office Suite

In this counterfactual, we increase the price of three technologies — Excel, Word, and PowerPoint — so as to decrease their adoption rates to zero. The impact of increasing their price is larger and opposite to that of Unix. To begin, these ICTs are used by many occupations and groups, and thus they are more widespread than Unix (or other specialty ICTs). Furthermore, unlike in the previous Unix exercise, a counterfactual elimination of Microsoft Office software would lead to an increase in the economy-wide nonroutine analytic task content by 0.8 percent, and a decline in nonroutine interactive task content by 0.9 percent.

The bottom right panel of Figure 5 shows that reducing the availability of the Microsoft Office suite decreases average earnings and increases inequality. The earnings decrease is least severe for workers with moderate levels of education: Earnings of workers without a high school degree would decline by 2.3 percent, while the earnings of high school graduates, workers with some college education, college graduates, and post-graduates would decline by 2.0, 1.8, 1.8, and 1.8 percent, respectively. Unlike Unix, there is a noticeable difference between female and male workers. The earnings of female workers decrease by about 0.2 percentage points more in a counterfactual world without Microsoft Excel, PowerPoint, and Word (i.e., close to a 13 percent larger drop than for males). The intuition for this finding is that, according to our Section 4.2 estimation, male workers have a comparative advantage over women in producing nonroutine analytic tasks. Since the Microsoft Office technologies are substitutes with these tasks, these technologies have attenuated the gender wage gap.

Figure 4: The Impact of Unix on Occupations' Tasks and Groups' Earnings

Notes: See the notes for Figure 4.

4.5.3 The Impact of All Observed ICTs

In this counterfactual, we increase the price of all ICTs so as to reduce adoption rates to essentially zero. Such a large shock has important macroeconomic implications, the most important of which is to reduce earnings across the board. In the counterfactual equilibrium, the ratio of nonroutine analytic to nonroutine interactive aggregate task content is approximately 8 log points lower, while the ratio of nonroutine analytic to routine manual task content is also approximately 8 log points lower. The bottom right panel of Figure 6 shows that earnings drop by 15 percent on average in a counterfactual without ICTs. However, the reduction is unevenly distributed across workers of different demographic groups. The removal of ICTs is associated with a 4.0 percentage point decline in the earnings of college graduates relative to those of high school graduates. This counterfactual reduction in the college premium is 5.2 percentage points for males and 2.7 percentage points for females. In this way, the introduction of ICTs accounts for approximately 17 percent of the 23 log point increase in the college to high school premium observed from 1960 to 2000.²⁹

This 17 percent figure is substantially smaller than that 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 in [Burstein, Morales, and Vogel \(2015\)](#) worker groups' comparative advantage in using computers is based on idiosyncratic shocks, our model also contains a comparative advantage component based on how ICTs change occupational tasks. In applying the hat algebra approach, however, we both condition on observed shares of workers across occupations and technologies. Therefore, our approaches will yield differing results to the extent that we calibrate to different groups of technologies — ICTs in our exercises as opposed to computer equipment in [Burstein, Morales, and Vogel \(2015\)](#) — or that we have different methods to compute the baseline shares.

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. In turn, this leads to a smaller role that lower capital prices can play in shaping occupations' task content

²⁹To compute this 23 log 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 and then comparing the coefficient estimates on the college and high school category dummies.

Figure 6: The Impact of All 48 ICTs on Occupations' Tasks and Groups' Earnings

Notes: See the notes for Figure 4.

and workers' earnings.

5 Extensions

We now consider three extensions of our model. First, we relax the rather severe imposition that counterfactual ICT price changes are so large as to completely eliminate technology adoption in our counterfactual equilibrium, by extracting changes by decade in ICT prices from observed adoption rates. Next, we break down the total effect we have measured in our Section 4 exercises into a component that comes from technology changes and a component that comes from worker sorting. We do so by considering counterfactual scenarios in which workers are fixed in their occupations. In a final extension, we augment our model to have a non-employment margin.

5.1 Finite Price Changes

In Section 4, we assessed the impact of technologies on the labor market by examining a counterfactual equilibrium in which the 48 technologies in our data set were unavailable. This counterfactual is a useful approximation of the long-run impact of these technologies: The frequencies at which employers mention our 48 ICTs is an order of magnitude smaller at the beginning of our sample than at the end. In this section, we aim to explore the impact of ICTs at shorter horizons, with more moderate shifts in ICT prices.

In Section 4.2, we have already estimated the changes in ICT prices that best explain demographic groups' wages, occupational choices, and average ICT adoption rates across each decade. The top left panel of Figure 7 presents the shifts in ICT prices from 1970 to 2000. For the median ICT, prices declined between 1970 and 2000 by approximately 6 log points per year. Among the ICTs we have highlighted in our counterfactual exercises, the price of Unix declined by 8 log points per year, with the largest decrease occurring in the 1980s. The price of Microsoft Excel, PowerPoint, and Word decreased by 16 log points, 27 log points, and 12 log points annually during the 1990s. In sum, our data on technology usage rates indicate a relatively sharp decline in the price of ICTs.

In the remaining panels of Figure 7, we consider counterfactual equilibria we would obtain if different combinations of ICT prices were changed from their year 2000 values. In the top right panel, we consider the effect of increasing Microsoft Office prices from their 2000 levels to their 1990 levels. For these prices, the effect on groups' earnings is similar to the changes we report in Figure 5. In other words, a large portion of the impact of the Microsoft Office suite on the distribution of earnings is due to shifts which occurred in the 1990s. In the bottom

Figure 7: The Impact of All 48 ICTs on Groups' Earnings: Finite Price Changes

Notes: The top left panel presents the change in ICT prices between 1970 and 2000, estimated in Section 4.2. The top right panel computes groups' earnings in the counterfactual equilibrium, in which Microsoft Office prices are set to the values associated with the year 1990. As of 1980, the Microsoft Office suite had not yet been introduced.. All other ICT prices are fixed to their 2000 levels. The bottom left panel and bottom right panel present earnings corresponding to counterfactual equilibria associated with 1970 and 1990 Unix prices, respectively. All other ICT prices are fixed to their 2000 levels.

Figure 8: The Impact of All 48 ICTs on Tasks and Groups' Earnings and Occupations': Baseline and Fixed-Shares Counterfactual

Notes: The left panel depicts the relationship between counterfactual changes in our 40 groups' earnings, according to the benchmark equilibrium in which workers are allowed to sort across ICT-occupation groups (x-axis), versus the equilibrium in which workers are fixed to their ICT-occupation (y-axis). The right panel depicts the changes in the value of occupations' routine cognitive tasks in the two counterfactual equilibria. In both panels, we also plot the 45-degree line.

two panels, we depict the counterfactual earnings which would result from an increase in the price of Unix to their 1970 levels (left panel) or their 1990 levels (right panel). From these panels, we conclude that similar to Microsoft Office, much of the Unix's impact on the labor market occurred due to 1990s price declines.

5.2 Short-run Adjustment

We now compute an equilibrium, which we also interpret as a short-run one, in which we limit workers ability to sort across occupations and technology uses in response to shocks. (We revert to the Section 4 assumption that ICT price changes are so large as to remove technology entirely.) In particular, we fix employment shares, λ_{gjk} , at their levels in 2000. Our goal is to break down the changes in inequality and task content that we measure in Section 4 into a component coming from re-sorting of workers across occupations and technology uses, and a component coming from changes in worker productivity (associated with the unavailability of ICTs).³⁰

For the sake of brevity, we focus again on our third counterfactual, in which we remove

³⁰Appendix E spells out this new notion of equilibrium and the associated hat-algebra equations. An intermediate counterfactual in which workers retain their occupations but are free to adjust across ICTs is conceptually possible, but is difficult to implement properly: Optimal re-sorting, conditional on optimal sorting at the baseline, removes most of the tractability of our framework.

all ICTs. The left panel of Figure 8 compares the changes in group average earnings in our baseline to those we obtain with fixed labor shares. When workers cannot re-sort, wages adjust more strongly to clear markets. Two results stand out. First, the resulting changes in inequality are larger than in the baseline since high income workers, who are more hurt by the absence of ICTs, cannot redeploy their skills in different occupations. For example, the standard deviation of the changes in earnings is approximately 26 percent smaller in the baseline, relative to this fixed labor counterfactual. Second, since the economy as a whole is less able to adjust to these changes, real wages fall more on average.

Next, we compare changes in task content. The first thing to note is that, given our specification of technology (Equation 2), changes in task *quantities* can only happen when workers move at the extensive margin — i.e., when they re-sort across occupations and ICTs. The reason is that, conditional on an occupation and ICT choice, the time allocation and task output, q_{hgjk} , is independent of ICT prices. Therefore, for the purpose of this exercise, we compare changes in the *value* of task content. As an example, the right panel of Figure 8 compares changes in the total value accrued to routine cognitive tasks in the baseline to those in the fixed-labor counterfactual. Again, there are two main takeaways from this comparison. First, most of the variation in changes in task value come from workers sorting across occupations, as the dispersion in changes is quite small relative to the full-adjustment baseline. Second, there is essentially no correlation between these changes, across occupations, suggesting that these intensive margin changes do not point in the same direction as the extensive margin ones.

5.3 Non-employment

Finally, we extend our model to allow workers to vary their total labor supply in response to market conditions. The motivation for this extension is the difference in trends in groups' labor force participation. In particular, labor force participation has declined for men, especially for low skilled men, throughout our sample period.

In our extension, we associate non-employment with a group-specific nonmarket benefit; within groups, workers are heterogeneous in their ability to take advantage of this benefit. The purpose of this exercise is to assess whether ICTs, by reducing the demand for certain tasks, have contributed to movements out of employment of groups that have a comparative advantage in producing those tasks.³¹

³¹Appendix F spells out how we extend our model. It also explains how most of our hat algebra expressions remain unchanged, since information on the fraction of non-employed in each group, which is directly observable, is a sufficient statistic for the benefits of non-employment. For simplicity, we assume that heterogeneity in the idiosyncratic benefits of non-employment is governed by the same θ as before.

We start by replicating our third counterfactual, in which we make all ICTs unavailable, and examine how our inequality results change. These results are governed by two opposing forces. First, as we have shown before, low earning individuals tend to lose less from the disappearance of ICTs, because these demographic groups adopt ICTs less. Their labor market prospects are less sensitive to the presence of ICTs in the workplace. As a result, low earning individuals’ labor market participation should be relatively unaffected by the removal of ICTs. Second, in this new extension, workers select into participating in the labor market. If a group has a relatively low baseline equilibrium labor force participation rate, the same drop in market wages will induce an exceptionally large drop in this group’s participation rate.

In Appendix F, we show that the second effect tends to dominate in our model. As a result, removing all ICTs leads to large drops in employment in low income groups, perhaps unintuitively. To isolate the effects of the first force, we examine a counterfactual in which we set a constant baseline rate of non-employment across all groups. This counterfactual shows that the lowest earning workers leave employment about 2.0 percentage points less (relative to their original employment shares) than the highest income workers.

6 Conclusion

This paper contributes to the literature on the labor market effects of the information and communication technology 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 4.2 million ads appearing between 1960 and 2000 in the *Boston Globe*, *New York Times*, and *Wall Street Journal*. This new and publicly available data set is, as far as we are aware, the most comprehensive available that includes time-varying information on tasks and technologies at the occupation level.

With this rich source of data, we first show that, for the most part, technology adoption is associated with an increase in nonroutine analytic tasks and — in conjunction with high income workers’ comparative advantage in occupations rich in these types of tasks — an increase in earnings inequality. However, there are important exceptions: Office software tends to substitute for nonroutine analytic tasks, and leads to an attenuation of the male-female earnings gap. We view our characterization of these types of differences in the impact of ICTs as a first step towards understanding how new technologies may impact the labor market of the future.

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.
- AUSLANDER, M. A., D. C. LARKIN, AND A. L. SCHERR (1981): “The Evolution of the MVS Operating System,” *IBM Journal of Research and Development*, 25(5), 471 – 482.
- 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.
- BAER, G. D. (2003): “IBM 1400 Series,” in *Encyclopedia of Computer Science*, pp. 826–828. John Wiley and Sons Ltd., Chichester, UK.
- 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.
- BERNERS-LEE, T., AND D. CONNOLLY (1993): “Hypertext Markup Language: A Representation of Textual Information and Metainformation for Retrieval and Interchange,” Discussion paper.

- BEZIER, P. (1974): “Mathematical and Practical Possibilities of UNISURF,” in *Computer Aided Geometric Design*, ed. by R. E. Barnhill, and R. F. Riesenfeld. Academic Press.
- BLACK, S. E., AND A. SPITZ-OENER (2010): “Explaining Women’s Success: Technological Change and the Skill Content of Women’s Work,” *Review of Economics and Statistics*, 92(1), 187–194.
- BROBECK, W., J. GIVINS, P. MEADS, AND R. THOMAS (1976): “Point-of-sale system and apparatus,” US Patent 3,946,220.
- BRONSON, G., AND D. ROSENTHAL (2005): *Introduction to Programming using Visual Basic .NET*. Jones and Bartlett Publishers.
- 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.
- CERF, V. G., AND R. E. KAHN (1974): “A Protocol for Packet Network Intercommunication,” *IEEE Transactions on Communications*, 22(5), 637 – 648.
- CERUZZI, P. E. (2003): *A History of Modern Computing*. MIT Press.
- CHAMBERLIN, D. D., AND R. F. BOYCE (1974): “SEQUEL: A Structured English Query Language,” in *Proceedings of the 1974 ACM SIGFIDET (Now SIGMOD) Workshop on Data Description, Access and Control*, SIGFIDET ’74, pp. 249–264, New York, NY, USA. ACM.
- CLARK, D. D., K. T. PORGAN, AND D. P. REED (1978): “An Introdution to Local Area Networks,” *Proceedings of the IEEE*, 66(11), 1457–1517.
- CORTES, G. M., N. JAIMOVICH, AND H. E. SIU (2018): “The ‘End of Men’ and Rise of Women in the High-Skilled Labor Market,” Working Paper 24274, 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.

- EPSTEIN, B. (2013): “History of Sybase,” *IEEE Annals of the History of Computing*, 35(2), 31–41.
- EVANS, D. S., A. L. NICHOLS, AND B. J. REDDY (1999): “The Rise and Fall of Leaders in Personal Computer Software,” Discussion paper.
- FALKOFF, A. D., AND K. E. IVERSON (1978): “The Evolution of APL,” *SIGAPL APL Quote Quad*, 9(1), 30–44.
- GOODALL, S. D. (1992): “Online Help: A Part of Documentation,” in *Proceedings of the 10th Annual International Conference on Systems Documentation*, SIGDOC ’92, pp. 169–174, New York, NY, USA. ACM.
- 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.
- HADERLE, D. J., AND C. M. SARACCO (2013): “The History and Growth of IBM’s DB2,” *IEEE Annals of the History of Computing*, 35(2), 54–66.
- HAIGH, T. (2006): “Remembering the Office of the Future: The Origins of Word Processing and Office Automation,” *IEEE Annals of the History of Computing*, 28(4), 6–31.
- HAISKEN-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.
- HAMILTON, W. L., J. LESKOVEC, AND D. JURAFSKY (2016): “Diachronic Word Embeddings Reveal Statistical Laws of Semantic Change,” in *Association for Computational Linguistics*.
- 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.
- KETTERING, C. F. (1906): “Cash Register,” US Patent 1,005,555.

- KIRKMAN, B. L., B. ROSEN, C. B. GIBSON, P. E. TESLUK, AND S. O. MCPHERSON (2002): “Five Challenges to Virtual Team Success: Lessons from Sabre, Inc.,” *The Academy of Management Executive (1993-2005)*, 16(3), 67–79.
- 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.
- MAJOR, D., G. MINSHALL, AND K. POWELL (1994): “An Overview of the NetWare Operating System,” in *Proceedings of the USENIX Winter 1994 Technical Conference on USENIX Winter 1994 Technical Conference*, WTEC’94, pp. 27–27, Berkeley, CA, USA. USENIX Association.
- MANN, F. C., AND L. K. WILLIAMS (1960): “Observations on the Dynamics of a Change to Electronic Data-Processing Equipment,” *Administrative Science Quarterly*, 5(2), 217–256.
- MAY, F. T. (1981): “IBM Word Processing Developments,” *IBM Journal of Research and Development*, 25(5), 741–754.
- 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.
- PREGER, R. (2012): “The Oracle Story, Part 1: 1977-1986,” *IEEE Annals of the History of Computing*, 34(4), 51–57.
- PUGH, E., L. R. JOHNSON, AND J. H. PALMER (1991): *IBM’s 360 and Early 370 Systems*. MIT Press.
- RANGASWAMY, A., AND G. L. LILIEN (1997): “Software Tools for New Product Development,” *Journal of Marketing Research*, 34(1), 177–184.

- ROSS, D. T. (1978): “Origins of the APT Language for Automatically Programmed Tools,” in *History of Programming Languages I*, ed. by R. L. Wexelblat, pp. 279–338. Association for Computing Machinery,.
- 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.
- SRINIVASAN, R., G. L. LILIEN, AND A. RANGASWAMY (2004): “First in, First out? The Effects of Network Externalities on Pioneer Survival,” *Journal of Marketing*, 68(1), 41–58.
- STARK, R., AND S. SATONIN (1991): *FoxPro: The Master Reference*. McGraw-Hill Companies.
- STROUSTRUP, B. (1996): “A History of C++: 1979-1991,” in *History of Programming Languages-II*, ed. by J. Thomas J. Bergin, and J. Richard G. Gibson, pp. 699–769. Association for Computing Machinery.
- WIRTH, N. (1971): “The Programming Language Pascal,” *Acta Informatica*, 1(1), 35–63.
- XEROX (1987): “Xerox 1987 Fact Book,” Discussion paper.

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 the data do not allow us to measure long-run patterns 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 (Version 22.1) file, we compute the average number of mentions of eighteen ICTs (per 8-digit SOC) in each 4-digit SOC code.³² In Figure 9, 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 0.84; the unweighted correlation is 0.40. In the right panel, we display the same comparison for a second ICT, Microsoft Excel. Here, the analogous weighted and unweighted correlations are 0.50 and 0.25. Overall, averaging across the 21 ICTs for which we can compare occupations’ technology adoption rates, the median weighted and unweighted correlations are 0.66 and 0.38. 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 over a short horizon.

As a second check, we compare computer usage across occupations in the October CPS and the sum of the 48 technology mentions in our newspaper data. Here, we apply three editions of the October CPS — 1989, 1993, and 1997.³³ In the left panel of Figure 10, we plot the fraction of occupations’ full time workers who directly use computers at work, according

³²These are the 21 (among the 48 in our original data set) which are measured in O*NET. They are C++, CAD, CNC, Foxpro, HTML, Java, LAN, Lotus Notes, Microsoft Excel, Microsoft PowerPoint, Microsoft Word, Novell NetWare, Oracle, Point of Sale, PowerBuilder, Quark, SQL, Sybase, Unix, VisualBasic, and WordPerfect.

³³For 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 9: 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 y-axis) and our newspaper data (on the x-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 0.84 and 0.50, respectively. Among the other ICTs, the same correlations range from -0.02, 0.10, and 0.38 (for Point of Sale technologies, Lotus Notes, and Novell) to 0.93, 0.96, and 0.96 (for Unix, Java and SQL).

to the CPS, to the number of technology mentions per ad in our newspaper text. 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, the fraction of full-time CPS respondents who report using computers on the job increases from 37 percent in 1989 to 50 percent in 1997. Similarly, in the newspaper data, the number of ICT mentions per ad increases from 0.09 in 1989 to 0.15 in 1993, and then to 0.23 in 1997. Exploiting the time variation in the October CPS and in our newspaper data, we next de-mean (within occupations) the technology measures in the two data sets. In the right panel, we plot the result of this 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 48 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).

Figure 10: 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.

B Representativeness of Boston and New York Job Ads

A key limitation of our newspaper data is that they draw on text from New York City and Boston metro area newspapers. We assess the potential scope of this limitation by comparing online vacancy postings from the New York City and Boston metro areas to vacancy postings from the rest of the United States. The underlying assumption behind this exercise is that the non-representativeness of these two metro areas in the early 2010s is informative of the non-representativeness of our sample of newspaper text during the earlier 1960 to 2000 period. To preview the results from this section, we find some mixed evidence for the representativeness of New York City and Boston: The relationship between tasks and technologies is similar for New York City and Boston compared to the rest of the U.S. for nonroutine analytic, nonroutine interactive, and nonroutine manual tasks; significantly stronger in New York City and Boston for routine manual tasks; and somewhat weaker in New York City and Boston for nonroutine analytic tasks.

Within this section, we draw on a 5 percent sample of the ads which were collected by Economic Modeling Specialists International (EMSI) between January 2012 and March 2017. We restrict attention to the 5.4 million ads (out of the 7.6 million from the 5 percent sample) for which we could map the posting's job title to an SOC code. For each of these ads, we count the number of task-related words, as well as the words related to technologies from

our 1960 to 2000 sample.³⁴

We examine whether the relationships between tasks and technologies are substantially different for the Boston and New York City metro areas. We estimate regressions described by Equation 19 below:

$$\begin{aligned} \text{task}_{ajt}^h = & \beta_{1h} \cdot 1_{a \in \{\text{Boston, New York}\}} + \text{technology}_{ajt} \cdot (\beta_{2h} + \beta_{3h} \cdot 1_{a \in \{\text{Boston, New York}\}}) \\ & + \iota_{jh} + \iota_{th} + \iota_{sh} + \epsilon_{ahjt}. \end{aligned} \quad (19)$$

Here, h refers to one of five task categories; task_{ajt}^h gives the number of mentions of task h (relative to the number of words in the ad) in a , published in year t , for an occupation j ; technology_{ajt} equals the mentions of one of the 14 technologies from footnote 34 (again relative to the number of words in the ad); and ι_{jh} , ι_{th} , and ι_{sh} respectively refer to occupation (4-digit SOC) fixed effects, year fixed effects, and fixed effects for the job message board from which EMSI procured the data. The coefficient of interest is β_{3h} , characterizing the difference in the slope of the task-technology relationship within the Boston and New York metro areas, relative to the rest of the U.S.

Table 3 gives our regression coefficients. Consistent with a similar set of regressions we perform in [Atalay, Phongthientham, Sotelo, and Tannenbaum \(2017\)](#) — in which we are interested in differential task content in Boston and New York City job ads relative to those in the rest of the U.S. — our β_{1h} estimates suggest that Boston and New York City job ads contain a higher frequency of nonroutine analytic, nonroutine interactive, and routine manual task words, and fewer nonroutine manual and routine cognitive words. New to this paper, the estimates of β_{3h} are statistically insignificant from zero for two of the five task measures, slightly negative for nonroutine analytic tasks, slightly positive for nonroutine manual tasks, and substantially negative for routine manual tasks. These coefficient estimates suggest that our benchmark estimates (estimated using newspaper data from New York City and Boston) may be overstating the strength of the negative relationship between technologies and routine manual tasks. In turn, such an over-estimation may be causing our Section 4.5 counterfactual exercises to overstate the increase in inequality due to the introduction of ICTs.

³⁴We search for mentions of 14 technologies: CAD, C++, HTML, JAVA, LAN, Microsoft Excel, Microsoft Power Point, Microsoft Word, SQL, Sybase, TCP, Unix, Visual Basic, and VSAM. The remaining 34 technologies in our benchmark set of calculations are essentially never mentioned in the 2010s in our online job ads.

Table 3: Estimates from Equation 19

Coefficient	Nonroutine Analytic	Nonroutine Interactive	Nonroutine Manual	Routine Cognitive	Routine Manual
β_{1h}	0.159 (0.011)	0.202 (0.011)	-0.083 (0.004)	-0.016 (0.004)	0.072 (0.003)
β_{2h}	0.186 (0.001)	-0.029 (0.001)	-0.035 (0.001)	0.001 (0.000)	-0.005 (0.000)
β_{3h}	-0.009 (0.004)	-0.001 (0.004)	0.004 (0.002)	0.000 (0.001)	-0.011 (0.001)

Notes: Each column contains coefficient estimates and standard errors, estimated from Equation 19, for a given h .

C Robustness Checks Related to Section 3

In Section 3, we interpret our β_{hk} coefficients as evidence of complementarity between tasks and technologies. The main concern for this interpretation is the endogeneity of technology adoption at the ad level. In addition, our Section 3 regressions impose time invariance in the relationship between task and technology mentions. But it is conceivable that, as technologies mature, the complementarity or substitutability of technologies and worker-performed tasks may evolve. In this section, we explore these two issues.

In this appendix, we consider three additional exercises related to our Section 3 investigation of the relationship between ads’ task and technology mentions. We start by reassessing these relationships, first by controlling for increasingly detailed occupation fixed effects and second by controlling for year-by-occupation fixed effects. Finally, we assess whether the relationships between tasks and technologies vary over time.

In a first exercise, we adopt specifications which include occupation-level fixed effects more detailed than those in our benchmark regressions: first, at the 6-digit SOC level (Figure 11), second at the job title level (Figure 12), and third at the 4-digit SOC by year level (Figure 13). The coefficient estimates given in these three figures are similar to those given in Figure 3. Whereas the median estimate (across the 48 technologies) of the relationship between technology mentions and nonroutine analytic task mentions is 0.061 when using 4-digit SOC fixed effects, the analogous coefficient is 0.061 when using 6-digit SOC fixed effects, 0.072 when using fixed effects for each job title, and 0.075 when using fixed 4-digit SOC by year fixed effects. (See Table 4 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: If, for example, job titles with the highest nonroutine analytic task content were more likely to adopt ICTs, then controlling for job title fixed effects would diminish our main estimates, as they would be partially driven by the composition

Figure 11: Relationship between Task and Technology Mentions

Notes: See the notes for Figure 3. Compared to this figure, here we apply fixed effects at the 6-digit SOC code level, as opposed to the 4-digit level. Horizontal, dashed lines separate technologies into the following groups: general software, office software and other technologies, networking software/hardware, other hardware, and database management systems.

Figure 12: Relationship between Task and Technology Mentions

Notes: See the notes for Figure 3. Compared to this figure, here we apply fixed effects at the job title level, as opposed to the 4-digit SOC code level. Horizontal, dashed lines separate technologies into the following groups: general software and other technologies, office software, networking software/hardware, and database management systems.

Figure 13: Relationship between Task and Technology Mentions

Notes: See the notes for Figure 3. Compared to this figure, here we apply fixed effects at the 4-digit SOC x year level, as opposed to the 4-digit level and year level separately. Horizontal, dashed lines separate technologies into the following groups: general software and other technologies, office software, networking software/hardware, other hardware, and database management systems.

Table 4: Technologies and Tasks: Sensitivity Analysis

	(1)	(2)	(3)	(4)	(5)	(6)
Nonroutine Analytic	0.061	0.061	0.072	0.075	0.061	0.071
Nonroutine Interactive	-0.125	-0.104	-0.083	-0.127	-0.119	-0.222
Nonroutine Manual	0.004	0.002	0.004	0.006	0.005	-0.003
Routine Cognitive	-0.017	-0.016	-0.011	-0.016	-0.018	-0.011
Routine Manual	-0.011	-0.013	-0.010	-0.011	-0.014	-0.001
Time Period	1960-2000	1960-2000	1960-2000	1960-2000	Early	Late
Fixed Effects	4-Digit SOC	6-Digit SOC	Job Title	4-Digit SOC × Year	4-Digit SOC	4-Digit SOC

Notes: This table summarizes the coefficient estimates given in Figures 3, 11, 12, and 13. Each cell gives the median coefficient estimate across the 48 technologies. “Early” refers to the portion of the sample, within each technology, on or before the year at which half of the mentions of the ICT have occurred. “Late” refers to the remainder of the sample period.

of job titles across occupations. That the estimates are not diminished by including 4-digit SOC by year fixed effects indicate that our benchmark estimates are not spuriously reflecting unobserved factors coincident with the introduction of new technologies.

Finally, we explore differences across time in the relationships between tasks and technologies. For each technology, we begin by splitting the sample into two halves, depending on the timing of mentions of the technology within our newspaper text. For example, half of the mentions of Lotus Notes occurred in ads on or before 1989; the other half occurred in ads after 1989. For FORTRAN, this median date is 1978. Given this, we estimate the relationships between task and technology mentions for Lotus Notes for 1960-89 and 1990-2000, separately, and for 1960-1978 and 1979-2000, separately, for FORTRAN, and so on.

Figure 14 presents the result of this exercise. Two patterns emerge from these plots. First, when looking across technologies, the relationships between tasks and technologies are largely stable: Technologies that have a relatively strong association with a given task in the beginning of the sample also tend to have a relatively strong association with the same task during the latter half of the sample. Second, the association between technologies and nonroutine analytic and interactive tasks (averaging across technologies) is larger during the latter portion of the sample, while the relationship between technologies and routine tasks is weaker during the second half of the sample.

Figure 14: Relationship between Task and Technology Mentions

Notes: Each panel plots the relationship between coefficient estimates, the β_{hk} , in the “Early” (x-axis) and the “Late” (y-axis) parts of the sample. The caption of Table 4 defines “Early” and “Late.” The caption within each panel gives the median β_{hk} in the two halves of the sample, as well as the correlation across the two portions of the sample.

D Baseline Model Derivations

D.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 :

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

where ϵ is the worker's idiosyncratic efficiency term, which varies across occupations and ICTs.

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 condition

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

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 so as to maximize his payoff:

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

subject to his unit time endowment

$$\sum_{h=1}^H l_{hgjk} = 1.$$

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

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

Using the optimal time allocation in the program above, we get that the worker's payment per efficiency unit of labor, conditional on working in occupation j , is

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

Note that earnings are $w_{gjk}\epsilon$.

D.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

$$\lambda_{gjk} = \frac{w_{gjk}^\theta}{\sum_{k'=0}^K \sum_{j'=1}^J w_{gj'k'}^\theta}.$$

We aggregate this labor supply at different levels, in order 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:

$$\lambda_{gj} = \sum_{k=0}^K \lambda_{gjk} = \sum_{k=0}^K \frac{w_{gjk}^\theta}{\sum_{k'=0}^K \sum_{j'=1}^J w_{gj'k'}^\theta}.$$

D.3 ICT Market Clearing

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

amount of workers, $\lambda_{gjk}L_g$. 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.

D.4 Analytical Results for the Simple Model in Section

We simplify the environment to: (i) two occupations, with $a_j = 1/2$; (ii) two ICTs with $\bar{\alpha}_1 = \bar{\alpha}_2 = \bar{\alpha}$; and (iii) two types of workers (with $L_g = L$ for each group). We assume that parameters are such that at the baseline, $\bar{W}_g = \bar{W}$, $\forall g$.

Market clearing dictates that

$$a_j \left(\frac{p_j}{P} \right)^{1-\sigma} E = \frac{1}{\bar{\alpha}} \sum_{g=1}^G \sum_{k=1}^2 \bar{W}_g \lambda_{gjk} L_g,$$

which implies relative prices relate to relative supplies according to

$$\left(\frac{p_j}{p_{j'}} \right)^{1-\sigma} = \frac{\sum_{g=1}^G \bar{W}_g L_g \sum_{k=1}^2 \lambda_{gjk}}{\sum_{g'=1}^G \bar{W}_{g'} L_{g'} \sum_{k'=1}^2 \lambda_{g'j'k'}}.$$

Optimal sorting across occupations and ICTs states

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

Substituting optimal sorting into market clearing, and using $\bar{W}_g = \bar{W}$ and $L_g = L$, we obtain

$$\left(\frac{p_j}{p_{j'}} \right)^{1-\sigma} = \frac{\sum_{g=1}^G \sum_{k=1}^2 \left(p_j^{1/\bar{\alpha}} c_k^{(\bar{\alpha}-1)/\bar{\alpha}} \prod_{h=1}^H S_{gh}^{(\alpha_{hjk}/\bar{\alpha})} \right)^\theta}{\sum_{g'=1}^G \sum_{k'=1}^2 \left(p_{j'}^{1/\bar{\alpha}} c_{k'}^{(\bar{\alpha}-1)/\bar{\alpha}} \prod_{h'=1}^H S_{g'h'}^{(\alpha_{h'j'k'}/\bar{\alpha})} \right)^\theta},$$

which finally yields equilibrium relative prices given by:

$$\frac{p_j}{p_{j'}} = \left[\frac{\sum_{g=1}^G \sum_{k=1}^2 \left(c_k^{(\bar{\alpha}-1)/\bar{\alpha}} \prod_{h=1}^H S_{gh}^{(\alpha_{hjk}/\bar{\alpha})} \right)^\theta}{\sum_{g'=1}^G \sum_{k'=1}^2 \left(c_{k'}^{(\bar{\alpha}-1)/\bar{\alpha}} \prod_{h'=1}^H S_{g'h'}^{(\alpha_{h'j'k'}/\bar{\alpha})} \right)^\theta} \right]^{\frac{\bar{\alpha}}{\bar{\alpha}(1-\sigma)-\theta}}.$$

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.

D.5 Derivations of Hat Algebra

1. Occupational-output markets clear

$$\begin{aligned} \left(\hat{p}_j / \hat{P} \right)^{1-\sigma} \hat{E} a_j (p_j / P)^{1-\sigma} E &= \sum_{g=1}^G \widehat{W}_g \bar{W}_g \sum_{k=0}^K \lambda_{gjk} \hat{\lambda}_{gjk} \hat{L}_g L_g + \sum_{g=1}^G \sum_{k=1}^K \hat{c}_k \hat{\Omega}_{jk} c_k \Omega_{gjk} \\ \left(\hat{p}_j / \hat{P} \right)^{1-\sigma} \hat{E} \Psi_j &= \frac{1}{E} \sum_{g=1}^G \sum_{k=0}^K \widehat{W}_g \hat{\lambda}_{gjk} \hat{L}_g \bar{W}_g \lambda_{gjk} L_g + \frac{1}{E} \sum_{g=1}^G \sum_{k=1}^K \hat{c}_k \hat{\kappa}_{jk} \hat{\lambda}_{gjk} \hat{L}_g c_k \Omega_{gjk} \\ \left(\hat{p}_j / \hat{P} \right)^{1-\sigma} \hat{E} \Psi_j &= \Xi \sum_{g=1}^G \sum_{k=0}^K \widehat{W}_g \hat{\lambda}_{gjk} \hat{L}_g \chi_{gjk} + (1 - \Xi) \sum_{g=1}^G \sum_{k=1}^K \xi_{gjk} \hat{c}_k \hat{\kappa}_{gjk} \hat{\lambda}_{gjk} \hat{L}_g. \end{aligned}$$

where Ψ_j is the share of 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. 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

$$\begin{aligned} c_k \kappa_{gjk} &= (1 - \bar{\alpha}_k) \frac{\bar{W}_g}{\bar{\alpha}_k} \\ \hat{c}_k \hat{\kappa}_{gjk} &= \widehat{W}_g. \end{aligned}$$

which implies

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

Recall, too, that our definition of c_k implies

$$\hat{c}_k = \hat{P}\hat{\bar{c}}_k.$$

3. Income

$$\begin{aligned} E &= \sum_{g=1}^G \left(\bar{W}_g L_g + \sum_{j=1}^J \sum_{k=1}^K c_k \Omega_{gjk} \right) \\ E\hat{E} &= \sum_{g=1}^G \left(\widehat{\bar{W}}_g \hat{L}_g \bar{W}_g L_g + \sum_{j=1}^J \sum_{k=1}^K \hat{c}_k \hat{\kappa}_{gjk} c_k \Omega_{gjk} \right) \\ \hat{E} &= \Xi \sum_{g=1}^G \widehat{\bar{W}}_g \hat{L}_g \zeta_g + (1 - \Xi) \sum_{g=1}^G \sum_{j=1}^J \sum_{k=1}^K \hat{c}_k \hat{\kappa}_{gjk} \hat{\lambda}_{gjk} \hat{L}_g \xi_{gjk}. \end{aligned}$$

where ζ_g is the share of group g in total payments to labor (i.e., $\zeta_g \equiv \sum_{j=1}^J \sum_{k=0}^K \chi_{gjk}$). That is, changes in income reflect changes in all factor payments.

4. Employment shares

$$\begin{aligned} \hat{\lambda}_{gjk} \lambda_{gjk} &= \frac{\hat{w}_{gjk}^\theta w_{gjk}^\theta}{\sum_{j'=1}^J \sum_{k'=0}^K \hat{w}_{gj'k'}^\theta \hat{w}_{gj'k'}^\theta} \Rightarrow \\ \hat{\lambda}_{gjk} &= \frac{\hat{w}_{gjk}^\theta}{\sum_{j'=1}^J \sum_{k'=0}^K \hat{w}_{gj'k'}^\theta \lambda_{gj'k'}}. \end{aligned}$$

5. Wages per efficiency unit of labor

$$\begin{aligned} w_{gjk} &= p_j^{\frac{1}{\bar{\alpha}_k}} (c_k)^{-\frac{1-\bar{\alpha}_k}{\bar{\alpha}_k}} \prod_{h=1}^H S_{gh}^{\frac{\alpha_{hjk}}{\bar{\alpha}_k}}, \text{ and} \\ \hat{w}_{gjk} &= (\hat{p}_j)^{\frac{1}{\bar{\alpha}_k}} (\hat{c}_k)^{-\frac{1-\bar{\alpha}_k}{\bar{\alpha}_k}}. \end{aligned}$$

6. Average wages

$$\begin{aligned} \bar{W}_g^\theta &= \sum_{j=1}^J \sum_{k=0}^K w_{gjk}^\theta \\ \widehat{\bar{W}}_g^\theta &= \sum_{j=1}^J \sum_{k=0}^K \lambda_{gjk} \hat{w}_{gjk}^\theta. \end{aligned}$$

7. Price index

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

8. Changes in aggregate task content

$$T_{hj} \equiv \sum_{g=1}^G \sum_{k=0}^K \frac{\alpha_{hjk}}{\bar{\alpha}_k} \cdot L_g \lambda_{gjk}, \text{ and}$$

$$\hat{T}_{hj} = \frac{\sum_{g=1}^G \sum_{k=0}^K \frac{\alpha_{hjk}}{\bar{\alpha}_k} \cdot L_g \lambda_{gjk} \hat{\lambda}_{gjk} \hat{L}_g}{\sum_{g=1}^G \sum_{k=0}^K \frac{\alpha_{hjk}}{\bar{\alpha}_k} \cdot L_g \lambda_{gjk}}.$$

D.6 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, $\lambda_{gj} = \sum_{k=0}^K \lambda_{gjk}$, (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, we estimate (iv) α_{hjk} using the newspaper data, as explained in Sections 3 and 4.4.

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

$$\left(\frac{\lambda_{gjkt}}{\lambda_{gj0t}} \right)^{1/\theta} = \left(\frac{c_{kt}}{p_{jt}} \right)^{1 - \frac{1}{\bar{\alpha}_k}} \prod_{h=1}^H (S_{gh})^{\frac{\alpha_{hjk}}{\bar{\alpha}_k} - \alpha_{hj0}}.$$

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:

$$\frac{\left(\frac{\lambda_{gjkt}}{\lambda_{gj0t}} \right)^{1/\theta}}{\left(\frac{\lambda_{g'jkt}}{\lambda_{g'j0t}} \right)^{1/\theta}} = \prod_{h=1}^H \left(\frac{S_{gh}}{S_{g'h}} \right)^{\frac{\alpha_{hjk}}{\bar{\alpha}_k} - \alpha_{hj0}}$$

$$\frac{\left(\frac{\lambda_{gjkt}}{\lambda_{gj0t}} \right)}{\left(\frac{\lambda_{g'jkt}}{\lambda_{g'j0t}} \right)} = \prod_{h=1}^H \left(\frac{S_{gh}}{S_{g'h}} \right)^{\theta \frac{\alpha_{hjk}}{\bar{\alpha}_k} - \theta \alpha_{hj0}}.$$

Because $\lambda_{gjkt} = \Pr(j, k|g, t) = \Pr(j|g, t) \cdot \Pr(k|j, g, t) = \lambda_{gjt} \cdot \pi_{gjkt}$, we can take logs and

re-arrange to write an expression for $\log \left(\frac{\pi_{gjk t}}{\pi_{\bullet jk t}} \right)$, which we define as the (log) ratio of ICT k adoption within occupation j for group g workers relative to the average ICT k adoption rate within occupation j across all workers:

$$\begin{aligned}
\log \left(\frac{\pi_{gjk t}}{\pi_{gj0t}} \right) - \log \left(\frac{\pi_{g'jk t}}{\pi_{g'j0t}} \right) &= \theta \sum_{h=1}^H \left[\frac{\alpha_{hjk}}{\bar{\alpha}_k} - \alpha_{hj0} \right] [\log S_{gh} - \log S_{g'h}] \\
\log \left(\frac{\pi_{gjk t}}{\pi_{gj0t}} \right) - \log \left(\frac{\pi_{\bullet jk t}}{\pi_{\bullet j0t}} \right) &= \theta \sum_{h=1}^H \left[\frac{\alpha_{hjk}}{\bar{\alpha}_k} - \alpha_{hj0} \right] \left[\log S_{gh} - \sum_{g'=1}^G \frac{L_{g'} \lambda_{g'jt}}{\sum_{g''=1}^G L_{g''} \lambda_{g''jt}} \log S_{g'h} \right] \\
\log \left(\frac{\pi_{gjk t}}{\pi_{gj0t}} \right) &= \log \left(\frac{\pi_{\bullet jk t}}{\pi_{\bullet j0t}} \right) + \theta \sum_{h=1}^H \left[\frac{\alpha_{hjk}}{\bar{\alpha}_k} - \alpha_{hj0} \right] \left[\log S_{gh} - \sum_{g'=1}^G \frac{L_{g'} \lambda_{g'jt}}{\sum_{g''=1}^G L_{g''} \lambda_{g''jt}} \log S_{g'h} \right] \\
\frac{\pi_{gjk t}}{\pi_{gj0t}} &= \frac{\pi_{\bullet jk t}}{\pi_{\bullet j0t}} \cdot \exp \left[\theta \sum_{h=1}^H \left[\frac{\alpha_{hjk}}{\bar{\alpha}_k} - \alpha_{hj0} \right] \left[\log S_{gh} - \sum_{g'=1}^G \frac{L_{g'} \lambda_{g'jt}}{\sum_{g''=1}^G L_{g''} \lambda_{g''jt}} \log S_{g'h} \right] \right].
\end{aligned}$$

The terms on the right hand side are directly observable or estimated. The $\frac{L_{g'} \lambda_{g'jt}}{\sum_{g''=1}^G L_{g''} \lambda_{g''jt}}$ come from the decennial census, the $\frac{\alpha_{hjk}}{\bar{\alpha}_k}$ from our Section 3 regressions, and the $\log S_{gh}$ come from our Section 4.2 model estimation. We use these expressions to impute π_{gjk} , on the basis of π_{jk} , which we actually observe.

Expenditure on ICT k Next we use the data to build total expenditure in ICT k , using the market clearing equation:

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

Manipulating the right-hand side, we get

$$\begin{aligned}
c_k \Omega_{gjk} &= (1 - \bar{\alpha}_k) \frac{\bar{W}_g}{\bar{\alpha}_k} \times \lambda_{gjk} L_g \\
&= (1 - \bar{\alpha}_k) \frac{\bar{W}_g}{\bar{\alpha}_k} \times \underbrace{\frac{\lambda_{gjk}}{\sum_{k'=0}^K \lambda_{gjk'}}}_{=\pi_{gjk}} \left(\sum_{k=0}^K \lambda_{gjk} \right) L_g \\
&= \frac{(1 - \bar{\alpha}_k)}{\bar{\alpha}_k} \bar{W}_g L_g \lambda_{gj} \pi_{gjk},
\end{aligned}$$

where we remove λ_{gjk} and instead use π_{gjk} , a variable which we now observe. Furthermore, the calibration of $\bar{\alpha}_k$ is discussed in footnote 24. Finally, $\bar{W}_g L_g \lambda_{gj}$ equals the wage bill of

group g workers in occupation j . This object is observable in the decennial census. 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:

$$\begin{aligned}
E &= \sum_{g=1}^G \left\{ \bar{W}_g L_g + \sum_{j=1}^J \sum_{k=1}^K c_k \Omega_{gjk} \right\} \\
&= \sum_{g=1}^G \bar{W}_g L_g + \sum_{j=1}^J \sum_{k=1}^K \sum_{g=1}^G c_k \Omega_{gjk} \\
&= \sum_{g=1}^G \bar{W}_g L_g + \sum_{j=1}^J \sum_{k=1}^K \frac{(1 - \bar{\alpha}_k)}{\bar{\alpha}_k} \sum_{g=1}^G \pi_{gjk} \bar{W}_g L_{gj},
\end{aligned}$$

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 a value for $1 - \Xi$.

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

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

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

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

Occupational shares Now we compute shares related to the importance of each occupation.

Recall that the total payment to occupation j firms that employ group g workers is

$$\sum_{k=0}^K \bar{W}_g \lambda_{gjk} L_g + \sum_{k=1}^K c_k \kappa_{gjk} \lambda_{gjk} L_g.$$

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

$$\begin{aligned}\frac{\sum_{k=0}^K \bar{W}_g \lambda_{gjk} L_g}{\sum_{k=0}^K \lambda_{gjk} L_g} + \frac{\sum_{k=1}^K c_k \kappa_{gjk} \lambda_{gjk} L_g}{\sum_{k=0}^K \lambda_{gjk} L_g} &= \bar{W}_g + \frac{\sum_{k=1}^K (1 - \bar{\alpha}_k) \frac{\bar{W}_g}{\bar{\alpha}_k} \pi_{gjk} L_g \lambda_{gj}}{\sum_{k=0}^K \lambda_{gjk} L_g} \\ &= \bar{W}_g + \bar{W}_g \sum_{k=1}^K \frac{(1 - \bar{\alpha}_k)}{\bar{\alpha}_k} \pi_{gjk}.\end{aligned}$$

Total payments to occupation j (both workers and ICT) is given by the following expression, where we use Λ_{gj} to denote the number of workers from group g who work in occupation j

$$\begin{aligned}\psi_j &= \sum_{g=1}^G \Lambda_{gj} \times \text{average payment to occupation } j, \text{ group } g \\ &= \sum_{g=1}^G \Lambda_{gj} \left\{ \bar{W}_g + \bar{W}_g \sum_{k=1}^K \frac{(1 - \bar{\alpha}_k)}{\bar{\alpha}_k} \pi_{gjk} \right\} \\ &= \sum_{g=1}^G \Lambda_{gj} \bar{W}_g \left\{ 1 + \sum_{k=1}^K \frac{(1 - \bar{\alpha}_k)}{\bar{\alpha}_k} \pi_{gjk} \right\}\end{aligned}$$

The share we are looking for is

$$\Psi_j = \psi_j / \sum_{j'=1}^J \psi_{j'}.$$

To calibrate this share, note that $\Lambda_{gj}\bar{W}_g$ equals the wage bill of group g workers in occupation j , which is observable in the decennial census. The $\frac{(1-\bar{\alpha}_k)}{\bar{\alpha}_k}\pi_{gjk}$ terms can be computed using calculations we have described above.

E Model Extension I: Fixed Occupational-ICT Shares

The idea is to start from the data as an equilibrium of the model with fully flexible labor supply, and then compute changes in the reaction to shocks if occupation-ICT shares, λ_{gjk} , are fixed at their original equilibrium values. We start with the equations that describe the new equilibrium.

E.1 Equilibrium

1. Occupational output markets clear

$$a_j \left(\frac{p_j}{P} \right)^{1-\rho} E = \sum_{g=1}^G \sum_{k=0}^K \bar{W}_{gjk} \bar{\lambda}_{gjk} L_g + \sum_{g=1}^G \sum_{k=1}^K c_k \kappa_{gjk} \bar{\lambda}_{gjk} L_g,$$

where now $\bar{W}_{gjk}$ is the average wage of workers who were sorted in cell g, j, k in the baseline equilibrium, under the new prices.

2. ICT markets clear

$$\kappa_{gjk} = \frac{1 - \bar{\alpha}_k}{\bar{\alpha}_k} \frac{\bar{W}_{gjk}}{c_k}.$$

3. Income definition

$$E = \sum_{j=1}^J \sum_{g=1}^G \sum_{k=0}^K \bar{W}_{gjk} \bar{\lambda}_{gjk} L_g + \sum_{g=1}^G \sum_{k=1}^K c_k \kappa_{gjk} \bar{\lambda}_{gjk} L_g.$$

4. Efficiency wages

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

5. Average wages

$$\begin{aligned}\bar{W}_{gjk} &= \int w_{gjk} \epsilon_{gjk} dG \left(\epsilon_{gjk} | \epsilon_{gjk} \tilde{w}_{gjk} = \max_{j', k'} \{ \epsilon_{gj'k'} \tilde{w}_{gj'k'} \} \right) \\ &\Rightarrow \\ \bar{W}_{gjk} &= \frac{w_{gjk}}{\tilde{w}_{gjk}} \cdot \bar{W}_g(\tilde{w}_{gjk}),\end{aligned}$$

where tildes denote the baseline equilibrium, and where

$$\bar{W}_g(\tilde{w}_{gjk}) = \Gamma(1 - 1/\theta) \cdot \left(\sum_{j=1}^J \sum_{k=0}^K \tilde{w}_{gjk}^\theta \right)^{1/\theta}.$$

6. Price index

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

In this system, the endogenous variables are $\{p_j\}_{j=1}^J$, $\{\kappa_{gjk}\}_{g,j,k}$, $E, \{w_{gjk}\}_{g,j,k}$, $\{\bar{W}_{gjk}\}_{g,j,k}$, P . Note that as opposed to only G average wages (as in the fully flexible model), we instead have $G \cdot J \cdot K$ average wages, reflecting that, given that quantities cannot adjust, wages must.

E.2 Derivations of hat algebra

1. Occupational-output markets clear

$$\left(\hat{p}_j / \hat{P} \right)^{1-\sigma} \hat{E} \Psi_j = \Xi \sum_{k=0}^K \sum_{g=1}^G \widehat{W}_{gjk} \chi_{gjk} \hat{L}_g + (1 - \Xi) \sum_{k=1}^K \sum_{g=1}^G \hat{c}_k \hat{\kappa}_{gjk} \xi_{gjk} \hat{L}_g.$$

2. ICT markets clear

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

with

$$\hat{c}_k = \hat{P} \hat{\tilde{c}}_k.$$

3. Income

$$\hat{E} = \Xi \sum_{k=0}^K \sum_{g=1}^G \sum_{j=1}^J \widehat{W}_{gjk} \chi_{gjk} \hat{L}_g + (1 - \Xi) \sum_{k=1}^K \sum_{g=1}^G \sum_{j=1}^J \widehat{W}_{gjk} \xi_{gjk} \hat{L}_g.$$

4. Efficiency wages

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

5. Average wages

$$\widehat{W}_{gjk} = \hat{w}_{gjk}.$$

6. Price index

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

Here, the endogenous variables are the changes of the endogenous variables defined above.

In addition, we exploit the following formulas for changes in task content quantity and value:

$$\begin{aligned} \hat{T}_{hj} &= \frac{\sum_{k=0}^K \sum_{g=1}^G \frac{\alpha_{hjk}}{\hat{\alpha}_k} \cdot L_g \bar{\lambda}_{gjk} \hat{L}_g}{\sum_{k=0}^K \sum_{g=1}^G \frac{\alpha_{hjk}}{\hat{\alpha}_k} \cdot L_g \bar{\lambda}_{gjk}}. \\ \widehat{VT}_{hj} &= \frac{\sum_{k=0}^K \sum_{g=1}^G \frac{\alpha_{hjk}}{\hat{\alpha}_k} \cdot \widehat{W}_g \hat{L}_g \bar{W}_g \bar{\lambda}_{gjk} L_g}{\sum_{k=0}^K \sum_{g=1}^G \frac{\alpha_{hjk}}{\hat{\alpha}_k} \cdot \bar{W}_g \bar{\lambda}_{gjk} L_g}. \end{aligned}$$

E.3 Results

In Section 5.2, we present the main results of the exercise wherein ICT-occupation shares are held fixed in response to a decline in ICT prices. There, we explored shifts in groups' earnings and in occupations' task content, contrasting our benchmark analysis with the equilibrium in which demographic groups are fixed in their ICT-occupation choice.

Building on this analysis, we present an additional comparison in Figure 15. We contrast changes in occupations' task value (incorporating both changes in the quantity of tasks performed by workers and the price associated with these tasks) in our extension (with fixed allocations across ICT-occupation pairs) and in our benchmark specification. The main takeaway from this figure is that changes in the value of tasks performed within each occupation is substantially less dispersed (across occupations) when workers do not reallocate.

Figure 15: Counterfactual Changes in Occupations' Task Content

Notes: For each of the five task measures, the panels plot the relationship between changes in task value in the benchmark specification (x-axis) against changes in task value when workers do not reallocate across ICT-occupation pairs (y-axis).

F Model Extension II: Non-employment Margin

F.1 Model

We now allow for a non-market occupation. We still denote market occupations by $j = 1, \dots, J$, and, when convenient, we denote non-employment by occupation $j = 0$.

We assume that non-employment generates a non-market benefit b_g , which we do not observe directly, and which rationalizes the fraction of people of each group in non-employment. To retain as much as possible from our original framework, we assume that the reward to non-employment is also proportional to efficiency units and is expressed in units of the numeraire, but does not depend directly on ICT availability. Hence, the worker unobserved efficiency vector has now dimension $J \cdot K + 1$, where we denote the new element by ϵ_{g0} .

Most results will carry, after appropriately redefining our variables. As before, let $W_g(\epsilon)$ denote a random variable which is the total compensation (or earnings) that a person from group g , with draws ϵ obtains from *market occupations*:

$$W_g(\epsilon) = \max_{j,k} \{w_{gjk} \epsilon_{gjk}\}.$$

We know W_g is a Fréchet random variable with mean $\mathbb{E}[W_g(\epsilon)] = \Gamma(1 - 1/\theta) \cdot \left(\sum_{k=0}^K \sum_{j=1}^J w_{gjk}^\theta\right)^{1/\theta}$, and shape parameter θ .

Employment and non-employment shares To capture the non-employment margin, start by defining $Z_g(\epsilon) \equiv \max\{W_g(\epsilon), b_g \epsilon_{g0}\}$. Since W_g and $b_g \epsilon_{g0}$ are Fréchet random variables, so is $Z_g(\epsilon)$. It follows that the fraction of people in group g in non-employment is given by:

$$\eta_{g0} = \frac{b_g^\theta}{\sum_{k'=0}^K \sum_{j'=1}^J w_{gj'k'}^\theta + b_g^\theta}.$$

Letting all people in group g – i.e., both employed and not, be N_g , the mass of workers from group g , which we denote by L_g as before, is now $L_g = (1 - \eta_g) N_g$. In turn, the fraction of *all people* (not only workers) in group g that sort into market cell j, k is

$$\begin{aligned} \eta_{gjk} &\equiv \frac{w_{gjk}^\theta}{\sum_{j'=1}^J \sum_{k'=0}^K w_{gj'k'}^\theta + b_g^\theta}, \\ &= \frac{w_{gjk}^\theta}{\sum_{j'=1}^J \sum_{k'=0}^K w_{gj'k'}^\theta} \times \frac{\sum_{j'=1}^J \sum_{k'=0}^K w_{gj'k'}^\theta}{\sum_{j'=1}^J \sum_{k'=0}^K w_{gj'k'}^\theta + b_g^\theta}. \end{aligned}$$

Note for future reference that we can also write $\eta_{gjk} = \lambda_{gjk}(1 - \eta_{g0})$, where λ_{gjk} is, as before, the fraction of g employed workers that sort into cell j, k ; and η_{g0} is the fraction of individual in group g who are not employed.

Market compensation We now obtain an expression for observed *market earnings*. Our data on $\bar{W}_g$ corresponds to the average market compensation, conditional on individuals being on the market. By properties of the Fréchet distributions, we know that $\mathbb{E}[W_g | W_g > \epsilon_b b_g] = \mathbb{E}[Z_g]$, so we conclude that

$$\bar{W}_g = \Gamma(1 - 1/\theta) \cdot \left(\sum_{j=1}^J \sum_{k=0}^K w_{gjk}^\theta + b_g^\theta \right)^{1/\theta}.$$

F.2 Estimation

We take the stance that we can estimate our parameters using the same moments as before

$$\begin{aligned} \tilde{\lambda}_{gj} &= \sum_{k=0}^K \left[\frac{w_{gjk}^\theta(\Theta)}{\sum_{j=1}^J \sum_{k'=0}^K w_{gj'k'}^\theta(\Theta)} \right], \\ \tilde{\pi}_{jk} &= \sum_{g=1}^G \frac{\lambda_{gjk}(\Theta) \tilde{L}_{gj}}{\sum_{g'=1}^G \tilde{L}_{g'j}}, \text{ and} \\ \tilde{\bar{W}}_g &= \Gamma(1 - 1/\theta) \cdot \left(\sum_{j=1}^J \sum_{k=0}^K w_{gjk}^\theta(\Theta) + b_g^\theta \right)^{1/\theta}. \end{aligned}$$

This would amount to choosing parameters b_g for each decade as to perfectly fit the data on the non-employment margin, summarized by a new set of G moments

$$\tilde{\eta}_{g0} = \frac{b_g^\theta}{\sum_{k'=0}^K \sum_{j'=1}^J w_{gj'k'}^\theta(\Theta) + b_g^\theta}.$$

By doing so, we retain our original $\hat{S}_{gh}$ estimates. For the purposes of counterfactual calculations, as we show below, all the information about b_g is contained in the perfectly observable non-employment shares, η_{g0} .

F.3 Simulation

With this formulation, we can retain most of our previous hat algebra equations:

$$\left(\hat{p}_j/\hat{P}\right)^{1-\sigma} \hat{E}\Psi_j = \Xi \sum_{g=1}^G \sum_{k=0}^K \widehat{W}_g \hat{\lambda}_{gjk} \hat{L}_g \chi_{gjk} + (1 - \Xi) \sum_{g=1}^G \sum_{k=1}^K \xi_{gjk} \hat{c}_k \hat{\kappa}_{gjk} \hat{\lambda}_{gjk},$$

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

$$\hat{c}_k = \hat{P} \hat{\hat{c}}_k,$$

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

$$\hat{\lambda}_{gjk} = \frac{\hat{w}_{gjk}^\theta}{\sum_{j'=1}^J \sum_{k'=0}^K \hat{w}_{gj'k'}^\theta \lambda_{gj'k'}}, \text{ and}$$

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

But we need to modify our equation for observed wages:

$$\begin{aligned} \widehat{W}_g &= \left(\sum_{j=1}^J \sum_{k=0}^K \eta_{gjk} \hat{w}_{gjk}^\theta + \eta_{g0} \hat{b}_g^\theta \right)^{1/\theta}, \\ &= \left((1 - \eta_{g0}) \sum_{j=1}^J \sum_{k=0}^K \lambda_{gjk} \hat{w}_{gjk}^\theta + \eta_{g0} \hat{b}_g^\theta \right)^{1/\theta}. \end{aligned}$$

Since we assume b_g is expressed in units of the final good, we must add the equation

$$\hat{b}_g^\theta = \hat{P}.$$

And we must also explicitly account for the non-employment margin

$$\hat{L}_g = (\widehat{1 - \eta_{g0}}) \hat{N}_g,$$

where³⁵

$$\widehat{(1 - \eta_{g0})} = \sum_{j=1}^J \sum_{k=0}^K \hat{\eta}_{gjk} \lambda_{gjk},$$

and

$$\hat{\eta}_{gjk} = \frac{\hat{w}_{gjk}^\theta}{\sum_{j'=1}^J \sum_{k'=0}^K \hat{w}_{gj'k'}^\theta \eta_{gj'k} + \eta_{g0} \hat{b}_g^\theta}.$$

F.4 Results

In this section, we compare our baseline results of removing all ICTs to those we obtain when we add an additional extensive margin. The left panel of Figure 16 shows, as we assert in the main body of the paper, that low income groups transition more frequently into non-employment in response to the shock. In fact the model generates quite sizable proportional transitions for the lowest income groups. However, as we have explained before, the model builds in a very strong force for this to happen, based on selection. To assess the strength of earnings falling less for low income workers, the right panel of Figure 16 plots the results of the same simulation, this time assuming (counterfactually) that the baseline year-2000 employment share in each group is the same and equal to 0.1. The simulation shows that, because low earning workers' returns from working in the market fall less in response to a removal of ICTs, they move less frequently into non-employment. In sum, the arrival of ICTs imposes a strong force pushing low income workers into non-employment.

³⁵To see why, note that

$$\begin{aligned} \widehat{(1 - \eta_{g0})} &= \frac{1 - \eta'_{g0}}{1 - \eta_{g0}} \\ &= \frac{\sum_{j=1}^J \sum_{k=0}^K \eta'_{gjk}}{1 - \eta_{g0}}. \end{aligned}$$

Since

$$\eta'_{gjk} = \hat{\eta}_{gjk} \eta_{gjk},$$

we conclude that

$$\begin{aligned} \widehat{(1 - \eta_{g0})} &= \sum_{j=1}^J \sum_{k=0}^K \hat{\eta}_{gjk} \frac{\eta_{gjk}}{1 - \eta_{g0}} \\ &= \sum_{j=1}^J \sum_{k=0}^K \hat{\eta}_{gjk} \lambda_{gjk}. \end{aligned}$$

Figure 16: The Impact of All 48 ICTs on Occupations' Tasks and Groups' Earnings (Active Non-employment Margin)

Notes: Each panel plots the relationship between groups' baseline equilibrium earnings (x-axis) and the counterfactual percentage point change in employment rates (y-axis). The left panel incorporates observed employment rates. In the right panel, the baseline equilibrium employment for all 40 g groups is set to the same level.