LABOR DEMAND AND THE MACROECONOMY

A Dissertation
Presented to the Faculty of the Graduate School
of Cornell University
in Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy

by
Zongyang Li
May 2024

LABOR DEMAND AND THE MACROECONOMY

Zongyang Li, Ph.D.

Cornell University 2024

My research lies at the intersection of macroeconomics and labor economics, with a particular focus on the demand side of the labor market. Broadly, my work can be characterized as that of an applied macroeconomist studying labor economic issues by applying macroeconomic models and applied labor tools. I combine rich empirical analyses with models of heterogeneous workers that incorporate labor market dynamics, firm dynamics, and macroeconomic outcomes to investigate the research questions. I discuss three chapters of my ongoing projects in this dissertation.

In the first chapter, using worker surveys and online job posting data, I document that the U.S. economy has seen a substantial increase in the mixing of skill requirements from 2005-2018, both for incumbent jobs and newly posted vacancies. American workers increasingly work in occupations that demand mixtures of analytical, computer, and interpersonal skills rather than specializing in one of them, even within granular occupations. This change occurred primarily in low- to medium-wage occupations, and workers in occupations that increasingly mix non-routine skills, or those with a broader set of these skills earn a wage premium.

In the second chapter, I build a multi-dimensional directed search and matching model with two-sided heterogeneity and endogenous choices to understand the

sources of skill mixing shifts,. In this framework, firms optimally choose occupations' skill intensities before producing with a worker. Simultaneously, workers make decisions about their jobs as well as their life-time skill development trajectories. Counterfactual analysis shows that the rise in the complementarity of skills in production and in the cost of skills for occupation operation are the main drivers of skill mixing shifts and the corresponding wage and employment dynamics in this period.

In the third chapter, I study the spatial general equilibrium and redistribution effects of e-commerce on different local labor markets from a trade perspective based on the production technology change in the retail sector. Using a panel of products and retailers on Amazon, I document that online retailers are more agglomerated in space, particularly for those using Amazon's distribution and fulfillment centers, and their agglomeration is related to higher trade flows of the upstream goods. I then incorporate consumer search and retailer location choices into a multi-sector gravity trade model with an elastic supply of heterogeneous workers. The model implies that the increase in online shopping efficiency, the rise in online retailer agglomeration, and the reduction in shipping friction will induce greater industrial and occupational specialization. Quantitative analysis shows that the growth of Amazon from 2007-2017 had led to overall declines in retail prices, but also worker reallocation out of manufacturing sector, resulting in a 1 percent decrease in welfare. Non-employment increases by 2.3 percentage points and the Gini index on employment across regions increases by near 20 percent, exacerbating regional inequality.

BIOGRAPHICAL SKETCH

Zongyang ”Elmer” Li is a Ph.D. Candidate in Economics at Cornell University, where his research endeavors focus primarily on macroeconomics and labor economics, with secondary interests in industry dynamics and trade. Born and raised in rural China before moving to an urban setting, Li’s journey is marked by his resilience in overcoming challenges associated with being a first-generation immigrant, including cultural and language barriers. This unique background has infused him with a profound empathy for others facing similar struggles and a dedicated commitment to fostering inclusivity within academic and professional environments.

Li’s educational path has been nothing short of exemplary, obtaining a Master of Arts in Economics from Cornell University, and earlier, a Master of Public Administration from New York University. His research delves into the implications of modern economic forces on diverse populations, aiming to contribute to more equitable economic outcomes. His work has been recognized at several prestigious conferences, including the European Economic Association and the ASSA/AEA Annual Meeting.

Beyond his academic achievements, Li has accumulated a wealth of experience through his roles as a research fellow at the NYC Mayor’s Office for Economic Opportunity, a data analyst at UNDP, and a graduate researcher at the NYU Furman Center for Real Estate and Urban Policy. These positions allowed him to directly engage with policy-making processes, enriching his understanding of economic disparities and contributing to his research’s real-world applicability.

Li’s commitment to service is evident in his extensive mentoring and teaching experiences. As a mentor for the DEI Program at the Cornell Brooks School, he has guided Ph.D. students from diverse backgrounds, and as a Teaching Assistant at Cornell University, he has fostered an inclusive learning environment for his students. His dedication to community service, particularly his mentoring work with underrepresented students at the NYC Urban Assembly Academy, highlights his belief in the transformative power of education and mentorship.

With his unique blend of academic excellence, personal resilience, and commitment to social justice, Zongyang ″Elmer″ Li is poised to make significant contributions to the field of economics and to the broader goal of creating a more equitable and inclusive society.

ACKNOWLEDGEMENTS

As I approach the conclusion of my PhD journey, I find myself overwhelmed with gratitude towards the many individuals who have played pivotal roles in my academic and personal growth. Their support, guidance, and inspiration have been so indepensible for me. I simply could not imagine how can I would complete my PhD journey without their presence in my life.

First, I am heartfeltly grateful for my PhD committee: Professors Philipp Kricher, Michael Lovenheim, and Julieta Caunedo. I met Mike during the AEA meeting in Philly in 2018, an encountering of which ultimately led me to Cornell, and his encouragement and breadth of knowledge have been instrumental for me to grow as a scholar. I have not met any economist who is as caring, patient, and kind as Philipp, whose knowledge and depth have always left me in awe. Philipp opened the door of macro labor for me. I am always blown-away by Julieta, who is an amazing scholar, a tough advisor (sometimes), and who has an incredibly warm and caring heart. During times of difficulty and uncertainty, Julieta has always been my beacon of light.

Second, my gratitude extends to Professors Mathieu Taschereau-Dumouchel, Evan Riehl, Ryan Chahrour, and Sergio Ocampo Díaz, who have devoted considerable time to me, shown tremendous kindness, and provided immense support.

Third, I thank my friends, Nikolai Boboshko, Luming Chen, Wentong Chen, Matthew Comey, Sam Dodini, Zhou Fan, Qiwei He, Zihan Hu, Kyle Yizhou Kuang, Mengwei Lin, Tianli Xia, Binglin Wang, David Wasser, Meredith Welch, Guanyi Yang, Yuanyuan Yang, Lilac Zihui Zhao, Julia Zhu, Zhiyang Zhu, etc. (I cannot name all), for their support and companionship during this journey. The discussions we have had, the suggestions I recieved from you, the support you've offered have been invaluable.

Lastly but importantly, I thank my parents, as well as my partner, Yusheng, whose love, care, and unwavering support was so pivotal to me.

This dissertation stands as a testament not solely to my own efforts but to the collective support and belief of everyone involved in this journey. To each and every one of you, I am eternally grateful.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

THE CHANGING SKILL MIXING IN THE US ECONOMY

1.1 Introduction

The nature of work in the United States has seen significant changes in recent decades. A vast literature documents the decline in the demand for "routine" tasks and associated worker skills due to technological shifts (i.e., Autor, Levy, and Murnane 2003; Acemoglu and Autor 2011) and the growing importance of social skills (Cortes, Jaimovich, and Siu 2021; Deming 2017). However, as occupational skill demands adjust, it remains unclear whether employers are leaning towards specific specialized skills or seeking a broad range of skills. Additionally, if there is a trend towards a *mix* of skill demands, how does this influence workers' returns to occupation and education choices? The degree of skill mixing among occupations carries important and distinct implications: if employers seek specific skills, indicating specialization in skill demand, then workers benefit from becoming experts in those particular skill dimensions; if, however, occupations increasingly require mixtures of different skills, indicating "skill mixing," then multidisciplinary schooling and training become more advantageous.

This paper studies the phenomenon of employer skill mixing, exploring its implications for workers, and seeks to understand the underlying sources of these shifts. The analysis begins with the aggregation of suitable data and the creation of measures to assess skill mixing. For this purpose, I primarily employ the Occupational Information Network (O*NET), which surveys incumbent workers of their current jobs and details the importance of different skill requirements in occupations. By considering extended time periods and focusing on continually updated occupations, I show that O*NET allows a credible analysis of longitudinal changes in skill demand. Supplementing this, Lightcast (formerly known as "Burning Glass") provides real-time skill demand from millions of online job vacancies, enabling the measurement of the extensive margin share of jobs that require specific skills. Equipped with these datasets, I evaluate the degree of skill mixing for each occupation by calculating the cosine similarity between an occupation's skill vector and the unit vector on which skills along several domains are equally important; consequently, this "mixing index" increases as an occupation's demand for different skills gets closer to each other.

Leveraging information about skill demand for both incumbent jobs and newly posted vacancies, this paper presents evidence that from 2005 to 2018, occupations in the United States increasingly demand mixtures of different skills. Using the O*NET dataset, I show that even at the 7-digit occupation level, there is a sizable increase in the degree of skill mixing, particularly for analytical, computer, and interpersonal skills that are considered non-routine.[1] Compared to 2005, the degree of mixing of these skills in 2018 as captured by the skill mixing indexes has increased by 9.2 percentiles on average. The growth of skill mixing is even starker in higher-level 4-digit occupations, by 12.4 percentiles on average and 11 percentiles for constantly updated occupations. For example, in 2005, "Maids and Housekeeping Cleaners" valued interpersonal skill four times over analytical and twice over computer skill. By 2018, analytical skill equaled, and computer skill reached two-thirds of interpersonal skill's importance. Conversely,

---

[1] O*NET's occupational classification is based on the Standard Occupational Classification (SOC) system but offers more granularity. For example, in 2010, O*NET lists 1,110 unique 7-digit occupations, which correspond to 868 unique SOC 7-digit occupations. For analysis at a higher occupational level using census data, I first crosswalk O*NET occupations to the SOC. Subsequently, I employ crosswalks between SOC and census occupations from Autor and Price (2013) and developed by Deming (2017).

for ”Insurance Appraisers, Auto Damage,” computer skill consistently led in importance, but by 2018, analytical skill doubled and interpersonal skill tripled, both surpassing 60 percent of computer skill’s importance.

I highlight two new facts about skill mixing. First, a shift-share decomposition of the rising trend in skill mixing attributes the majority of the increase to changes within occupations, rather than workers’ reallocation across occupations. This pattern distinguishes skill mixing from other labor market changes for which worker reshuffling plays a key role or for which the change is mainly across-occupation.[2] Further decomposition shows that the within-occupation increases in skill mixing persist accounting for workers’ gender, education, and experience, and are robust to alternative measures of skills and indexes of mixing. Second, the most pronounced rise in the mixing of the three non-routine skills appears in service and white-collar occupations, including roles like healthcare givers and housekeepers. Whereas blue-collar occupations, such as operators and machinists, have witnessed a more significant mixing of routine skill and the three non-routine skills. On the other hand, high-wage managerial and professional occupations show relatively limited skill mixing.

The phenomenon of skill mixing bears significant distributional consequences in the labor market. A notable structural shift in the U.S. labor market since the 1980s has been job polarization (Acemoglu and Autor 2011; Goos, Manning, and Salomons 2014), a trend that continues to be evident in the data from 2005 to 2018. Skill mixing emerges as a key factor in explaining these distributional dynamics. For occupations within similar wage ranks in 2005, it is observed

---

[2] For example, in Autor and Dorn (2013), the polarization of the labor market is attributed to the substitution of medium-skill workers in routine jobs and their flow into service jobs; in Deming (2017) across-occupation employment shift drive the rising importance of social skills. Dodini, Lovenheim, and Willen (2022) find that changes in employment concentration across existing occupations account for the skill intensity differential of unionized workers.

that those who have become more skill-mixed experience greater growth in both employment shares and wages. Remarkably, the growth in employment and wages during this period is almost exclusively attributed to occupations that have become more skill-mixed. Therefore, skill mixing provides a novel and multi-dimensional lens to understand these labor market transformations.

To evaluate the impact of skill mixing on workers' labor market outcomes, I estimate the wage returns to skill mixing by combining the National Longitudinal Survey of Youth 1979 and 1997 (NLSY 79 & 97), taking advantage of the rich information on participants' abilities, employment, and educational histories. I find a significant return to skill mixing for both occupational choices and worker skills. To assess the wage premium, I estimate a regression model that incorporates multiple skills and their degrees of mixing for both occupations and individual workers, with worker and occupation fixed effects in the spirit of Abowd, Kramarz, and Margolis (1999) (hereafter AKM). My preferred specifications indicate that workers in occupations that become a standard deviation more mixed in analytical, computer, and interpersonal skills gain a 1.5 percent wage premium; meanwhile, workers who are more mixed in these skills earn 6.5 percent more. I further show some additional returns to skill mixing, both in terms of employment and college major choices.

1.2 Literature Review

I study labor market dynamics emphasizing *skill mixtures* and explore new theoretical perspectives to explain them. The empirical objective aligns with the literature investigating the long-term trend of skill demand and skill-biased technological changes (i.e., Tinbergen 1974, 1975; Katz and Murphy 1992; Autor, Katz, and Krueger 1998; Autor, Levy, and Murnane 2003; Goldin and Katz 2010; Acemoglu and Autor 2011; Autor and Dorn 2013; Deming and Kahn 2018; Deming and Noray 2020).[3] My finding that the within-occupation changes drive skill mixing is consistent with other studies that find a major role played by within-occupation variation for aggregate job attributes (Autor and Handel 2013; Atalay et al. 2020; Freeman, Ganguli, and Handel 2020; Cortes, Jaimovich, and Siu 2021).[4] Unlike these studies, this paper studies skills in their conjunction, i.e., as mixtures, and show that employers do increasingly require mixtures of skills from workers, especially non-routine ones. This paper further finds that skill mixing has important distributional implications for wage and employment and for workers' return in occupation and education choices. The evidence on skill mixing leads to unique policy implications and broadens the understanding of the influence of technological change on the labor market.

Two papers closely related to the empirical phenomenon documented in this paper are Hershbein and Kahn (2018) and Deming (2017). The former illustrates that employers in metropolitan areas hit harder by great recession were more likely to post jobs demanding cognitive and computer skills, particularly in routine-cognitive occupations. My analysis differs by demonstrating that skill mixing occurs for a broad set of skills, within a wide array of granular

---

[3] The changes in relative efficiency of inputs is the focus of the skill-biased technological change (SBTC) literature, and has been shown to successful account for the major U.S. wage dynamics. See for example, Katz and Murphy (1992), Autor, Katz, and Krueger (1998), and Goldin and Katz (2010). This paper incorporates both changes in relative skill efficiency and changes in the skill complementarity, and show the latter's important role in determining skill mixing, wage shifts, and employment distribution post 2000s.

[4] Extracting task information from job ads, Atalay et al. (2020) reveal that the major change in job content during 1950-2000 occurred within-occupation, a pattern that is found to persist post-2000 by Freeman, Ganguli, and Handel (2020). Cortes, Jaimovich, and Siu (2021) show that from 1980 to post-2010, high-paying occupations in the United States require more social skills. Using worker-reported job tasks, Autor and Handel (2013) find that there is significant within-occupation variation in task requirements.

occupations, and is not specific to regions or economic downturns. Deming (2017) highlights that occupations requiring higher math and social skills based on O*NET 1997 have seen increased employment and wage growth from 1980 to 2012. In contrast, I use various versions of O*NET to capture longitudinal changes in skill demand and explore the wage and employment gains stemming from within-occupation skill mixing shifts.

1.3 Evidence of Skill Mixing

In this section, I examine the shifts in the extent of skill mixing in the economy. I start by showing that an angle-based index can effectively measure the magnitude of skill mixing of occupations within a multi-dimensional skill space. Using both O*NET and Lightcast data at varying levels of granularity, I explore the growth in skill mixing, decomposing it into across- and within-occupation changes. I further explore the primary sources of this variation and the differences across occupation groups. Lastly, I underscore the significance of the mixing of different skills by relating it to the changes in employment and wage distributions.

1.3.1 Measures and Data

The Degree of Skill Mixing: To evaluate the degree of skill mixing in an occupation, one can analyze the angular difference between the occupation's skill vector and the unit vector, on which different skill requirements are equivalent. Figure 1.1 illustrates this in the two-dimensional skill space of analytical and interper-

Figure 1.1: Illustrating Skill Mixing

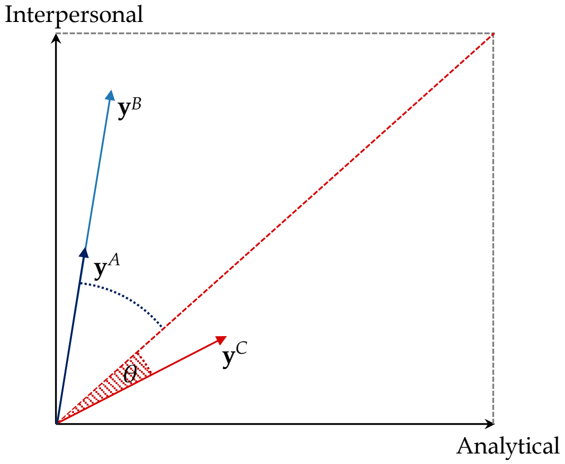

Notes: This figure contrasts three occupations—A, B, and C—in the two dimensional skill space of analytical and interpersonal skills. Each occupation is characterized by its skill vector ($\mathbf{y}^A$, $\mathbf{y}^B$, and $\mathbf{y}^C$), as well as by the angle (θ) between the skill vectors and the 45-degree line.

sonal skills, showing three occupations represented by vectors ($\mathbf{y}^A$, $\mathbf{y}^B$, and $\mathbf{y}^C$). Occupations A and B exhibit greater specialization towards interpersonal skill, as their vectors diverge from the diagonal in the direction of the interpersonal skill axis. Despite varying skill intensities — with occupation B's vector ($\mathbf{y}^B$) being notably longer than occupation A's ($\mathbf{y}^A$) — they share a similar degree of skill mixing, evident from the angle (θ) between their vectors and the diagonal line are the same. The emphasis of skill mixing is on the proportionate use of different skills (indicated by the angle) rather than the overall skill intensity (represented by the vector's length). In contrast, occupation C demonstrates a higher degree of skill mixing, evident from its smaller angle (θ) to the diagonal line.

As θ decreases, indicating a higher degree of skill mixing, $\cos(\theta)$ increases, making it an suitable measure for skill mixing. Building on this concept, I transition from the two-dimensional representation in Figure 1.1 to a multi-dimensional space. I accomplish this by employing the cosine similarity between an occupation's skill vector and a multi-dimensional norm vector.[5] Specifically:

Definition 1 (Degree of Skill Mixing of an occupation) *The skill mixing index for an occupation j in a K-dimensional space characterized by the skill intensity vector* $\mathbf{y}^j = \{y_1^j, ..., y_k^j, \ldots, y_K^j\} \in S \subset \mathbb{R}^{K+}$ *is the cosine similarity between its skill vector and the norm* $\hat{\mathbf{v}}$ *in the skill space.*

$$Mix(\mathbf{y}^j) = \frac{\mathbf{y}^j \hat{\mathbf{v}}}{||\mathbf{y}^j|| \cdot ||\hat{\mathbf{v}}||}, \text{where } \hat{\mathbf{v}} = [1, 1, ..., 1]' \subseteq \mathbb{R}^{K+}. \tag{1.1}$$

The mixing index, as described in equation (1.1), captures the same intuition presented in Figure 1.1. It evaluates the multi-dimensional angular similarity between a skill vector $\mathbf{y}^j$ of dimension K and the multi-dimensional norm $\hat{\mathbf{v}}$. As different skills in $\mathbf{y}^j$ get closer to each other, the value of $Mix(\mathbf{y}^j)$ will rise accordingly. There are three key advantages of the skill mixing index defined using cosine similarity that are worth noting. First, it easily accommodates occupations represented by multi-dimensional skills. Second, this index is independent of the length of the skill vector, and focuses on the proximity between a skill vector and the norm, which indicates the degree of skill mixing. Lastly, this measure is inherently normalized, as the cosine of an angle in the first quadrant (indicating positive skills) lies in [0,1].

Data Construction: In analyzing the extent of skill mixing within occupations over time, I primarily use the Occupational Information Network (ONET). This dataset provides detailed information about the importance of skill requirements

---

[5] Cosine similarity together with other measures, such as Euclidean distance and Manhattan distance, have been used to calculate the similarity between vectors (i.e., Xia, Zhang, and Li 2015). An angle-based measure is by no means the only measure of skill mixing, though it has the clearest graphical illustration of the trade-off among skills. Online Appendix A.1.6 discusses two alternative skill mixing indexes: inverse Herfindahl–Hirschman Index (HHI) and normalized absolute distance.

for various occupations, offering an intensive measure of skill demand. To complement the insights from ONET, I also use data from online job postings via Lightcast. This dataset captures whether certain skills are required for the job postings, offering an extensive measure on skill demand for unfilled vacancies specifically. Below I discuss the details of data construction.

Developed by the North Carolina Department of Commerce and administered by the U.S. Department of Labor, O*NET is a successor to the Dictionary of Occupational Titles (DOT). It has become a primary resource for analyzing occupational skill requirements and work environments (i.e., see Acemoglu and Autor 2011; Yamaguchi 2012; Deming 2017). O*NET offers an comprehensive picture of occupations, covering approximately 270 descriptors categorized into nine modules.[6] While the earlier versions of O*NET include legacy ratings from DOT analysts, a shift occurred in 2003 when O*NET began sourcing responses from random samples of workers (job incumbents). To ensure consistent measurement, I choose descriptors from questionnaires updated based solely on these worker surveys.[7]

A key challenge when using O*NET comes from employing the longitudinal variation in occupation descriptors. Specifically, while each version of O*NET contains roughly 970 7-digit occupations, an average of 110 occupations undergo updates annually.[8] Such a pattern of updates could introduce selection bias when constructing measures of skill demand based on the descriptors. Contrasting

---

[6] For a comprehensive overview of O*NET, refer to online Appendix A.1.1, and for a discussion on the descriptors employed, see online AppendixA.1.2.

[7] Specifically, I use descriptors from the Work Context, Work Activities, Knowledge, and Skills questionnaires. After 2003, O*NET still contain responses from job analysts for questionnaires that have small sample sizes from workers. I abstract from those questionnaires in this paper.

[8] The decision of occupation updating is based on analysts' evaluations of factors such as the size of employment, the demand for labor, and alterations in the type of work involved. See Tippins and Hilton (2010) for more details.

prior research, which often explores worker reallocation across occupations using a single O*NET version, this paper's emphasis on the dynamics of skill demand requires the examination of these longitudinal changes.

To examine these longitudinal shifts in skill demand via O*NET data, I employ two approaches, following works such as Ross (2017) and Freeman, Ganguli, and Handel (2020). First, I focus on broader year intervals. For the time period from 2005 to 2018, I analyze the differences in skill requirements between the start and end of this period, during which most occupations are updated at least twice. To capture more granular time patterns, I use 4-year intervals, ensuring updates to cover over half of the occupations within these intervals. Given that each O*NET version retains data from prior years, I make a distinction between the release year and the represented year when integrating O*NET with other datasets. Online Appendix A.1.1 shows the specific O*NET versions used, their release dates, and the corresponding years.[9] Second, 274 7-digit occupations consistently receive updates between 2005, 2011, and 2018. While these occupations do not represent the entire economy, their trends under continual updates supplements the broader occupation analysis.

Furthermore, I utilize data from online job postings from Lightcast (previously "Burning Glass") for the years 2007 and from 2010 to 2017 that offers insights into unfilled vacancies. Lightcast is a labor market analytics firm that collects and analyzes millions of online job postings and provides detailed education requirements and thousands of codified skills extracted from the posting text. The key advantage of Lightcast data is that it provides comprehensive and up-to-date information on labor demand, and many recent studies have used this dataset

---

[9] Specifically, O*NET versions 13.0, 18.0, 22.0, and 25.0 were released in 2008, 2013, 2017, and 2022, respectively. These versions are interpreted as representing the years 2005, 2009, 2013, and 2018, respectively.

to analyze trends in job skill requirements (see, i.e., Deming and Kahn 2018; Hershbein and Kahn 2018; Braxton and Taska 2023). It is essential to recognize that while O*NET gauges the level and importance of a skill (intensive margin), Lightcast identifies whether a skill is required for a vacancy (extensive margin).[10] I employ Lightcast as an additional source to complement the picture of skill mixing changes over time.

Skill Measures: Leveraging the O*NET occupation descriptors, I first derive skill measures in line with Acemoglu and Autor (2011) to focus the analysis on the degree of skill mixing. These measures are widely applied and are easily comparable with other studies. To have a feasible dimension of skills to understand their mixing, I consolidate the two routine skills (routine cognitive and manual) into one, which I call routine skill, while I keep the non-routine skills (non-routine analytical and interpersonal) separate.[11] To capture the rise of computer technology post-2000, I also construct a computer skill measure based on two components related to programming and interacting with a computer. As these work activities are not easily codifiable, computer skill is also considered to be non-routine.[12] Appendix Table A.2 shows the detailed composing descriptors for each of the skill measures.

While these four skill measures including both routine and non-routine skills (hereafter RNR) serve as the core of this study's analysis, to provide a more comprehensive perspective on evolving skill demands, I also introduce two

---

[10] Several caveats of Lightcast data are that it may not capture jobs advertised through other channels, possibly over-represents certain sectors that tend to advertise online, and inherently might favor growing firms (Davis, Faberman, and Haltiwanger 2013).

[11] Since I only use descriptors updated by job incumbents in this study, I do not use non-routine manual skill since part of the composing descriptors comes from surveys of job analysts exclusively.

[12] For subsequent references to specific non-routine skills, I use terms like analytical, interpersonal, or computer skill, excluding the prefix ”non-routine.”

additional skills that have not been analyzed in previous studies and that are relatively non-routine—leadership and design. To enhance the reliability of these skill measures, especially for granular longitudinal time patterns, I apply principal component analysis (PCA) on the chosen descriptors following Guvenen et al. (2020) and Yamaguchi (2012). The final skill measures are linearly rescaled to lie in [0,1].[13] As a check of validity, online Appendix Table A.3 shows that my constructed skills correlate highly with other similar skill measures used in the literature. In online Appendix A.1.2, I build "broader" skill measures that each include more relevant descriptors than Acemoglu and Autor (2011), which are also highly correlated with the benchmark ones. Along with the discussion of my empirical results, I demonstrate their robustness to using alternative measures of skills and indexes of skill mixing.

Regarding the Lightcast data, I directly use the measures from Braxton and Taska (2023), which in turn are based on the methodology of Hershbein and Kahn (2018). Specifically, for the years 2007 and 2010-2017 of Lightcast data that this study uses, a vacancy is defined to use analytical skill if any of the codified job skills contain keywords such as "research", "analy", and "decision". Similarly, a vacancy is defined to require interpersonal skill if the codified job skills contain keywords such as "communication" or "teamwork".[14] Each occupation's skill measure is then determined by the proportion of vacancies demanding that

---

[13] Based on Definition 1, it is crucial that skill vectors are in the positive real space for an angle-based measure to be appropriate. In that regard, normalization by standard deviation will not work unless with additional re-normalization, and linear transformation to a positive interval appears most desirable as it also retains the cardinal information that is likely to be useful for an easily interpretable skill comparison (i.e., Autor and Handel 2013; Deming 2017; Lise and Postel-Vinay 2020). Alternative measures of skills and skill mixing are discussed in online Appendixes A.1.6 and A.1.7.

[14] More specifically, the keywords used to capture analytical skill are: ″research″, ″analy″, ″decision″, ″solving″, ″math″, ″statistic″, and ″thinking″. The keywords used to capture interpersonal skill are ″communication″, ″teamwork″, ″collaboration″, ″negotiation″, and ″presentation″. The key words used for computer skill are "computer", or any skill flagged as software by Lightcast.

specific skill, capturing the extensive margin of firm skill demand. To classify occupations within the Lightcast data, I used a 4-digit consistent census occupation code, as developed by Autor and Dorn (2013) to ensure matching with other datasets.

1.3.2 Aggregate Trends

This subsection investigates the trend of skill mixing across various occupations in the U.S. economy from 2005 to 2018, highlighting a notable shift towards increased mixing, particularly in non-routine skills. The analysis begins at the most granular level using 7-digit O*NET data, and also illustrate the key skills driving this overall trend. I then examine the time patterns of skill mixing changes at the census SOC level, employing both O*NET and Lightcast data. Throughout this discussion, I examine the robustness of these findings regarding to the choice of skill mixing indexes and skill measures.

Figure 1.2 depicts the density and median values of two different skill mixing indexes for the years 2005 and 2018, derived from 7-digit O*NET data. The first index incorporates the four routine and non-routine (RNR) skills, while the second focuses on the three non-routine skills (Non-routine). Panel (1) first reveals a modest rightward shift in the density of skill mixing index for RNR skills during this period. The Kolmogorov-Smirnov (KS) test confirms the distributional difference as statistically significant at the 1 percent level. However, there is more noticeable shift in density for only non-routine skills (analytical, computer, and interpersonal) as shown in Panel (2). By 2018, the density of the skill mixing indexes of these non-routine skills peaks at a much higher value than 2005, and the distribution's shift to the right is more pronounced, indicating a substantial

Figure 1.2: Density for Skill Mixing Indexes (Cosine Similarities), 2005 vs. 2018

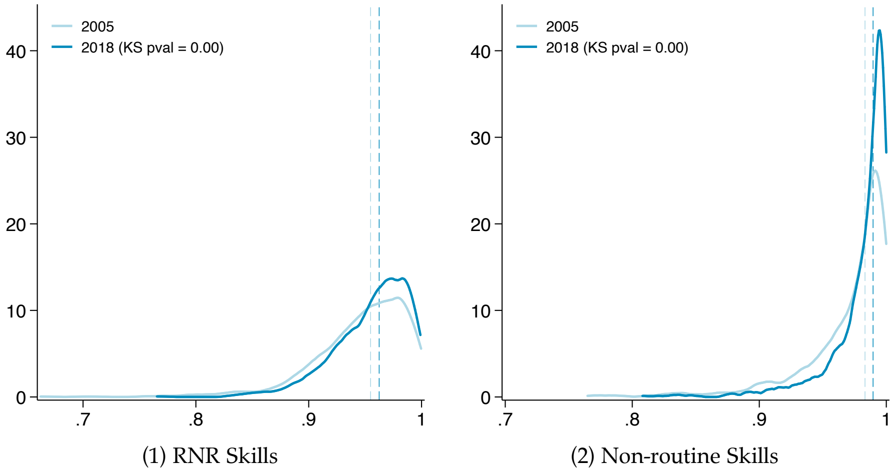

(1) RNR Skills

(2) Non-routine Skills

Notes: These figures plot the kernel density of different skill mixing indexes in 2005 (light blue line) and 2018 (dark blue line). The x-axis displays the value of skill mixing indexes with a maximum of 1 by construction. "RNR" stands for the one routine and three non-routine skills (analytical, interpersonal, computer). "Non-routine" skills only include the three non-routine skills. Specific composing descriptors of the skills are in online Appendix Table A.1. These plots are created using O*NET at 7-digit occupations unweighted by employment.

growth in occupations demanding a high level of mixing of non-routine skills.[15]

This pattern of growth in skill mixing is not unique to the choice of non-routine skills and becomes even sharper when accounting for the composition of labor force across occupations. In online Appendix Figure A.2 Panel (1), I show that the rightward shift in the density of the mixing index remains consistent when including other non-routine skills (leadership and design). I also combine O*NET data with detailed employment weights from the Occupational Employment and Wage Statistics (OEWS) in online Figure Appendix A.2 Panel (2).[16] The rightward shift of all the skill mixing indexes becomes more pronounced when

---

[15] In addition to index-based evaluation of skill mixing, one can also non-parametrically examine occupation skill requirements in two-dimensional spaces. Online Appendix A.1.3 discusses and presents non-parametric plots for six skill pairs from both 2005 and 2018, confirming the observed increase in skill mixing, particularly for non-routine skills.

[16] The OEWS uses 6-digit SOC codes, while O*NET uses 7-digit occupation codes that are based on 6-digit SOC. I match OEWS with O*NET at a 6-digit SOC level and distribute the employment weight evenly for 7-digit O*NET occupations within a 6-digit occupation.

weighted by employment shares.[17]

Two exercises further highlights the *drivers of skill mixing* and the heterogeneity across different occupations. In online Appendix Figure A.2, Panel (3), I exclude each skill individually from the three non-routine skills to compute the skill mixing index. Regardless of which skill is left out, a noticeable rightward shift in density persists, suggesting that each skill has contributed to the overall increase. In Online Appendix Table A.4, I conduct a decomposition of changes in the mixing of non-routine skills using polynomial regressions.[18] The results reveal that computer skill is the most significant driver across all occupations. However, for higher-paid occupations (i.e., professionals, managerial, white-collar), interpersonal skill is more important, whereas for medium to lower-paid jobs (i.e., blue-collar, service), computer skill remains dominant.

Time Pattern: To examine the time profile of the shifts in skill mixing and understand the sources of variation, I combine the longitudinal variation in skill mixing from O*NET with worker employment and characteristics from the ACS using consistent census occupation codes from Autor and Dorn (2013).

I take three additional steps to illustrate the time pattern. First, I construct the trend at 4-year intervals so that more than half of the occupations (about 60 percent of employment) are updated between observations. Second, a limitation

---

[17] This result implies that occupations with larger employment shares have a more significant rise in skill mixing.

[18] Since the mixing index based on cosine similarity is a nonlinear function of the composing skills, a standard variance-decomposition can hardly partial out the variations clearly. Instead, I conduct a polynomial regressionof non-routine skills' mixing index and each composing skill's polynomials up to order N over the period from 2000 to 2020: $\text{Mix}(\mathbf{y})^{\text{percentile}}_{ijt} = \beta_1 y^1_{ijt} + \beta_2 y^2_{ijt} + ... + \beta_N y^N_{ijt}$, where $\text{Mix}(\mathbf{y})^{\text{percentile}}_{ijt}$ indicates the percentile rank of an individual's i mixing index of non-routine skills in occupation j at time t, and y_{ijt} is the measure of a specific composing skill for that individual's occupation at time t. The R-Squared values are then used to evaluate the degree to which each composing skill explains the variance in skill mixing.

of the skill mixing index is that it has a long left tail and concentrates at values near 1; moreover, it cannot be comfortably treated as cardinal. To address this concern, I transform skill mixing indexes into percentile values based on their rank in the 2005 distributions. Third, I weight each skill requirement from O*NET by the total hours of work in each sex-education-industry-occupation cell in the ACS to implicitly control for changes in task inputs due to variations in gender, education, industry, and occupation compositions in the U.S. economy (see Autor, Levy, and Murnane (2003) and Deming (2017) for other examples).

Table 1.1 presents the 3-digit Census SOC occupations that have the highest increase in skill mixing from 2005 to 2018, in percentile units. It also details the skill compositions for those occupations, as well as their broader 2-digit occupational categories (in parentheses). To ease the presentation, the table includes only those occupations that constitute a minimum of 0.2% of overall employment, though all occupation codes have been used in the analysis.

The occupations with the highest increase in non-routine skill mixing include service sector, particularly in housekeeping and cleaning, sales roles, and blue-collar jobs involving operation, production, and labor tasks. In these jobs, analytical and computer skills have become increasingly important. For example, in 2005, the role of a housekeeper heavily relies on interpersonal skills; by 2018, the importance of computer skills equals three-quarters of that of interpersonal skills, and analytical skills become just as crucial. Similarly, in the case of sales clerks, while interpersonal and computer skills are predominant in 2005, by 2018 analytical skills equal the importance of computer skills and surpass interpersonal skills. In contrast, blue-collar occupations such as packers, which primarily required routine skills in 2005, have seen a marked increase in all non-routine

Figure 1.3: Trend of Skill Mixing in the US Economy, 2005-2018

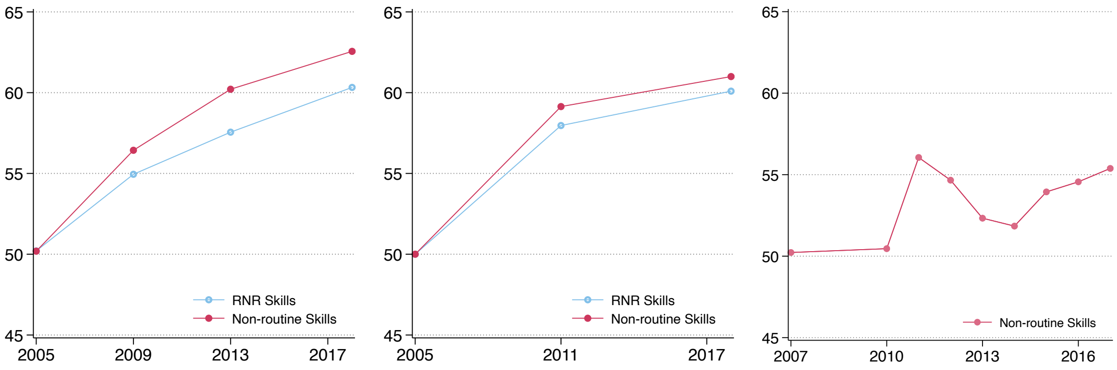

(1) Full O*NET (2) Constant Updates (3) Lightcast

Notes: These figures plot the employment-weighted skill mixing indexes in the U.S. economy from 2005-2018. The y-axis is the percentile of skill indexes in year 2005. By construction, each index has a mean of 50 percentiles in 2005; succeeding points are employment-weighted means mapped to its percentiles in 2005. Panel (1) and (2) combine O*NET and ACS data with consistent 4-digit occupation codes from Autor and Price (2013) and developed by Deming (2017). The matching of different O*NET releases and ACS years are detailed in online Appendix Table A.1. Panel (1) show the trend for the universe of occupations while Panel (2) only include 274 7-digit occupations that are constantly updated between 2005, 2011, and 2018. Panel (3) combines Lightcast job posting data and the ACS with the same occupation coding. Employment weights from ACS are the total hours of work aggregated to sex-education-industry-occupation cells.

skills, resulting in a significant rise in their RNR skill mixing index.

Figure 1.3 demonstrates that the degree of skill mixing has risen substantially and steadily between 2005 and 2018. By construction, each index has a mean of 50 percentiles in 2005; succeeding points are employment-weighted means of each index mapped to its percentiles in 2005. By 2018, the degree of mixing in non-routine skills for an average occupation in the US economy is 12.4 percentiles higher than its 2005 level. The degree of skill mixing in RNR skills has also increased steadily to a slightly lesser extent, averaging 10.1 percentiles higher.[19]

---

[19] The inclusion of routine skill decreases the magnitude of the rise in skill mixing implies that the speed of mixing of routine with other skills is slower than the speed of mixing among non-routine skills. Online Appendix Figure A.6 depicts the trend of skill mixing for specific skill pairs. The findings reveal a modest increase in the mixing of routine with computer skills at 2.9 percentiles from 2005 to 2018. Conversely, the degree of mixing between routine and other non-routine skills has remained stable.

A potential concern of using O*NET data to obtain the longitudinal variation of skill demand is that the trend could be affected by the inconsistency in occupation updating. In Panel (2) of Figure 1.3, I compute these trends focusing solely on the 274 7-digit occupations that are constantly updated between 2005, 2011, and 2018, thus reflecting a consistent updating of skill requirements among these occupations. The same qualitative pattern holds, that is, there has been a sharp increase in the degree of skill mixing, particularly of non-routine skills between 2005 and 2018. Nonetheless, for the constantly updated occupations, the shift is mostly pronounced before 2011.[20]

In Panel (3) of Figure 1.3, I complement the picture of changing degree of skill mixing using the Lightcast data through a similar pairing with O*NET data starting in 2007, the first year when the company starts to collect job postings. Overall, firms are more likely to post job requirements that contain more mixed-skill demands. By 2017, the degree of skill mixing in online posted vacancies averaged 5.2 percentiles higher compared to 2007. The time pattern of skill mixing among online job postings appears to be more volatile, first peaking in 2011, then sliding down until 2014, before rising dramatically afterwards. Despite the greater variance, the same qualitative pattern holds that occupations have a higher demand for the mixing of non-routine skills.[21]

Robustness of the trend: One may be concerned about that the overall pat-

---

[20] In online Appendix Figure A.3, I show employment percentages and hourly wages across various job categories in the full sample and the sample for constantly updated occupations. The hourly wage rates across the categories are fairly consistent between the full and selected samples, with minor discrepancies: the selected sample has less presence of professionals and sales occupations.

[21] The higher degree of volatility is partly driven by the nature of the measure and the data. The measures of skills from job postings are whether firms require a particular skill in the text of job ads, which are naturally noisier than the questions on level and importance from O*NET. Moreover, firm job posting is more influenced by firm entry and exit patterns.

terns shown so far are driven by the choice of skill measures or the choice of skill mixing index. To address this concern, in online Appendix A.1.6 and A.1.7, I demonstrate the robustness of these trends across various skill measures, alternative skill mixing indexes, as well as skill mixing indexes of distinct skill pairs. For example, using standardized (or broader) measures of skills, the increase in the degree of mixing of non-routine skills is 6 (or 13) percentiles from 2005 to 2018; using inverse Herfindahl-Hirschman Index, the increase in the mixing indexes of any given skill pair is above 10 percentiles during the same period. Across these checks, the qualitative picture remains consistent: there has been a notable rise in the degree of skill mixing, particularly for non-routine skills.

1.3.3 Decomposing the Sources

To gain a deeper understanding of the variations underlying changes in skill mixing, I undertake three exercises. First, I decompose the longitudinal changes in skill mixing in the U.S. economy, differentiating between intensive margin skill mixing index changes and extensive margin employment shifts across occupations. This analysis reveals that within-occupation skill mixing shifts play a more influential role in driving skill mixing than across-occupation employment shifts. Second, I perform a regression analysis that include extensive controls such as various skill supply measures, as well as gender, industry, and occupation fixed effects. I find that the pronounced trend of increasing skill mixing persists.

Table 1.2 shows a shift-share decomposition of the changes in the employment-weighted skill mixing indexes into within-occupation index shifts and across-occupation employment changes, at both 7-digit O*NET occupation and 4-digit census occupation levels using employment weights from the OEWS and ACS respectively. I conduct the analysis both for the full O*NET data and the subset of persistently updated occupations, alongside the Lightcast data. Irrespective of the dataset or skill groupings, within-occupation variation predominantly drives the rise in skill mixing. For example, for the 12.4 percentile increase in the mixing of non-routine skills in the full O*NET data at 4-digit occupation level, within-occupation increase contribute 9.7 percentiles while only 2.7 percentiles stem from worker reallocation; for the 5.2 percentile increase in the mixing of non-routine skills in Lightcast data, within-occupation increases account for 4.4 percentiles.[22] Interestingly, for the constantly updated occupations at 7 digits, worker reallocation actually contributes negatively to the increase in skill mixing. This pattern implies that for these granular occupations under regular updates, the contribution of within-occupation variation more than accounts the increase in skill mixing. At 4-digit occupations, worker reallocation does contribute positively to these increase in the mixing of non-routine skills, but the influence is still marginal compared to within-occupation variation; for RNR skills, the contribution remains negative.

An alternative explanation of the employer-side shifts in accounting for skill mixing could be that even within occupations, the supply of labor might have changed due to for example, rising human capital or labor force participation of female workers. To further shed light on the sources, Table 1.3 shows results from a regression of skill mixing indexes on a linear time trend (year indicator) across combinations of O*NET and Lightcast with the ACS data. I further control for the interaction between gender and education fixed effects, and between industry and occupation fixed effects; additionally and I include flexible polynomials and interactions of years of education and experience. The table shows a universal

---

[22] Online Appendix A.5 shows the decomposition results using skill mixing indexes for different skill pairs and a similar result holds.

increase in the degree of skill mixing at a magnitude of 0.65 to 0.75 percentiles per year using O*NET data and 0.33 percentiles per year using Lightcast data. This increase persists within gender, education, industry, and occupation groups and is unaffected by controls of worker's labor supply. Despite the varying trends among different industry and occupation groups, these did not alter the overall increase in skill mixing. This finding suggests that the increase in skill mixing withstands adjustments for worker composition, highlighting the important impact of demand-side forces.

Lastly, I investigate the relationship between changes in skill mixing and two industrial-level shifts: the increase in IT capital and the adoption of industrial robots. I obtain productive capital stock for "Total information processing equipment" from the Bureau of Labor Statistics Total Multifactor Productivity tables. Additionally, following Acemoglu and Restrepo (2020), I use data on the stock of robots from the International Federation of Robotics (IFR).[23] Online Appendix Table A.6 presents regressions of skill mixing indexes from 2005-2015 on IT capital levels and robot adoption changes, accounting for worker composition across gender, education, and industry groups, as well as flexibly controlling for years of experience and levels of education. The results indicate that a rise in productive IT capital stock by 10 billion is associated with a 0.1 percentile increase in non-routine skill mixing and a 0.1 percentile decrease in RNR skill mixing, driven by rising demand for computer skills and a marginal decline in the importance of interpersonal skills.[24] On the other hand, an increase of one

---

[23] Specifically, I focus on the average number of industrial robots per thousand workers in five European countries to isolate the impact of global technological advancements. These 5 countries are Denmark, Finland, France, Italy, and Sweden. Germany is omitted due to its growth in robotics way above other countries.

[24] The average IT capital stock across industries in year 2005 is 70 billion. The results on the association between IT capital and skill mixing is significant at 5 percent using O*NET data and not precisely estimated using Lightcast data.

industrial robot per thousand workers has no significant association with the mixing of non-routine skills but is associated with a 1.24 percentile decrease in mixing RNR skills.

1.3.4 Variation in Skill Mixing Changes by Subgroups

Beneath these general trends of skill mixing are diverse patterns among occupations and worker groups. Panel (1) of Figure 1.4 illustrates the changes of skill mixing indexes from 2005 to 2018 across four primary occupation categories, grouped by wage levels and encompassing all non-agricultural employment in the U.S.[25] The units of changes are in percentiles of the skill mixing indexes' 2005 distributions, similar to Figure 1.3.

Considering the occupational trends, for the mixing of non-routine skills, service and white-collar occupations see higher increases than other occupations. For the mixing of RNR skills, on the other hand, blue-collar followed by service jobs show most pronounced rise. In contrast, high-wage occupations show the least increase in skill mixing for both skill groups. This pattern highlights that the bulk of skill mixing occurs in medium- to lower-wage professions, especially within the service sector.

Panel (2) and (3) of Figure 1.4 depict changes in skill mixing for RNR skills and non-routine skills, respectively, across gender and education groups. For the mixing of RNR skills, workers with less than a college degree show a greater

---

[25] The categorization into four groups is based on Acemoglu and Autor (2011), which is derived from 10 1-digit occupational groups that cover the entirety of US non-agricultural employment. Specifically, ”High-wage” includes Managers, Professionals, and Technicians; ”White-collar” comprises Office/Administrative and Sales roles; ”Blue-collar” includes Production, as well as Operators/Laborers; and ”Service” consists of Protective Services, Food/Cleaning Service, and Personal Care occupations.

Figure 1.4: Skill Mixing Index Change by Occupation and Gender Groups, 2005-2018

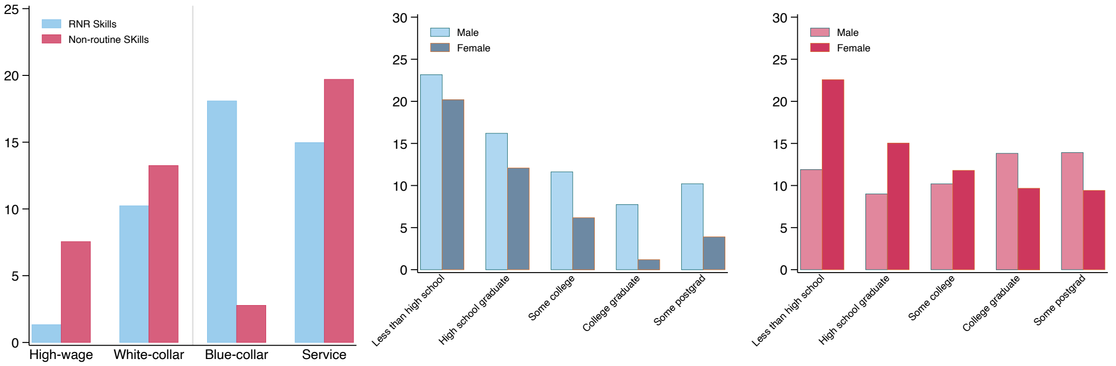

(1) Occupation groups

(2) Gender and education groups – RNR skills

(3) Gender and education groups – Non-routine skills

Notes: This figure plots the changes in skill mixing indexes across different occupation groups and gender-education combinations. The units of the index changes are percentiles of their distributions in 2005. The occupation groups (High-wage, White-collar, Blue-collar, Service) follow Acemoglu and Autor (2011). The education grouping aligns with the educational categories defined in the ACS. O*NET and ACS data are combined for these figures with consistent occupation codes from Autor and Price (2013) and developed by Deming (2017).

increase than other groups. However, gender difference appears; on average, male workers experience 5 percentiles higher skill mixing than their female counterparts, and the difference is particularly evident among highly educated workers. Regarding the mixing of non-routine skills, male workers across different education levels experience roughly similar increases in skill mixing, whereas female workers observe a slower increase in mixing as education level rises, less than their male counterparts for those with a college education or higher.

In online Appendix A.1.4, I show the decomposition of skill mixing by industries, and a similar pattern holds. The private service sector followed by retail trade and construction lead others in the growth of skill mixing, while public, social and professional services sectors demonstrate only modest increases, particularly for RNR skills.[26] I also show the decomposition of occupations'

---

[26] The sectors that have the least growth in skill mixing are public and education and social

Figure 1.5: Smoothed Employment and Wage Changes by Skill Percentile, 2005-2018

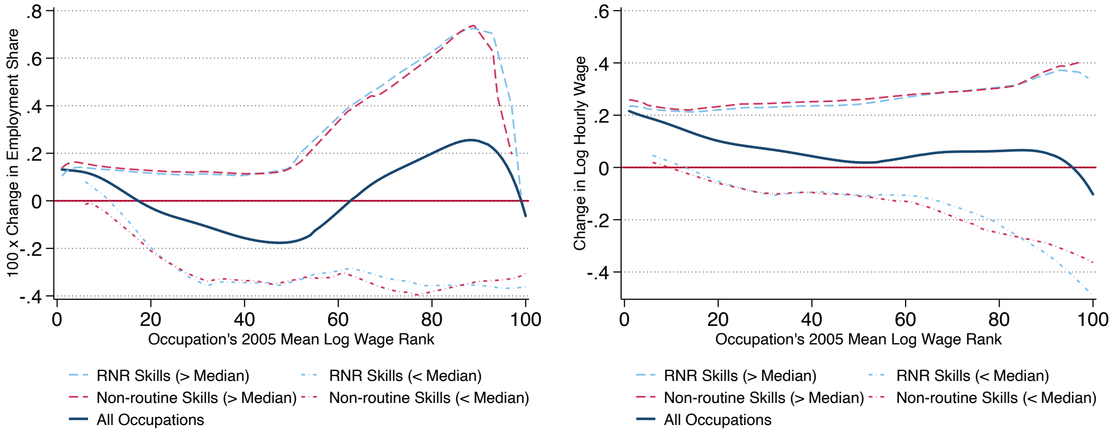

Notes: These figures plot the smoothed changes of share of total hours workded (Panel A) and hourly wage (Panel B) for occupations between 2005-2018. On the x-axis, occupations are ranked into 100 percentiles by the average log wages of workers in those occupations in 2005. The changes in the share of hours worked and percent wage growth are then calculated for each percentile, which fit into smoothed lines using cubit polynomial fit. Solid lines depict the smoothed employment/wage changes for all the occupations, while dashed (or dotted) lines depict the changes for occupations with above-median (or below-median) increases in the skill mixing indexes.

changing mixing of distinct skill pairs, which confirms that non-routine skills drive skill mixing in all occupations, while routine skill is only more mixed with other skills in blue-collar and service occupations.

Distributional Implications: One of the key structural changes in the U.S. labor market post-1980 is the pronounced job polarization or hollowing out of middle-skill employment and wage growth, due potentially to the routine-biased technological change and offshoring (Acemoglu and Autor 2011; Goos, Manning, and Salomons 2014). To see how much skill mixing can relate to these distributional dynamics, Figure 1.5 depicts the smoothed observed changes in both the share of total hours worked and log wage in 2005-2018 for occupations ranked

---

services. This result is consistent with Hershbein and Kahn (2018) that industries with locally consumed goods are more likely to change skill demand.

by their hourly wage percentiles in 2005. I reconstruct these smoothed employment/wage changes for two groups of occupations: those with above-median increases in skill mixing indexes and those below the median.

Figure 1.5 first confirms the inverted bell shape (polarization) of observed employment and to a lesser extent, the change of wages. Furthermore, it illustrates key differences for occupations that have become more skill mixed. For occupations within similar wage ranks in 2005, those that become more mixed in skills have a higher increase in employment share and wage growth. In fact, almost the entirety of employment and wage growth is accounted for by occupations that have become more skill mixed during this period. Therefore, relating to polarization, the differential growth in employment and wage among occupations at the top and bottom end of the 2005 wage distribution are entirely accounted for by skill-mixing occupations during this period. Besides being an important phenomenon for labor market dynamics, skill mixing also provides a unified and multi-dimensional perspective of the polarization changes.

1.4 Returns to Skill Mixing

In order to better understand the influence of skill mixing on workers' labor market outcomes, this section examines the wage returns associated with skill mixing in relation to occupational choices and inherent worker skills. Additionally, I discuss the return on investment for a college major with a more mixed skill set.

1.4.1 Data and Measurement

To assess wage returns associated with skill mixing, I use the National Longitudinal Survey of Youth (NLSY) datasets from both the 1979 and 1997 cohorts, which offer comprehensive records of the participant's employment and educational histories. I combine these two cohorts to increase the sample size, limiting to the period from 2005 to 2019 to align with the timing of my skill mixing measurements from O*NET as discussed in the previous section.[27] The NLSY data are connected with O*NET via the census occupation information in NLSY and the crosswalk formulated by Autor and Dorn (2013). My principal focus is the real log hourly wage, adjusted to 2013 dollars. As in Altonji, Bharadwaj, and Lange (2012), I trim values of the real hourly wage below 3 or above 200. The results of wage returns are robust to considering alternative sample constructions, such as excluding respondents over the age of 55 or using the unprocessed real hourly wage.

The key advantage of NLSY is that it is a worker-level panel, and also contains information on workers' pre-market abilities. This allows for the control of worker characteristics in assessing occupational wage returns to skill mixing and also facilitates the evaluation of returns to the worker-level degree of skill mixing. The selected measures of worker abilities are chosen to align well with the skill measures in O*NET: the Armed Forces Qualifying Test (AFQT) scores represent analytical skill, the social skills measure developed by Deming (2017) is employed to represent interpersonal skill,[28] and routine skill is measured by

---

[27] The NLSY 1979 and NLSY 1997 are nationally representative surveys of youth, capturing data from individuals aged 14 to 22 in 1979 and 12 to 16 in 1997, respectively. During my sample period, the median age is 37, and 91 percent of the sample is below 50.

[28] I use the AFQT scores constructed by Altonji, Bharadwaj, and Lange (2012) that are consistent across NLSY waves and account for age-at-test, test format, and other peculiarities. For interper-

the workers' Armed Services Vocational Aptitude Battery (ASVAB) mechanical orientation scores.[29] As NLSY offers scant information on workers' computer skills, I adopt the worker's occupation or college major's computer skill value in the year 2005 as a proxy for the worker's initial endowment of computer skill. Online Appendix Table B.1 lists the corresponding measures.

1.4.2 Wage Returns

To estimate the returns to skill mixing, I regress the log wage of workers on the levels of different skills required by their employed occupations, as well as the skill mixing indexes of these skills. Conditional on skill levels, the coefficients on skill mixing indexes identify the returns to working in occupations that are more mixed among the skills. To further examine the worker level returns to skill mixing, I add to the right-hand side the levels of the skills that workers have, and their degrees of mixing. The coefficients on worker-level mixing indexes then identify the wage premium to the mixing of worker skills conditional on occupational skill requirements. To streamline the discussion, I focus on the degree of mixing of the three non-routine skills (analytical, computer, interpersonal), given its significant increase during the observation period as in Section 1.3. Analysis of returns to mixing between routine and other skills as well as to individual skills is deferred to the online Appendix A.1.8.

Throughout all the specifications, I include ethnicity by gender, age,

---

sonal skills, I use the social skill measure developed by Deming (2017) assessing extraversion, which is constructed based on sociability in childhood and adulthood in NLSY79, and two questions from the Big 5 inventory in NLSY97 respectively.

[29] ASVAB test scores are only available for the NLSY79 survey. For NLSY97, I impute their ASVAB scores using a regression model with indicators for gender and ethnicity, and fixed effects that include age, year, census division, metropolitan area, and urbanicity.

metropolitan status, individual year, years of education, census region, and urbanicity fixed effects. I also include occupation fixed effect to control for time-invariant differences across occupations, which allows me to focus on how the changes in skill requirements within occupations are affecting wage returns, consistent with the empirical finding that this margin is the main driver of skill mixing. Standard errors are clustered at the individual worker level to account for within-group correlation and heteroskedasticity among repeated observations at the individual level.

Occupation and Worker Level Returns: Table 1.4 shows the wage returns to skill mixing at both the occupation and individual levels, indicating a positive premium for mixing non-routine skills. Column (1) reveals that workers in occupations that become one standard deviation more mixed among analytical, computer, and interpersonal skills earn a wage premium of 1.7 percent per year, significant at the 1 percent level. In column (2), I incorporate the mixing index of worker abilities, which enhances the precision of occupation-level wage premiums and estimates the return to skill mixing at the worker level. The results suggest that on the worker side, workers who are a standard deviation more mixed among the non-routine skills earn a wage premium of 6.5 percent. Meanwhile, the wage premium for the three non-routine skills remains at 1.5 percent per year at the occupational level. In column (3), I further restrict the analysis to within-worker variation by adding worker fixed effects; combined with the occupation fixed effects, this specification closely aligns with an AKM model.[30] The magnitude of the returns to skill mixing presented in column (3) is similar to

---

[30] Using within worker variation to study wage growth has been discussed and applied in i.e., Neal (1999); Gibbons et al. (2005); Lazear (2009) and Deming (2017). Choné and Kramarz (2021) found that under a worker assignment model with bundled skills, the implied wage equation also has an AKM form.

that in column (2). Workers in occupations that become one standard deviation more mixed among non-routine skills experience a wage increase of 1.4 percent.

Given the positive wage premium associated with mixing non-routine skills at both occupation and worker levels, below I discuss some additional returns to skill mixing. I first examine the robustness of the wage return results. Next, I explore the returns of skill mixing in for employment and college major choices.

Discussions:

Robustness: To gain a more detailed view of the drivers of the positive premium of skill mixing, online Appendix Table A.8 uses mixing indexes of skill pairs instead of a high-dimensional mixing index for non-routine skills, which indicates the positive wage returns primarily arise from the mixing of analytical with computer and analytical with interpersonal skills. Robustness checks in online Appendix A.9 show that the occupational returns to skill mixing conditional on worker fixed effects are robust to alternative skill measures and indexes of mixing. Specifically, the results consistently suggest a wage premium of 1 to 2.5 percent in occupations that mix these skills. Further, in online Appendix Table A.8 I show that there is also a positive employment premium for workers with a more mixed skill set: workers with a more mixed level of all the skill pairs except for routine and interpersonal are also more likely to exit unemployment.[31]

Additional Returns: Moreover, using the college education information from NLSY, I assess the returns of pursuing college majors with different degrees

---

[31] Throughout my analysis, I classify a worker as employed if the worker earns a wage greater than zero and has held one or more jobs since the last NLSY interview, consistent with Altonji, Bharadwaj, and Lange (2012) and Deming (2017). Further, workers without a paying job for 24 months are considered to be out of the labor force.

of skill mixing.[32] I calculate for workers who studied a particular major, the employment weighted average of skill intensities of their occupations in O*NET to compute skill mixing index for that major.[33] Online Appendix Table A.10 highlights majors based on their skill mixing levels and changes. Notably, Architecture and Environmental Design stands out in mixing the three non-routine skills, with Computer and Information Sciences, and Communications following closely. Additionally, Social Sciences and Agriculture and Natural Resources are among the top majors in mixing routine and non-routine skills. Online Appendix Table A.8 column (4) quantifies workers' human capital by their majors' skill contents, and shows a positive return of around 3 percent studying a college major that is associated with a standard deviation higher mixing of non-routine skills.

---

[32] There are some inconsistencies in NLSY's field of study coding: NLSY79 uses its own major codes that contain 25 two-digit categories, while NLSY97 uses another set codes for years leading to 2010 and transfers to National Center for Education Statistics (NCES)'s 2010 College Course Map (CM10) for years after 2010. For consistency, I map the two different types of major codes in NLSY97 to the 25 two-digit major categories in NLSY79. Online Appendix Table A.11 shows the crosswalk of different types of major field of study codes.

[33] I take the first field within a year as representing a worker's major field in the case of multiple fields.

Table 1.1: Top Occupations in Skill Mixing Growth

Top Occupations	Year	Analytical	Computer	Inter-personal	Routine	Mixing Index	Percentile
Mix of Non-routine Skills							
Packers, fillers, and wrappers	2005	0.58	0.44	0.16		0.915	1
(Operators/Fabricators/Laborers)	2018	0.52	0.40	0.42		0.994	99
Housekeepers, maids, cleaners	2005	0.00	0.10	0.24		0.753	0
(Personal Care and Services)	2018	0.28	0.20	0.25		0.990	96
Sales counter clerks	2005	0.13	0.32	0.30		0.946	7
(Sales)	2018	0.50	0.52	0.39		0.993	99
Recreation facility attendants	2005	0.24	0.18	0.39		0.947	7
(Personal Care and Services)	2018	0.38	0.40	0.35		0.998	99
Janitors	2005	0.10	0.07	0.21		0.913	1
(Food Prep/Buildings and Grounds)	2018	0.15	0.16	0.21		0.987	93
Carpenters	2005	0.50	0.14	0.44		0.915	1
(Production/Craft/Repair)	2018	0.59	0.38	0.53		0.985	90
Cashiers	2005	0.08	0.41	0.33		0.892	0
(Sales)	2018	0.31	0.41	0.49		0.984	87
Packers and packagers by hand	2005	0.16	0.16	0.30		0.951	12
(Operators/Fabricators/Laborers)	2018	0.49	0.40	0.54		0.992	99
Data entry keyers	2005	0.56	0.77	0.27		0.935	3
(Office/Admin)	2018	0.55	0.66	0.43		0.985	90
Sales supervisors and proprietors	2005	0.40	0.39	0.79		0.943	6
(Sales)	2018	0.49	0.57	0.74		0.985	92
Mix of RNR Skills							
Packers and packagers by hand	2005	0.16	0.16	0.30	0.71	0.824	0
(Operators/Fabricators/Laborers)	2018	0.49	0.40	0.54	0.70	0.979	99
Cashiers	2005	0.08	0.41	0.33	0.71	0.863	2
(Sales)	2018	0.31	0.41	0.49	0.61	0.973	99
Assemblers of electrical equipment	2005	0.35	0.25	0.34	0.82	0.894	5
(Operators/Fabricators/Laborers)	2018	0.44	0.43	0.40	0.65	0.979	99
Equipment cleaners	2005	0.23	0.24	0.26	0.63	0.896	5
(Operators/Fabricators/Laborers)	2018	0.41	0.32	0.52	0.54	0.981	99
Cooks	2005	0.24	0.16	0.34	0.59	0.899	6
(Food Prep/Buildings and Grounds)	2018	0.46	0.33	0.46	0.64	0.974	99
Painters, construction, maintenance	2005	0.29	0.12	0.28	0.56	0.892	5
(Production/Craft/Repair)	2018	0.53	0.30	0.56	0.72	0.962	94
Hairdressers and cosmetologists	2005	0.52	0.10	0.38	0.35	0.912	11
(Personal Care and Services)	2018	0.51	0.30	0.44	0.46	0.985	99
Accounting and auditing clerks	2005	0.26	0.72	0.33	0.28	0.905	7
(Office/Admin)	2018	0.40	0.69	0.33	0.44	0.960	93
Packers, fillers, and wrappers	2005	0.58	0.44	0.16	0.84	0.900	6
(Operators/Fabricators/Laborers)	2018	0.52	0.40	0.42	0.82	0.954	90
Punching and stamping operatives	2005	0.25	0.29	0.15	0.86	0.809	0
(Operators/Fabricators/Laborers)	2018	0.42	0.37	0.35	0.74	0.948	81

Notes: This table presents specific O*NET occupations at census SOC levels that have the greatest growth in skill mixing from 2005 to 2018. It provides details on compositions of skills within these occupations and the corresponding changes in skill mixing indexes. The last column translates skill mixing levels into percentiles relative to their 2005 distributions.

Table 1.2: Shift-Share Decomposition of Skill Mixing Index Changes

	Skill Groups	7-digit Occupations			3-digit Occupations		
		total	within	across	total	within	across
Full O*NET	RNR Skills	6.78	4.93	1.85	10.12	9.46	0.66
	Non-routine Skills	9.21	5.62	3.59	12.37	9.72	2.65
Constant Updates	RNR Skills	5.59	6.73	-1.14	10.09	10.74	-0.65
	Non-routine Skills	4.05	5.33	-1.29	11.00	9.69	1.31
Lightcast	Non-routine Skills				5.16	4.37	0.78

Notes: This table shows a shift-share decomposition of changes in the average level of different mixing indexes between 2005-2018 in percentile units. Specifically, for a change in the percentile of a mixing index over two periods t and τ, its change $\Delta T_\tau = T_\tau - T_t$ which can be decomposed to $\Delta T = \sum_j \left(\Delta E_{j\tau}\alpha_j\right) + \sum_j \left(E_j \Delta\alpha_{j\tau}\right) = \Delta T^a + \Delta T^w$ where $E_{j\tau}$ is employment weight in occupation j in year τ, and $\alpha_{j\tau}$ is the level of mixing index h in occupation j in year τ, $E_j = \frac{1}{2}(E_{jt} + E_{j\tau})$ and $\alpha_j = \frac{1}{2}(\alpha_{jt} + \alpha_{j\tau})$. ΔT^a and ΔT^w then represent across-occupation and within-occupation change.

Table 1.3: Annual Changes in Skill Mixing Indexes (in Percentiles)

	RNR Skills		Non-routine Skills	
	(1)	(2)	(3)	(4)
*A. Full O*NET, 2005-2018*				
Year indicator	0.77^{***}	0.70^{***}	0.81^{***}	0.71^{***}
	[0.14]	[0.07]	[0.08]	[0.06]
Observations	237,885	237,885	237,885	237,885
R-squared	0.10	0.83	0.08	0.83
*B. O*NET Constant Updates, 2005-2018*				
Year indicator	0.77^{***}	0.75^{***}	0.68^{***}	0.65^{***}
	[0.12]	[0.11]	[0.11]	[0.11]
Observations	107,956	107,956	107,956	107,956
R-squared	0.29	0.81	0.15	0.82
C. Lightcast, 2007-2017				
Year indicator	—		0.42^{**}	0.33^{**}
	—		[0.19]	[0.15]
Observations	—		532,636	532,636
R-squared	—		0.25	0.87
Experience and edu controls	X	X	X	X
Gender × education FE	X	X	X	X
Industry × occupation FE		X		X

Notes: This table provides regression results on the relationship between the percentile values of RNR skills and Non-routine skills, based on their distributions in the year 2005, and a time trend variable (year values). The analysis incorporates data from the full O*NET, constantly updated O*NET, and Lightcast datasets combined with ACS. See online Appendix A.1.1 and A.1.6 for the data construction. The regressions include controls for gender-education fixed effects, industry-occupation fixed effects, polynomials of years of work experience up to power 4, and the interaction of experience polynomials and education fixed effects and gender. Education fixed effects include 5 categories (no high-school, high-school graduate, some college, college graduate, post-college). *** p¡0.01, ** p¡0.05, and * p¡0.1.

Table 1.4: Return to Skill Mixing: Occupations and Workers

Dependent: ln (hourly wage)	(1)	(2)	(3)	(4)
Mix (non-routine skills): Occupation	0.017^{***}	0.015^{***}	0.001	0.014^{***}
	[0.005]	[0.005]	[0.006]	[0.005]
Mix (non-routine skills): Worker		0.065^{***}	0.070^{***}	
		[0.017]	[0.017]	
Interaction			0.032^{***}	
			[0.008]	
Ethnicity, gender, age/year, region, edu FE	X	X	X	X
Occupation FE	X	X	X	X
Worker FE				X
Observations	88, 391	79, 343	79, 343	88, 391
R-squared	0.41	0.43	0.43	0.76

Notes: This table reports the result of estimating wage equations using pooled NLSY79&97 data for employed workers from 2005-2019. Log hourly wages are the outcome variables and person-year is the unit of observation. The occupational skill and skill mixing measures come directly from O*NET and are merged to NLSY79&97 based on census occupation codes. The worker-level skill measures are constructed to correspond to occupation-level measures as in Table B.1 and skill mixing indexes are then calculated accordingly. All measures of skill and skill mixing are normalized to have mean 0 and standard deviation 1. Ethnicity-by-gender, age, year, census region, urbanicity, and a 5-category (no high-school, high-school graduate, some college, college graduate, post-college) education fixed effects are included for all regressions, with additional fixed effects as indicated in the table. Standard errors are clustered at the individual level. *** p¡0.01, ** p¡0.05, * p¡0.10.

CHAPTER 2

A MULTI-DIMENSIONAL SKILL DIRECTED SEARCH MODEL WITH OCCUPATION DESIGN

2.1 Introduction

The rich empirical findings on skill mixing pose challenges in understanding their underlying forces. I build a directed search model with several novel features to investigate the mechanisms of skill mixing. First, the model represents both firms and workers through multi-dimensional skills, laying the basis for an examination of skill mixing. Second, before producing with workers, firms of both vacant and incumbent jobs will need to design their occupations, incurring a cost payable upon operating the occupation that depends on their skill demand choices, as in Acemoglu (1999).[1] This endogenous occupation design is crucial in delivering the dynamic choices of skill mixing based on the skill distribution in the labor market. Third, the model incorporates non-linear production and cost technologies, departing from the common assumption of linear production functions in standard search models. This non-linearity allows the model to capture the varying degrees of skill complementarity in production and the increasing marginal costs of combining skills in occupations.

The model provides insights into changes in skill mixing, wages, and employment that are tied closely to the empirical observations. Central to its insights

---

[1] The endogenous choices of the intensity of inputs were first studied in the appropriate technology literature (Atkinson and Stiglitz 1969; Basu and Weil 1998; Acemoglu and Zilibotti 2001; Jones 2005; Caselli and Coleman 2006; León-Ledesma and Satchi 2019). Several studies in the labor literature allow firms to adjust labor usage as well as the quantity margin. In Lazear (2009), firms choose the weight on the skills workers supplied; in Eeckhout and Kircher (2018), firms trade-off between more versus higher quality workers; allows firms to choose appropriate skills given equilibrium skill prices.

is the idea that, as skills become more complementary in production or as their marginal costs increase, firms find it more profitable to mix skills than to specialize. Further, in designing the occupations, firms take into account the skills different workers bring and the likelihood of employing those workers. The model further links the production and cost technology, as well as worker skill supply adjustment to wage and employment distributions.

I then quantitatively evaluate the model to assess the relative importance of various channels' contributions to the observed skill mixing and to investigate their implications for wages and employment. Using two periods of NLSY data, I calibrate the model parameters by targeting the wage and employment distribution across different occupation and worker types, as well as the degree of skill mixing of occupations. Besides matching these targeted moments closely, the model replicates well the wage returns of skill mixing. The calibration results reveal that in a multi-dimensional matching framework, skills are substitutable in production, and firms face increasing marginal costs in operating occupations. Notably, sizable technology shifts have occurred: from the early 2000s to the late 2010s, there has been an increase in the complementarity of skills in production and also in firms' cost of skills for occupation operation. Meanwhile, the efficiency of analytical, computer, and interpersonal skills has increased but has declined substantially for routine skill.

Counterfactual analyses further illustrate that the technology shifts reflected in the increase in skill complementarity in production and in the cost of skills for occupation operation appear as the main drivers of the increase in skill mixing. Specifically, two-thirds of this adjustment in skill mixing is attributed to enhanced skill complementarity, while the remaining third is due to changes in occupational skill costs. In contrast, the changing skill efficiencies contribute negatively to skill mixing, and the shifts in worker skill supply play a negligible role.

The forces driving skill mixing also significantly influence shifts in wage and employment distributions. For the wage premium in high-wage relative to low-wage occupations, the increasing complementarity of skills and cost of skills together account for 74 percent, while the changing skill efficiencies contribute 26 percent. Conversely, in terms of employment gains in high-wage occupations, skill efficiencies play a more crucial role, accounting for 62 percent. These results indicate that while skill efficiency, a traditional focus of the task-biased technological change (TBTC) literature, is important in driving wage and employment dynamics, skill complementarity and cost are also pivotal factors. Additionally, a counterfactual training program that increases the mixing of non-routine skills compresses the wage disparities between skill specialists and non-specialists.

2.2 Literature Review

Theoretically, I build a directed search model with multi-dimensional skills and endogenous occupation design, following the literature on directed search (i.e., Menzio and Shi 2010, 2011; Kaas and Kircher 2015; Schaal 2017; Baley, Figueiredo, and Ulbricht 2022; Braxton and Taska 2023). Two main contributions of this model are: First, I allow firms to have endogenous skill demand in the spirit of Acemoglu (1999), which delivers the comparative statics regarding skill mixing. Second, I model skills in a multi-dimensional environment with non-linearity technologies. As such, the model incorporates directed search on both the worker and firm sides with high-dimensional heterogeneity on the two sides, which departs from most search models, but allows me to analyze the changes in skill mixing and the contribution of skill complementarity and cost factors.

The foundational model for worker sorting can be traced back to the seminal work of Roy (1951). Within this framework, occupations are treated as distinct categories, each requiring a unique skill, and workers possess skills specific to particular occupations, preventing the exploration of skill mixing.[2] An earlier tradition, including theoretical work by Shi (2001) and empirical investigations such as Hagedorn, Law, and Manovskii (2017), adopt a single-dimensional index to represent worker heterogeneity. By design, these models preclude discussions on skill mixing. A burgeoning literature explores the multidimensional matching of workers and firms that features two-sided heterogeneity and skill transferability (i.e., Yamaguchi 2012; Lindenlaub 2017; Lise and Postel-Vinay 2020). While much of this literature focuses on the assortative nature of worker-firm matching and the evolution of worker skills[3], this study instead examines firms' endogenous skill demand trade-offs in response to technological advancements or shifts in skill supply.

A related literature, inspired by Rosen (1983), Murphy (1986), and Heckman and Sedlacek (1985), features skill indivisibility or bundling, allowing for nonlinear wage schedules and a flexible degree of occupational specialization. Choné and Kramarz (2021) introduce a skill bundling framework featuring heterogeneous firms and using Swedish matched employer-employee data, they find

---

[2] In Roy or Ricardian type of models, workers will also specialize in a particular skill based on comparative advantages, making it harder to study skill mixing's implications for workers.

[3] A notable exception is Ocampo (2022), which introduces the optimal combination of tasks, leading to endogenous occupational heterogeneity.

that generalist workers earn more over time. In a separate study, Edmond and Mongey (2021) show that when skill are priced differently across occupations, firms tend to adopt technologies that reflect these skill prices, leading to opposing within-occupation changes in inequality. A critical aspect of these models is the need to take a stance on the aggregation of worker skills within firms, as discussed by Eeckhout and Kircher (2018). Different from this approach, I apply a matching model to address the indivisibility of skills and endogenous skill demand at the worker level, inherently delivering nonlinear wages and skill mixing.

Quantitatively, I provide model-based identification of the elasticity of substitution parameters among a number of different skills and the relevant occupation operation cost parameters under a tractable general equilibrium model of the labor market with endogenous skill intensities. These results contribute to the recent work on task-based models that has typically assumed exogeneity of the elasticity of substitution among different types of skills (i.e., Autor, Levy, and Murnane 2003; Autor and Dorn 2013), and also relates to studies on the elasticity of substitution among different types of workers (Johnson 1997; Heckman, Lochner, and Taber 1998; Krusell et al. 2000).

2.3 A Directed Search Model with Occupation Design

The rich empirical findings on skill mixing pose challenges in understanding their driving forces. In what follows, I attempt to provide an overarching framework to investigate the mechanisms. For this purpose, I build a directed search model with several novel features: First, both firms and workers are represented by multi-dimensional skills; Second, firms must make decisions about occupation design before producing with workers, a process that involves a cost payable upon operating the occupation as in Acemoglu (1999);[4] Third, the model incorporates non-linear production and operation cost technologies. Despite the rich setup, the model remains tractable satisfying Block Recursivity as in Menzio and Shi (2011). Under these specifications, the model offers clear insights regarding changes in skill mixing, wages, and employment that are linked to the empirical findings.

2.3.1 Environment

Workers: Time is discrete. At each period, there is a unit measure of heterogeneous workers that lives forever. Each worker of type i is characterized by a vector of multi-dimensional skills $\mathbf{x}^i = \{x^i_1, ..., x^i_k, \ldots, x^i_K\} \in S \subset \mathbb{R}^{K+}$, where K is the dimension of a closed skill space S. Workers draw their initial skill vectors at the beginning of the period from an exogenous distribution $G(\mathbf{x})$. Workers are risk-neutral, have linear utilities over consumption, and discount the future with a factor β.

Firms: On the other side of the market, there is a mass of risk-neutral firms each running one vacancy. Firms pay a cost c to post their vacancies across different occupations $j = \{1, ..., J\}$, with $J \geq 2$. Each occupation is characterized in the same multi-dimensional skill space as workers' skills, $\mathbf{y}^j = \{y^j_1, ..., y^j_k, \ldots, y^j_K\} \in S \subset \mathbb{R}^{K+}$, which has the interpretation of a vector of skill requirements or skill importance for each of the worker skills. Firms share workers' discount factor β.

---

[4] As such, the model incorporates directed search on both the worker and firm sides with high-dimensional heterogeneity on the two sides.

The production function of each worker-firm pair takes a CES form of the skill inputs of workers and skill requirements of an occupation that the firm operates:

$$f(\mathbf{x}^i, \mathbf{y}^j) = \left[\sum_{k=1}^{K} (x_k^i \alpha_k y_k^j)^{\sigma^j} \right]^{\frac{1}{\sigma^j}}, \tag{2.1}$$

where α_k controls the efficiency between worker skill and job skill requirements for a particular skill k, and σ^j controls the elasticity of substitution among different skills for an occupation j.[5] This production technology represents an extension of the production technology used in the multi-dimensional skill matching literature (i.e., Lise and Postel-Vinay 2020; Lindenlaub 2017; Ocampo 2022), where worker and firm attributes take a multiplicative form for their output associated with efficiency α_k differing by skill, which I term as "skill efficiency". I also allow complementarity across skills regulated by σ^j. When multi-dimensional skill distributions of workers and firms are considered, such a production technology gives a clear portrayal of the interaction between skill demand and supply, as well as the interaction among different skills.[6] Due to the one-to-one matching nature of the model, I omit the superscripts for worker and occupation skill vectors i and j in the exposition below.

A unique feature of this model is that I allow firms to actively *design* the jobs before producing with the worker (Acemoglu 1999), delivering endogenous demand specialization and the degree of skill mixing. Specifically, firms with both filled and unfilled vacancies design their occupations by optimally choosing the occupational skill requirements $\mathbf{y}$ in each period.[7] Such an intensity choice of

---

[5] Since labor is the only input in the model, it can be understood as "equipped" labor, and occupations' skill requirement or importance $\mathbf{y}$ takes a factor augmenting form, essentially acting as demand shifters.

[6] As such, the model explores both the role of changes in relative input efficiency that is the focus of task-based literature, and changes in skill complementarity.

[7] This feature is consistent with the empirical finding that both incumbent jobs and vacancies

Figure 2.1: Model Timing

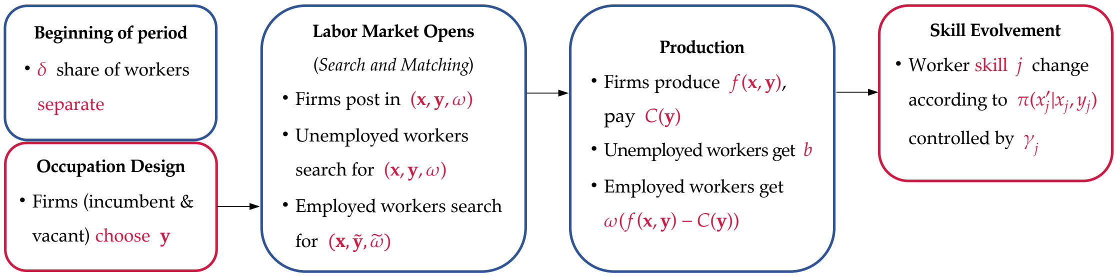

skills in an occupation alters the efficiency of that skill will and essentially leads to different production technologies for that occupation, capturing the overall quality of the occupation and the optimal degree of skill mixing. In designing the occupation, both worker skill profiles and skill complementary play an important role, as firms would want to exploit what skills the workers supply given the technology.[8]

Nonetheless, such a job design incurs a cost $C(\mathbf{y})$ that is payable upon producing with a worker. This cost is convex and strictly increases in the skill level that the firm chooses ($\frac{\partial C(\mathbf{y})}{\partial y_k} > 0, \frac{\partial^2 C(\mathbf{y})}{\partial y_k^2} > 0, \forall k$), and represents the necessary expenses to operate an occupation for a given skill requirement choice. This cost can be understood as an operation cost as in Hopenhayn (1992) that increases in the skill level of the occupation.[9]

Labor Market: There is a continuum of submarkets that are indexed by worker

---

have changing degrees of skill mixing.

[8] For example, in designing an occupation (i.e., salespersons) for lower-skill workers who might have a greater supply of interpersonal than analytical skill, firms may want interpersonal skill to be more intensively used to take advantage of the labor supply; on the contrary, if online marketing has increased the complementarity between analytical and interpersonal skills, firms may adjust accordingly to let the analytical skill to be more intensive. In equilibrium, this endogenous intensity will depend on other forces in the model.

[9] For example, to design and operate an occupation that employs high-skill workers, a firm will need to incur higher expenses in terms of better offices and equipment rentals.

and occupation skill profiles $(\mathbf{x}, \mathbf{y})$, as well as the share of worker-firm surplus ω that firms promise to workers.[10] Workers with skill profile $\mathbf{x}$ direct their search towards different occupations and surplus shares, meeting one vacancy at a time. Matching between workers and firms is frictional and is regulated by a standard constant to scale matching function. Under this directed search environment, each submarket has a separate tightness (vacancy-unemployment ratio), denoted by $\theta(\mathbf{x}, \mathbf{y}, \omega)$. In each submarket, workers find job with probability $p(\theta(\mathbf{x}, \mathbf{y}, \omega))$ and firms fill the vacancy with probability $q(\theta(\mathbf{x}, \mathbf{y}, \omega)) = p(\theta(\mathbf{x}, \mathbf{y}, \omega))/\theta(\mathbf{x}, \mathbf{y}, \omega)$.[11]

The timing of the model evolves as follows. At the beginning of each period, a fraction δ of worker-firm pairs separate exogenously. Before the labor market opens and unlike standard search models, firms of both unposted vacancies and incumbent jobs will first need to design the occupations at this stage. The labor market that is comprised of different submarkets then opens, and both unemployed and employed search for unfilled vacancies and form matches with firms under the constant return to scale matching technology. The labor market then closes, firms produce the output and pay the occupation operation cost as well as the wage, which is a share of the surplus. Unemployed workers receive a transfer with a value of b. Lastly, workers are able to learn by doing, and their skills evolve according to the Markov process depending on their employment status, as described below.

Aggregate and Individual State: The aggregate state of the economy is the distribution of workers across employment status, skill profiles, occupational

---

[10] This arrangement can be considered as an employment contract simply specifies the surplus share ω promised to the worker contingent on the state for the current period, as well as the continuation value of the match in the subsequent period (see next section). The contract is assumed to be fully committed by both the workers and firms.

[11] Functions p and q also satisfy usual regularity conditions: twice continuously differentiable; $p'(\theta) > 0, p''(\theta) < 0, p(0) = 0; q'(\theta) < 0, q''(\theta) > 0, q(0) = 1$.

skill requirements, and surplus shares, denoted as $\psi \in \Psi$. I subsume aggregate state in the exposition of model equilibrium in the next section and show that in fact, the model equilibrium is independent of the aggregate state.

Nonetheless, in the model, I allow workers to learn on the job, and their subsequent skill profiles are contingent upon their current employment status, as in Lise and Postel-Vinay (2020). Specifically, considering each skill j in the worker's skill profile $\mathbf{x}$ as an element of the finite set S, the evolution of this skill follows a Markov process $\pi(x'_j|x_j, y_j)$, conditional on the worker's current skill level and employed occupation. If a worker is matched with an occupation that requires a skill level exceeding his or her own ($x_j < y_j$), the worker's skill j will adjust upward in the next period: $x'_j > x_j$, and the inverse applies for a worker whose skill is lower than the requirements of their current occupation. The probability, or the speed of skill adjustment, is contingent upon the specific skill j. For unemployed workers, they are treated such that their present occupation demands a zero level for all skills. The calibration of the skill adjustment probability is discussed in Section 2.4.

2.3.2 Model Equilibrium

I will now characterize the optimal strategies for workers' job search and firms' job creation and continuation. The value functions for workers are described at the point of the production stage when the labor market comes to a close, while for firms I also consider the job design stage before the labor market opens.

Worker's Problem: Let $U(\mathbf{x})$ denote the value of being unemployed and searching for a job for worker $\mathbf{x}$. Similarly, let $W(\mathbf{x}, \mathbf{y}, \omega)$ be the total discounted returns from holding a job of skill requirements $\mathbf{y}$ and surplus share ω at time t. These values can be written as:

$$\begin{aligned}
U(\mathbf{x}) &= b + \beta E\Big\{\max_{\mathbf{y}',\omega'}\; p(\theta(\mathbf{x}',\mathbf{y}',\omega'))W(\mathbf{x}',\mathbf{y}',\omega') \\
&\quad +[(1-p(\theta(\mathbf{x}',\mathbf{y}',\omega'))]U(\mathbf{x}')\Big\} \\
W(\mathbf{x},\mathbf{y},\omega) &= \omega(f(\mathbf{x},\mathbf{y})-C(\mathbf{y}))+\delta U(\mathbf{x}') \qquad (2.2) \\
&\quad +\beta(1-\delta)E\Big\{\max_{\tilde{\mathbf{y}}',\tilde{\omega}'}\; p(\theta(\mathbf{x}',\tilde{\mathbf{y}}',\tilde{\omega}'))W(\mathbf{x}',\tilde{\mathbf{y}}',\tilde{\omega}') \\
&\quad +[(1-p(\theta(\mathbf{x}',\tilde{\mathbf{y}}',\tilde{\omega}'))]W(\mathbf{x}',\mathbf{y}',\omega)\Big\}
\end{aligned}$$

Unemployed workers gain a utility b through the current period's transfer. In the subsequent period, their skills may transition to $\mathbf{x}'$, which are likely to depreciate due to their unemployed status. Meanwhile, within the submarket that aligns with their skill profiles, workers engage in the search for vacancies that span a variety of occupations $\mathbf{y}$ and surplus shares ω, looking for the highest continuation value. In choosing $\mathbf{y}$ and ω, workers face the tradeoff between the value of employed and the success probability of a match $p(\theta(\mathbf{x}', y', \omega')$, both of which hinge on the occupation and surplus share that the workers target. Should the match prove successful, the workers enjoy the continued value that employment offers; otherwise, their status of unemployment persists.

Workers currently employed in a firm characterized by $(\mathbf{y}, \omega)$ receive a wage equivalent to the share ω of the output from their match. When the subsequent period arrives, they face a probability δ of an exogenous separation, in which case they become unemployed with a value $U(\mathbf{x}')$ and engage in job search immediately. Employed workers perform on-the-job searches in their current match for new occupations and surplus shares $(\tilde{\mathbf{y}}', \tilde{\omega}')$, on the premise that there is a positive probability $p(\theta(\mathbf{x}', \tilde{\mathbf{y}}', \tilde{\omega}'))$ that the continuation value from the new match offers exceeds that of the original firm. In the absence of such possibilities or if the transition is not successful, the worker remains with the initial firm.

Firm's Problem: Consider a firm running occupation $\mathbf{y}$, offering surplus share ω, and employing worker $\mathbf{x}$. Let $J(\mathbf{x}, \mathbf{y}, \omega)$ denote the total discounted profits to this firm:

$$\begin{aligned} J(\mathbf{x}, \mathbf{y}, \omega) = \; & \max_{\mathbf{y}} (1 - \omega)(f(\mathbf{x}, \mathbf{y}) - C(\mathbf{y})) \\ & + \beta(1 - \delta)E\Big\{(1 - p(\theta(\mathbf{x}', \tilde{\mathbf{y}}', \tilde{\omega}')) J(\mathbf{x}', \mathbf{y}', \omega)\Big\} \end{aligned} \tag{2.3}$$

In the current period, firms receive a portion $(1 - \omega)$ of the worker-firm surplus, after paying the workers their wages. In the production process, firms also need to cover the occupation operation cost $C(\mathbf{y})$, which depends on the skill levels required by the occupation that the firms designed. The labor market operates under free entry for firms, hence, maintaining a vacancy bears no value. In the case of exogenous separation, or with a probability $p(\theta(\mathbf{x}', \tilde{\mathbf{y}}', \tilde{\omega}'))$ that the worker finds another job at an optimal occupation $\tilde{\mathbf{y}}'$ and surplus share $\tilde{\omega}'$ through on-the-job search, the firm accrues no profits. In the case where the match persists, the firm continues to acquire discounted profits from the match.

$$c = \beta E\Big\{q(\theta(\mathbf{x}, \mathbf{y}, \omega)) J(\mathbf{x}, \mathbf{y}, \omega)\Big\} \tag{2.4}$$

The free-entry condition further highlights firms' choice of optimal degree of skill mixing and the tradeoff that agents face in the model. Prior to the opening of the labor market in each period, firms of incumbent jobs and unfilled vacancies re-design the occupation, taking into consideration the overall production technology and worker skills within their respective submarkets.[12] Given that the value of a vacancy is zero, firms will opt for an optimal skill mixing that equates the firm's anticipated discounted profits to the cost of vacancy posting as

---

[12] Considering that incumbent firms and new entrants utilize identical production technologies and confront the same worker skills within each submarket, their choices align.

in equation (2.4). This condition implicitly pins down market tightness $\theta(\mathbf{x}, \mathbf{y}, \omega)$. If an occupation for a specific worker type becomes more profitable, the number of vacancies posted will increase, leading to a rise in market tightness but at the same time a reduction in the job-filling rate.[13]

The free entry condition also reflects the tradeoff faced by workers. Since workers receive the remaining surplus claimed by the firms, in markets with higher job-finding probabilities (i.e., tighter markets), the value of employment is likely to be lower. Workers' job-finding probability also feeds back to firms' discounted profits through worker on-the-job search and the chance that the firm attracts other employed workers.

Block-recursive Equilibrium: Despite the multi-dimensional skill setup, the model still achieves analytical tractability by relieving the dependence on the entire distribution of agents across aggregate states in characterizing agents' value functions and market tightness. Such a convenient feature was coined as "block-recursive" in Menzio and Shi (2010) and Menzio and Shi (2011) for a broad range of directed search models.[14] This is a result of two features of the model. First, as search is directed and workers choose optimally the occupation and surplus share, their life utility does not depend on their outside options, and workers do not need to forecast the wage depending on the entire distribution of employment. Second, there are separate markets for workers of different profiles, and workers search for jobs within their own submarket, in which firms carry different occupations. This additional degree of separability implies that

---

[13] As in other directed search models, only a portion of submarkets may open in equilibrium, depending on firm's value and corresponding market tightness in different markets

[14] Block recursivity allows not only analytical tractability but also enables standard numerical techniques to solve the model. The framework considered in this paper involves more heterogeneity and requires an additional degree of directness, as discussed.

the market tightness of a submarket is independent of the worker distribution in other markets, relieving the burden of workers and firms to forecast other markets in making their decisions.[15] In online Appendix B.1.2, I formally define a Block-recursive equilibrium for the economy and show its existence and uniqueness.

Skill Mixing, Wages and Employment: The model yields several predictions regarding changes in skill mixing, wages, and employment that align closely with the empirical findings detailed in Sections 1.3 and 1.4. These predictions emphasize the role of skill complementarity within a production framework that features indivisible skills. The formal propositions and proofs of these outcomes can be found in the online Appendix B.1.1, and a concise discussion is provided here.

Under the production technology described in equation (2.1) and the occupational design cost $C(\mathbf{y})$ that is strictly increasing and convex in skill requirements, increased complementarity in production or a higher degree of increasing marginal costs leads firms to find it more profitable to employ a mixture of different skills rather than specializing in one, leading to increased skill mixing. Additionally, the supply of skills by workers influences these outcomes: as workers supply a more diverse set of skills, it becomes more efficient to design jobs that require this mix of skills. In terms of wages and employment, if skills become more complementary in production or less costly to combine, the output of the worker-firm match rises, leading to wage increases. Through the

---

[15] Such additional directness implies that, i.e., computer scientists only confront other computer scientists in job search, while sales clerks only compete with other sales clerks. In reality, the degree of separability will depend on specific occupations and the overall economic condition. As reported by Osberg (1993), search directedness is procyclical and is higher when the market is tight. In bringing the model to the data, I use economic recovery periods and more coarse occupations to be consistent with the model.

free entry condition specified in equation (2.4), this increased joint worker-firm value results in a tighter labor market and an elevated job-finding probability for unemployed workers. I quantitatively calibrate the model and test these predictions in the next section.

2.4 Model Quantification

I will now calibrate the model to evaluate the quantitative significance of various channels contributing to skill mixing and examine their implications for wages and employment. First, I outline the data construction and measurement, followed by a discussion on calibration strategy and estimated parameters. Next, I analyze worker sorting and job ladder under the baseline calibration. In Section 2.5, I will perform counterfactual analyses to decompose the shifts in skill mixing as well as employment and wage distributions.

2.4.1 Measurement and Calibration

I apply the same combination of NLSY 79 & 97 along with O*NET data as in Section 1.4 to calibrate the model. The datasets provide counterparts to the model variables: the worker abilities correspond to worker skills ($\mathbf{x}$) as discussed in Section 1.4, while O*NET provides occupational skill requirements ($\mathbf{y}$). NLSY also provides information on employment distribution and wage levels. The model is calibrated to two periods of data from the early 2000s to late 2010s separately, which coincides with a substantial shift in skill mixing and abstracts from the great recession. Specifically, the steady state of the model is fitted to the data from 2005–2006 and 2016–2019 to ensure comparability of sample sizes across these two periods, and I restrict to those workers with information on their skills.[16] Finally, for both worker and job skill profiles, I consider the same set of skills (analytical, computer, interpersonal, routine) as in Section 1.3 and 1.4, only that I combine analytical and computer skills to have a three-dimensionality feasible for quantitative analysis ($K = 3, k = \{$analytical/computer (a), interpersonal (p), routine (r)$\}$).[17]

Considering the potential influence of skill supply variation on skill mixing, I calibrate two key aspects of it. First, the distribution of worker skills $G(\mathbf{x})$ varies *across* the two data periods to align with the workers' choice of occupations and college majors (if attended) as in the NLSY data, following Lise and Postel-Vinay (2020). Specifically, a worker accumulates γ_j times the gap between the worker's endowment and an occupation's or college major's requirement of skill j in each year, with γ_j depending on upward or downward accumulation. Table 2.2 panels B shows the calibrated γ_j using the estimates from Lise and Postel-Vinay (2020).[18] Second, the Markov process $\pi(x'_j|x_j, y_j)$ for worker skill adjustment *within* a model steady state is calibrated to make a worker's skill level adjust upward or downward in the next model period if the worker's occupation requires a higher or lower skill level than the worker has. The Markov adjustment probability is equal to the annual skill adjustment rates *gamma*$_j$ scaled by the gap between the worker's skill set and the occupation's demands.[19] Online Appendix B.1.4

---

[16] NLSY 1997 was conducted annually during 2005-2006, but only biannually in 2016-2019, so does NLSY 1979 for the later period. The same sizes for the two selected periods are 30,654 and 43,340 respectively.

[17] As I merge analytical and computer skills into one for calibration using their average values, I denote this combined skill as "analytical/computer".

[18] Workers' skills can be lost when not employed but cannot be lower than their initial endowments. For skill changes while in school, I specify that workers spend on average 3 years learning the skills of their majors.

[19] Specifically, the Markov probability of upward adjustment is determined by: $\frac{x_j^{up}-x_j}{y_j-x_j}\mathbf{1}(x_j^{up} <$

Figure 2.2: Worker Skill Distribution Shifts

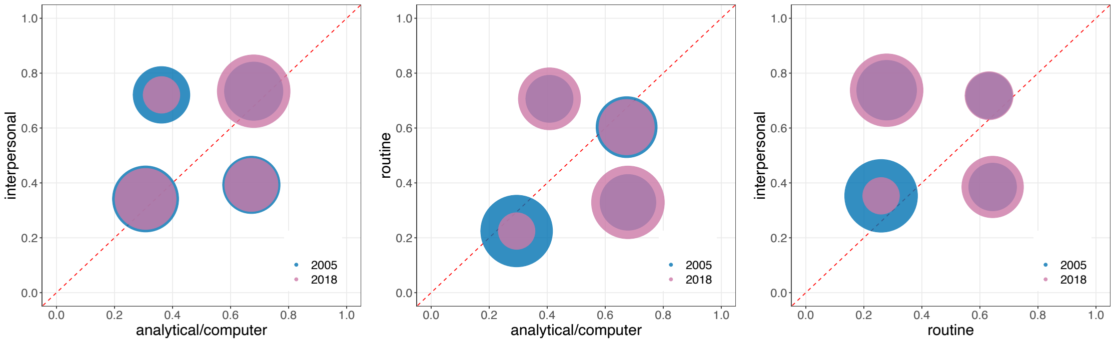

Notes: This figure illustrates the evolution of the skill distribution for different types of workers over the years 2005 (shown in blue) and 2018 (shown in cranberry), across three distinct two-dimensional skill spaces. These worker skills are measured using data from NLSY79&97, with the specific skill measure discussed in the online Appendix Table B.1. Skill variations of these worker types are calibrated based on the skill accumulation and depreciation rates associated with different occupations and college majors, using the estimates of by Lise and Postel-Vinay (2020).

provides further details of the skill supply calibration.

Figure 2.2 presents the calibrated variation in skill distributions from 2005 (in blue) to 2018 (in cranberry) for four worker types within each two-dimensional skill space. The circle sizes represent the probability of corresponding skill combinations. In the analytical/computer and interpersonal skill space, which are non-routine skills, there is a noticeable shift towards greater skill mixing. This is evidenced by an increase in workers possessing both high analytical/computer and interpersonal skill, indicated by greater areas representing these skill combinations. This shift towards more mixed skill sets is due to the rising mixing of skill demand and the worker learning by doing. In contrast, in skill spaces

---

$y_j) \times \gamma_j^{up}$, and of downward adjustment is given by: $\frac{x_j^{down}-x_j}{y_j-x_j}\mathbf{1}(y_j < x_j^{down}) \times \gamma_j^{down}$. Here, x_j represents the current grid value of worker skill j, while x_j^{up} or x_j^{down} denotes the value of worker skill j up or down a grid, respectively. The indicator variables $\mathbf{1}(y_j < x_j^{down})$ or $\mathbf{1}(x_j^{up} < y_j)$ evaluates whether the skill j grid value of the worker's current employed occupation is greater or smaller than the value of the worker's skill j grid. This means that a worker will only adjust up or down a grid if the occupation's skill is larger or smaller than the corresponding up or down grid value for the worker's skill.

involving routine skill, there is a clear trend towards specialization, indicated by an increased area in off-diagonal skill combinations. This shift is similarly driven by a growing demand for specialized routine skill relative to others. The implications of these changes in skill supply on labor market dynamics are explored in the next section.

To map occupations and workers in the model to the data, I set grid points as follows. I classify occupations into high- and low-wage, as in Section 1.4, with the former group including managerial, professional, and white-collar occupations, and the latter blue-collar and service occupations. The grid point for an occupation's requirement of a skill y_j is set such that moving up one grid corresponds to 50 percent of the average observed value of y_j for that occupation.[20] On the worker side, workers are classified based on their skill level x_j: those with skills above the average are deemed high type and assigned the mean of the above-average values; those below the average are considered low type and assigned the mean of the below-average values.[21]

Functional Forms: The model is parameterized as follows. The multi-dimensional skill production function is defined as in equation (2.1), which accommodates cross-skill complementarity controlled by σ and enables a sensible interaction between skill demand and supply, in line with the multi-dimensional matching literature (i.e.,Lise and Postel-Vinay 2020; Lindenlaub 2017; Ocampo 2022). I allow a flexible form of occupation operation cost $C(\mathbf{y})$ as in equation (2.5), where ρ regulates the convexity of the occupation operation cost function

---

[20] As the model calibration uses data of two periods with a consistent grid, I determine grid points by averaging the occupation's median values across both periods.

[21] With three chosen skills, there are 8 worker types in the model.

with respect to skill requirements, and τ governs the scale of the cost.[22]

$$C(\mathbf{y}) = \tau[\sum_{k=1}^{K}(y^k)^\rho] \tag{2.5}$$

The matching function assumes a standard Cobb-Douglas form, $M(s, v) = \mu s^\eta v^{1-\eta}$, indicating that η is the elasticity of matches concerning total search effort, and μ is the matching efficiency. This function form leads to the job finding rate being $p(\theta) = \mu\theta^{1-\eta}$ and the vacancy filling rate being $q(\theta) = \mu\theta^{-\eta}$.

2.4.2 Calibration Strategy

The calibration of parameters falls into three categories. For parameters that regulate the search environment, I follow closely the conventions of the search and matching literature. I rely on estimates from the multi-dimensional matching literature for the skill adjustment and skill efficiency parameters. Lastly, for the parameter regulating elasticity of substitution across skills and relating to costs, I estimate them internally through Simulated Methods of Moments (SMM).

External Calibration: The model period is a year. Given that all agents are risk-neutral, the discount rate β is assigned a value of 0.96, corresponding to an annual interest rate of 4 percent. The job separation rate δ is set at 10 percent as in Shimer (2005). For employed workers, their share of output ω is set at 0.6, mirroring the labor share of GDP in 2005. For unemployed workers, the unemployment benefits b is set at 41.5 percent of the earning loss of lowest-paid occupations, following the estimates of Braxton, Herkenhoff, and Phillips (2020). The elasticity of the matching function η is set at 0.5 as is standard, and the

---

[22] Besides technical convenience, the functional form (2.5) also implies that for a given cost, firms need to trade off the choice of altering different skill intensities.

Table 2.1: Moments and Model Match

	First Period		Second Period	
	Data	Model	Data	Model
A. Worker moments				
Relative wage of high type				
Analytical/computer	1.46	1.62	1.60	1.78
Interpersonal	1.05	1.09	1.20	1.25
Routine	1.12	1.23	0.92	1.21
Wage return of skill mixing (untargeted)	0.07	0.04	0.07	0.04
Unemployment Rate	0.05	0.03	0.04	0.04
B. Occupation moments				
Relative wage of high skill	1.30	1.07	1.56	1.38
Corr. wage & abilities (low-wage)	0.23	0.23	0.49	0.49
Corr. wage & abilities (high-wage)	0.35	0.32	0.60	0.71
Employ. share (low-wage)	0.43	0.31	0.37	0.09
Employ. share (high-wage)	0.57	0.69	0.63	0.91
100 × Skill mixing (low-wage)	97.54	95.11	98.96	98.82
100× Skill mixing (high-wage)	95.74	96.03	94.12	94.60

Notes: This table reports the average values of the targeted moments both in the data and through model simulation. The data used for the moment calculation and for SMM estimation are two periods of pooled NLSY79&97 for employed workers: period 1 from 2005–2006 and period 2 from 2016–2019. Two types of moments are included. The worker moments include the relative wage of high type workers as well as the unemployment rate. The occupation moments include the relative wage of high skill occupations, the employment share and the skill mixing index of RNR skills in low and high skill occupations.

matching efficiency μ is set to 0.65, as in Mercan and Schoefer (2020). Table 2.2 panel A summarizes these externally calibrated parameters.

I calibrate the speed of skill adjustment (γ_j) and the skill efficiencies (α_k) following Lise and Postel-Vinay (2020) and Lindenlaub (2017), as detailed in Table 2.2 panels B and C. The calibration aligns the adjustment of analytical/computer, interpersonal, and routine skills with the cognitive, interpersonal, and manual skills detailed in Lise and Postel-Vinay (2020).[23] Analytical/computer skill

---

[23] Lise and Postel-Vinay (2020)'s estimates are presented on a monthly basis, which I have

Table 2.2: Parameter Estimates

Parameter	Description	Value	
	A. Externally calibrated - search		
β	Discount Rate	0.96	
δ	Job separation rate	0.10	
ω	Worker share of surplus	0.60	
b	Unemployment benefit as a share of output	0.42	
η	Elasticity of the matching function	0.50	
μ	Matching efficiency	0.65	
	B. Externally calibrated - skill adjustment	*Up*	*Down*
γ_a	Annual adjustment speed of analytical/computer skill	0.36	0.10
γ_p	Annual adjustment speed of interpersonal skill	0.05	0.00
γ_r	Annual adjustment speed of routine skill	1.00	0.36
	C. Externally calibrated - skill efficiency	*Period 1*	*Period 2*
α_a	Skill efficiency of analytical/computer skill	0.63	0.95
α_p	Skill efficiency of interpersonal skill	0.05	0.08
α_r	Skill efficiency of routine skill	0.14	0.06
	D. Internally estimated	*Period 1*	*Period 2*
σ^{low}	Elasticity parameter of skills in production (low-wage)	0.64	0.41
σ^{high}	Elasticity parameter of skills in production (high-wage)	0.60	0.36
τ	Scaler of occupation operation cost	0.74	0.53
ϕ	Convexity of occupation operation cost	3.63	4.90
c	Vacancy posting cost as a share of output	0.56	0.82

Notes: This table shows the exogenously calibrated as well as internally estimated parameters. The data used for the internal estimation are two periods of pooled NLSY79&97 data for workers with information on their pre-market abilities. Period 1 is from 2005–2006 and period 2 from 2016–2019.

adjusts upward two times faster than it depreciates, while interpersonal skill changes slowly in both directions. Routine skill adjusts most rapidly in either direction. I linearly interpolate Lindenlaub (2017)'s estimates of skill efficiencies for my period of analysis.[24] Between 2005 and 2018, the productivity of analytical/computer and interpersonal skill in matching worker abilities with job

---

adjusted to an annual scale.

[24] Lindenlaub (2017)'s estimates span from 1990 to 2010.

skill requirements increased by about 60 percent. In contrast, the productivity of routine skill saw a decrease of more than 50 percent.

Internal Estimation: For the internal estimation, the SMM procedure initiates by determining the agents' steady-state policies based on the model's parameters, simulating a cohort of workers. Each simulation results in a distribution of employment statuses and corresponding labor market outcomes. The parameters are then estimated minimizing the distance between simulated and empirical moments.[25] The estimation targets 11 moments as shown in Table 2.1 for both periods of data that include: i) the relative wage of the high-type worker for each skill; ii) the unemployment rate; iii) the relative wage of high-skill occupation; iv) the within-occupation correlation between wages and worker abilities; v) the share of employment across occupations; and vi) the skill mixing index of RNR skills of occupations.[26] The model does a decent job of matching all the moments.

The model parameters are jointly identified from the moments, for which a concise summary of the key information for identification is given below with a more detailed discussion in online Appendix B.1.3. I first identify the complementarity parameter of skills in production σ targeting the correlation of within-occupation relative wages and worker skills. The cost parameter ρ is then estimated by leveraging the firm's optimization conditions in skill mixing. Conditional on parameters estimated at the production side, the employment distribution and relative wages further aid in estimating τ. Lastly, the unemploy-

---

[25] Online Appendix B.1.5 provides further details on the numerical implementation.

[26] All moments are directly computed from the two periods of data from NLSY, except for unemployment, for which I use the statistics from the Bureau of Economic Analysis (BEA) to avoid the age composition effects present in NLSY. For example, by the late 2010s, a larger segment of the NLSY 79 cohort was above age 50, making them more likely to be out of the labor force. Additionally, the unemployment rate from NLSY, derived from the number of jobs held since the last survey, averages 9 percent, notably higher than BEA data. However, this decision primarily affects vacancy posting cost parameters.

ment rate disciplines the vacancy posting cost.

Table 2.2 panel D presents the internally estimated parameters, which indicate considerable technological shifts between the two periods. For the initial period, the estimated σ is 0.64 and 0.6 for low- and high-wage occupations respectively, suggesting that skills are substitutable in production. In the late 2010s, there was a significant rise in skill complementarity in production, reflected in the reduction of σ to 0.4 for both types of occupations. Firms also encounter rising costs of skills in occupation operation, as reflected in the increase of both the scale and the convexity of the cost function (τ and ρ). As discussed in Section 2.4, this increased complementarity as well as the cost of skills intensifies firms' incentives to mix skills. Lastly, the cost of posting vacancies has risen slightly post-2010s.

2.4.3 Job skills Over the Life Cycle

I now turn to discuss the paths of workers' job skill requirements or job ladder as simulated by the model, and compare them with empirical data. The model initiates with a job match determined by $G(\mathbf{x})$ and incorporates two main mechanisms that influence workers' movement across jobs overtime: active job searching, both while employed and unemployed, and the evolution of workers' skills, which is governed by a Markov process. These mechanisms provide workers with different occupational choices over time, guided by the parameters calibrated in Table 2.2. For empirical comparison, I apply the 2005 O*NET values to align with observations from the NLSY 79 and the year 2018 to align with observations from NLSY 97. I also impose age restrictions on the datasets to compare similar life stages, limiting the age range to 41 to 60 for NLSY 79 and 21

Figure 2.3: Predicted and Observed Occupational Skills Across Age Groups in 2005 vs. 2018

Panel A. Model Simulation

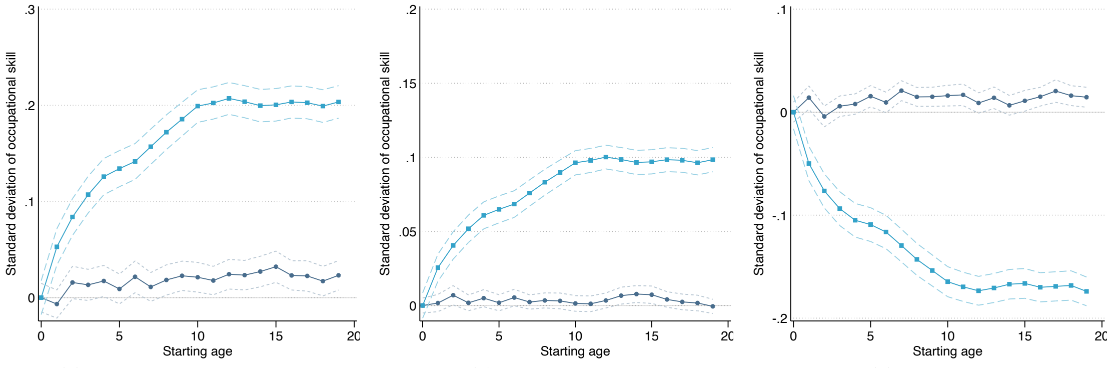

(1) Analytical/Computer (2) Interpersonal (3) Routine

Panel B. NLSY

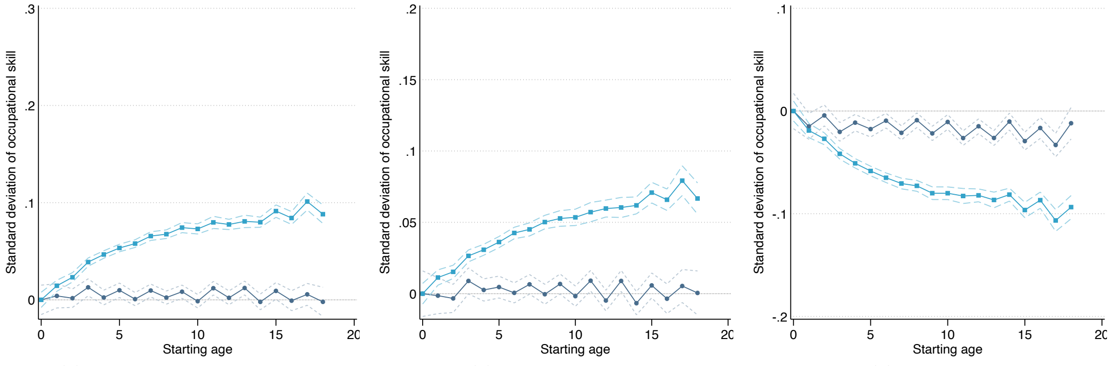

(1) Analytical/Computer (2) Interpersonal (3) Routine

Notes: These figures display the average occupational job skills over the life cycle of workers, for both model predictions and empirical data. Model simulated or emprical data points for the year 2005 are represented by circles, and for the year 2018 are depicted as squares. The emprical observations are drawn from the NLSY 79 cohort for year 2005 and NLSY 97 cohort for year 2018

to 40 for NLSY 97, similarly constraining the model-simulated paths.

Figure 2.3 illustrates the average occupational job skills over the life cycle of workers, as predicted by the model and observed empirically, depicting the years 2005 (represented with circles) and 2018 (shown with squares) across three different skills. In 2005, the job ladder for all three skills appears flat for workers aged 41 and above, as shown by both the model simulations and empirical data. However, by 2018, there is a notable increase in the job skill requirements for analytical/computer and interpersonal skills as workers age in the labor market, although the model tends to overestimate the depth of these job ladders. In contrast, both the model and empirical data indicate a decline in routine skills over time for workers in 2018. Thus, the model provides a reconciliation of the changing job ladder across these two periods throught the lens of underlying technology.

2.5 Counterfactual Analysis

What drives the observed increase in skill mixing, and what are their implications for labor market outcomes? In this section, I employ the model to perform a series of counterfactual experiments to assess the relative significance of each model channel in explaining the shifts in the degree of skill mixing. I then evaluate the influence of these channels on the changes in earnings and employment distributions.

For the counterfactual analysis, I take the 2018 economy and then sequentially remove shifts in calibrated parameters representing different channels, setting their values to that in the 2005 economy. Specifically, I examine the roles of changes in skill supply ($G(\mathbf{x})$), skill efficiencies (α_k), skill complementarity in production (σ), and occupation operation cost (τ, ϕ) in generating moment variations. Given the non-linear interplay of these forces, I remove these elements in different sequences and calculate the effect of each channel by averaging across those sequences.

Counterfactual Skill Mixing: I begin by assessing how different channels con-

Figure 2.4: Counterfactual Decomposition

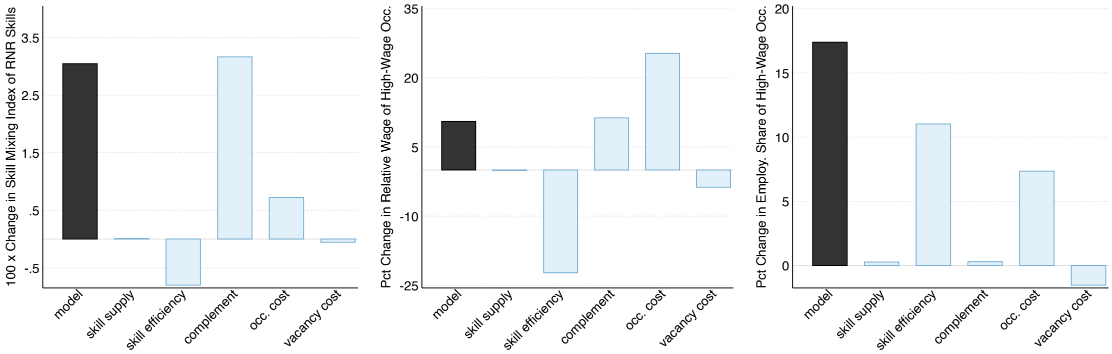

(1) Skill mixing

(2) Wages

(3) Employment

Notes: These figures plot the model generated changes in skill mixing in low-skill occupations (panel 1), changes in relative wage of high-wage occupation (panel 2), changes in employment share of high-wage occupation (panel 3), and different individual skills' contributions to relative wage and employment share changes (panel 4). Different model channels are shut down individually by eliminating the changes in calibrated values to highlight the contribution of each channel. The full model has all the model features. The values of skill complementarity in production, cost of skills in occupation operation, efficiency differential, and vacancy posting cost across the two periods are shown in Table 2.2. Worker skill supply distribution variation across the periods are calibrated according to Table B.1.

Table 2.3: Returns to Skill Specialization Decomposition

Decomposition	Analytical/ Computer	Interpersonal	Routine
Full model	15.45	15.16	-3.72
Skill supply	-2.60	-0.52	-3.13
Skill efficiency	26.59	1.60	-11.82
Complementarity	-23.86	11.01	12.33
Occ. cost	10.82	0.80	-7.42

Notes: This table shows the model-generated changes in relative wages of high-type workers for the three skills. The first row shows the changes with all model channels, corresponding to the first three rows of Table 2.1. The following rows then show the variation attributable to different model channels. See the footnote of Figure 2.4 for details.

tribute to the growth in skill mixing within low-wage occupations, which has been noticeably observed in the data.[27] The first panel of Figure 2.4 illustrates that the full model predicts a rise in skill mixing within low-wage occupations over the two periods consistent with the observed data. Changes in the supply

---

[27] Online appendix B.1.6 shows the results for high-wage occupations.

of worker skills and vacancy posting cost during these periods did not play a significant role in this rise. On the other hand, shifts in skill efficiency have had a negative contribution to the change in the degree of skill mixing. The latter result arises because, as the efficiency of routine skill declines and of interpersonal and analytical/computer skills improves drastically, firms are incentivized to redesign occupations to shift towards either analytical/computer or interpersonal skill away from routine skill, to a degree that it leads to a slight increase in specialization towards the skills that become more efficient.

The subsequent counterfactual results indicate that the rise in the complementarity of skills in production and occupational operation costs account for the increase in skill mixing. The increase in skill complementarity contributes to three-quarters of the increase, while changes in occupational operation cost account for another quarter. These results are consistent with the predictions in Section 2.1 and highlight the importance of skill complementarity and their cost in occupation operation in driving firms' endogenous skill demand specialization.[28]

Wage and Employment Effects: I proceed to investigate how the same model channels that influence skill mixing also impact wage and employment distributions. Column (2) of Figure 2.4 illustrates the changes in the relative wage between high-wage and low-wage occupations from 2005 to 2018, with the model predicting a 10 percentage point increase. As observed with skill mixing, changes in worker skill supply and vacancy posting costs have negligible impacts. In contrast, changes in skill efficiencies significantly reduce the relative wage gap, decreasing it by more than the net increase. However, the growing complemen-

---

[28] Further analysis of the the implications of τ and ϕ individually for skill mixing changes shows that ϕ plays a bigger role.

tarity and changing costs of skills notably increase wage disparities, contributing to 40 and 60 percent, respectively, of the overall rise in wage premiums for high-wage occupations.

In Table 2.3, I analyze the drivers of wage returns for specializing in different skills. The full model indicates a 15 percent increase in wage returns for specializing in analytical/computer and interpersonal skill, while returns for specializing in routine skill decrease by 3 percent. Variations in skill supply have negatively impacted these returns. In terms of remaining forces, for analytical/computer skills, the significant rise in its efficiency is the primary factor, leading to a 30 percent increase in returns, followed by changes in skill costs; the increase in skill complementarity has reduced these returns. In contrast, for interpersonal and routine skill, the increase in complementarity is the principal driver, leading to an 11 and 12 percent rise in returns, respectively. For routine skill, the drastic decrease in efficiency and the increase in skill costs have led to a reduction in its returns.

I further decompose the influence of various model factors on the increase in employment in high-wage occupations over two periods, as depicted in column (3) of Figure 2.4. The full model shows a 17 percentage points rise in the employment share of high-wage occupations. The skill supply has only a marginal impact, similar to its effect on relative wages. The most significant factor is the change in skill efficiencies, accounting for 62 percent of the overall increase. Changes in skill cost are also important, contributing to 37 percent of the increase. However, the rise in complementarity plays only a marginal role in this growth. In the online Appendix B.1.6, I conduct a further decomposition of the contributions of individual skills to changes in wages and employment.

The results demonstrate that the rise in analytical/computer skills is the most significant factor influencing employment, whereas the decline in routine skills drives the wage trends.

In summary, the counterfactual analysis demonstrates the significance of the growing complementarity of skills in production, the changing skill cost for occupation operation, and shifts in skill efficiency as primary drivers behind the observed changes in skill mixing and the increases in wage and employment gains in high-wage occupations.

Discussions:

Task-biased vs. Skill Complementarity & Cost: Biased technological change, especially task-biased technological change (TBTC), is shown to be a key driver of the recent trends in wage inequality in developed countries. Studies by Costinot and Vogel (2010) and Acemoglu and Zilibotti (2001) employ one-dimensional assignment models, while Lindenlaub (2017) uses a multi-dimensional assignment model to examine this phenomenon. This change is characterized by increasing complementarities and efficiency in cognitive tasks and a decline in routine tasks. Such a shift leads to the replacement of workers in medium-wage occupations and an increase in wages and employment in high-wage occupations.

In my model, I incorporate both changes in skill efficiency representing TBTC, as well as variations in skill complementarity and cost. My counterfactual analysis first confirms the significance of TBTC, which accounts for 62 percent of the employment gains in high-wage occupations, whereas skill complementarity and cost account for the rest. However, for the wage premium in high-wage occupations, skill complementarity and cost are more crucial, contributing three-quarters of the change, whereas TBTC accounts for only a quarter. For the wage return to skill specialization, TBTC emerges as the primary driver for analytical/computer skills; however, for interpersonal skill, skill complementarity plays a more significant role, and it also increases the return to specialization in routine skill, even though TBTC has a reduced it. The increase in skill mixing is entirely attributed to skill complementarity and cost. Overall, the results indicate that while TBTC is crucial for employment distribution, skill complementarity is more influential for wage distribution and also in shaping firms' endogenous skill specialization.

The Role of Education: In my model, although direct education investment is not explicitly included, the changing skill supply ($G(\mathbf{x})$) implies at the potential role of education in shaping labor market outcomes under skill mixing. As depicted in Figure 2.2, the calibrated skill supply variation based on occupational and major choices indicates an increase in the mixing of analytical/computer and interpersonal skill, while a rise in specialization in routine skill. While skill supply has not played a significant role in wage and employment distributions, it has reduced the returns to specialization in different skills by 0.5 to 3 percent. These results suggest that while education may have marginal effects on overall distributions, it significantly impacts the wage disparity between experts in specific skills and those who are not.

CHAPTER 3

E-COMMERCE AND REGIONAL INEQUALITY: A TRADE FRAMEWORK AND EVIDENCE FROM AMAZON'S EXPANSION

3.1 Introduction

As e-commerce is transforming the retail sector, regions across the United States face very different prospects. While a town in New Jersey might see expanding warehouses and manufacturers, another town in Wyoming may mostly suffer from the collapse of local brick-and-motor stores. Studies that examine the impact of e-commerce have noted its impact on the demand, productivity, and markup of physical retail stores (i.e., Goldmanis et al. 2010; Pozzi 2013; Ellison and Ellison 2018), as well as on consumer welfare (Fan et al. 2018; Dolfen et al. 2019). However, little work has thoroughly examined the regional inequality and redistribution effects of e-commerce in terms of economic activities and job opportunities. As the divergence in regional economies has key implications for life outcomes (see Chetty and Hendren 2018; Austin, Glaeser, and Summers 2018), understanding the consequences of e-commerce on regional inequality is important for policy making.

In this paper, I adopt a trade perspective to study e-commerce's impact on different local labor markets, taking into account trade and input-output linkages and regions' comparative advantages. A key feature of e-commerce is that online retailers don't have to be where the customers are, therefore having more mobility in their location. As in Krugman (1991) and Krugman and Venables (1995), the additional mobility will induce agglomeration in the online retail sector. In an environment where online retailers are the intermediary between the upstream producers and downstream consumers, online retailers would want to locate near the largest consumer or the cheapest producer, but also need to take into account the resulting rise in wages and land prices. The intermediary nature and agglomeration of online retailers will imply greater specialization in both the upstream and the online retail sectors.

Using a comprehensive panel dataset of products and retailers on Amazon, as well as Amazon's fulfillment and distribution facilities, I document four stylized facts that suggest online retailers are more agglomerated in space and their agglomeration is associated with greater trade flows of the upstream goods. First, online retail sales are more spatially concentrated than overall retail sector sales, and are and are less correlated with population and more correlated with manufacturing output; second, online retailers that use Amazon's fulfillment and distribution facilities more agglomerated than those that don't use the facilities; third, destination markets with more online retailers import more wholesale trade goods, whereas origin markets with more online retailers export less wholesale trade goods; fourth, regions near to Amazon's fulfillment and distribution facilities import and export less wholesale trade goods.

Taking these key features of online retailing into account, I build a multi-sector spatial trade framework of intra-regional retailing to analyze e-commerce's impact. The role of e-commerce is first reflected in that consumers have to conduct costly simultaneous search and matching of retailers as in Weitzman (1979), the efficiency of which is subject to the online retail platform. Moreover, I allow online retailers to optimally choose their locations where they import from the upstream sectors and ship to consumers, giving rise to agglomeration incentives. To better understand the impact on employment, I also let workers be heterogeneous in their productivity and optimally choose the sector of employment or to be unemployed. I show that despite the rich micro-foundation, this framework can still aggregate to a gravity trade model with CES demand, with the demand shifter reflecting online match efficiency and the iceberg cost influenced by the shipping cost of online retailers. The location probability of online retailer in a region directly scale up the gravity of trade flows in that region, highlighting the important role of online retailer agglomeration in the model.

I then estimate key fundamentals to take the model to the data, particularly the reduction in shipping friction and the increase in match efficiency related to the rise of e-commerce. I apply the datasets I obtained on Amazon retailers and sales, as well as Amazon facilities to conduct the estimation. The major challenge in identifying the impact of Amazon's expansion concerns its endogeneity to other factors, particularly from the demand side. To overcome this issue, I employ a spatial simulated instrumental variable strategy (Duflo and Pande 2007; Lipscomb, Mobarak, and Barham 2013; Faber 2014). Instead of using the actual location of Amazon's facilities to calibrate the shock, I build counterfactual distribution centers with the simulated location choices based solely on plausibly exogenous geographic and climatic factors. The shipping cost reduction due to these counterfactual centers is used to instrument the actual decline of shipping frictions and iceberg costs. Conditional on the estimated reduction in iceberg cost, the predicted changes in regional online retail expenditures identify the increase in online match efficiency. My estimation results show that Amazon's growth has led to a 3 percent decline in iceberg cost and a 29 percent increase in online matching efficiency from 2007 to 2017.

Equipped with the estimated shocks and calibrated model parameters, I evaluate Amazon's impact on regional economies in terms of total welfare and employment. I find that Amazon's growth in this period has led to a positive effect on total welfare due to the associated price decline, but meanwhile reallocation of workers out of the manufacturing sector, decreasing income. Taking these two forces, welfare has declined by 1 percent on average, but underlying this overall effect is huge regional dispersion. States with an initially small share of online retail consumption (Wyoming, South Dekoda) and states with a bigger market and diversified industrial composition (California, Washington) enjoy a welfare surplus, while middle-eastern states (Indiana, North Carolina) bear welfare losses. The non-employment rate has increased by 2.3 percentage points; in the meantime, the Gini index of non-employment also increased from 0.11 to 0.13 (20 percent), implying growing dispersion in employment opportunities in different regions.

The likely widening of gaps in economic outcomes across regions due to the rise of e-commerce as represented by Amazon creates a strong rationale for national-level policy interventions. To compensate for the growing trade imbalances across regions, leaving it to the local governments, they might impose domestic "tariffs" on non-local goods, which recover the first-best allocation, but also create welfare losses for consumers (Costinot et al. 2015; Antràs et al. 2022). Due to the spatial nature of the market failure, there is a need for a national-level revenue reallocation. Moreover, since the key aspect of e-commerce shock works through match and shipping friction, the government might directly intervene in the online retail market design. I will conduct counterfactual analyses with these policy experiments in the next step.

The rest of the paper is organized as follows. The ensuing section reviews the relevant literature in more detail and highlights this paper's contributions. Section 3.3 presents the stylized facts on the online retailers and associated trade flow. Section 3.4 presents the theoretical framework and how to use it to conduct comparative statics and welfare analysis. I discuss model quantification in section 3.5 and the estimation of Amazon shock. Section 3.6 shows the results on the impacts of Amazon.

3.2 Literature Review

The rise of e-commerce presents a salient case where technology progress redistributes economic opportunities not only across sectors, but also across spaces. This paper propose using a trade framework to study e-commerce, and particularly highlighting the agglomeration of online retailers. It contributes to the literature by applying and extending a standard trade framework to study the spatial general equilibrium effects of e-commerce with new data and identification strategy. Specifically, this paper closely relates to four strands of literature.

Firstly and most relevant, this work builds on the literature studying the market structure of the retail sector and the impact of e-commerce. Two important findings emerge from this literature. For the retail industry, it is found that e-commerce reduces the demand of the physical department stores, raising their productivity but reducing the mark-up in the consumer goods sector (Stanchi 2019;Goldmanis et al. 2010). This supports the modeling of e-commerce as a productivity shock to the retail sector as adopted in this paper. For consumers, Dolfen et al. (2019) finds that e-commerce increases consumer welfare mainly through substituting to online merchants. Fan et al. (2018) shows e-commerce increases domestic trade and benefit consumers in smaller cities and markets particularly. This paper instead studies e-commerce from a general equilibrium spatial trade framework and focuses on its impact on employment and GDP growth differentials across regions. In the welfare analysis I take into account the consumption channel and evaluate the trade-offs.

The theoretical framework of this paper builds on the large literature on of international trade and spatial equilibrium models, and presents a novel application of these theories to study e-commerce. In particular, I adopt the analogy to "globalization" and model e-commerce as a trade shock; for the geographic implications, I apply a Ricardian trade framework focusing on intra-regional and sectoral reallocation taking into account comparative advantages of localities for labor market outcomes (Caliendo et al. 2018; Caliendo, Dvorkin, and Parro 2019; Lee 2020; Adao, Arkolakis, and Esposito 2019). Theoretically, I add into a typical Eaton and Kortum (2002) framework with information frictions, transportation cost and worker sorting to more accurately depict the retail sector, as well as roles played by local and federal governments to discuss policy implications. Empirically, I use Amazon's expansion as the source of variation and present new estimation strategy that introduces simulated IV into a typical Bartick estimator.

This paper also relates to studies about the differential impact of technological changes on workers. The earlier discussion in this literature focuses on the wage premium for higher-skill workers, or "skill-biased technological change" (Autor, Katz, and Krueger 1998; Acemoglu 2007). It is also found that starting from 1980s, workers conducting "routine" tasks are more likely to be substituted, leading to the polarization of the labor market (Autor, Levy, and Murnane 2003; Acemoglu and Autor 2011; Autor and Dorn 2013). This paper contributes to this literature by focusing the spatial nature of a technological change (e-commerce) that has the feature of both an automation and a trade shock, and analyzes under a full general equilibrium spatial trade framework of its impact on workers across sectors and regions.

Lastly, this paper speaks to the literature that examines the differential economic opportunities across spaces. Kline and Moretti (2013) and Amior and Manning (2018) show that there is strong persistence of unemployment and labor force participation differences across regions; Amior and Manning (2018) argues that this is mainly due to the long adjustment to persistent local labor demand shocks. Also relevant is the large literature revealing the importance of neighborhood quality differences on one's life outcomes and hence place-based policies (i.e., Chetty, Hendren, and Katz 2016). Here I analyze a particular technology-induced local labor demand shock (e-commerce) that has strong spatial redistributive effects and explores place-based taxes and subsidies. A related literature on the mismatch between workers and jobs found that mismatch across industry and three-digit occupations could explain up to a third of the increase in unemployment (Şahin et al. 2014). This paper investigates a particular cause of mismatch from the labor demand side at the intersection of location, industry and occupation, and explores the tax policy implications.

3.3 Evidence on Online Retail Sales

In this section, I show that the empirical data patterns are consistent with the agglomeration of online retailers and corresponding trade shock. I first lay out

Figure 3.1: The Online Retail Business Model

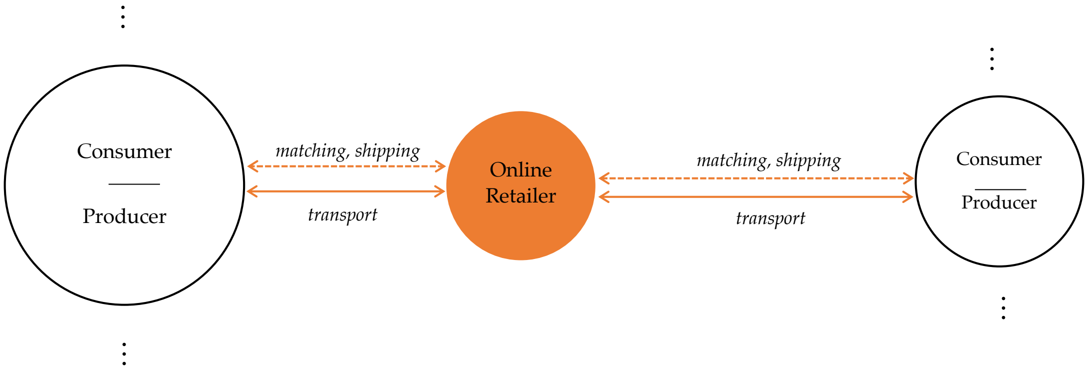

how online retailers engage in e-commerce based on direct and indirect industrial evidence. These observations generates implications for online sellers' locations choices, agglomeration, and for intra-regional trade flows. I then introduce the specific data regarding online sellers, products, as well as intra-regional trade. Finally, I conduct empirical analysis to test the implications.

3.3.1 The Online Retail Business Model

A distinguishing feature of conducting e-commerce relative to conventional retail sales is the decoupling of retailer and consumer location. Retailers don't have to physically present where the consumers are to sell their goods, and instead, they engage in online match making with consumers through online platforms. Figure 3.1 shows the e-commerce business model of a typical online retailer. Different from brick-and-mortar retailing where the consumers need to commute to the store, in e-commerce, consumers obtain their goods either directly from the online retailer or from the storage the retailer has in the fulfillment center, both incurring a shipping cost. Nonetheless, e-commerce share one common feature with brick-and-mortar retail: the retailers has to buy goods from producers in the wholesale market, and incur transport cost there.

The key assumption for the e-commerce business model in Figure 3.1 is that online retailers first purchase the goods and place in their locations, before shipping to consumers, either directly or through third-party fulfillment service. Despite that the academic literature has little to say about the shipping modes of online sellers, in the Amazon data that I obtained, 72 percent of Amazon sellers and 78 percent of products sold use the Amazon fulfillment service, implying the use of direct shipping from producer to consumer is not a huge part of the sample. Some indirect evidence, such as case studies of Amazon sellers also indicate that these sellers' physical location act mainly as inventory storage, acting as the relaying point between producer and consumer.

Implications: The greater flexibility of online retailers' locations creates strong incentives of agglomeration. As in Krugman (1991); Krugman and Venables (1995) and Puga (1999), the presence of both spatial frictions and input-output linkages creates pecuniary spillovers of co-location. Specifically in the setting of the e-commerce business model, faced with matching frictions and shipping cost in the downstream, as well as transport cost in the upstream, online retailers would want to locate closer to either their major consumer or producer to save the costs, the decision of which depends on the relative cost magnitude on the two sides.

Such an insight makes it clear how the drastic expansion of e-commerce affect the economy by altering the location motives of online retailers. A key feature of e-commerce platform expansion (i.e., Amazon) post-2005 is its improvement of online shopping experience and rolling out of fulfillment centers. These changes reduces the matching friction and shipping cost to the downstream consumers. While online sellers' transport cost with upstream producer is not reduced as fast.[1] Such asymmetric changes in the spatial friction should motivate online retailers to locate more agglomerated in space to major producers, and the agglomeration is likely to be stronger when the online retailer has better access to fulfillment centers, since the shipping burden would be reduced more.

The potential agglomeration of online retailers will also alter the trade flows across regions. As online retailers serve as the intermediary of selling upstream producers to the downstream consumers, their agglomeration in a region will direct more purchases of the upstream goods in that region. If the region happens to the destination market where consumers are, there will be more imports of the upstream goods into the region; if on the contrary, the region happens to be the origin market, there will be less exports since the online retailers source locally. In testing these implications, I will use intra-regional wholesale trade data from CFS to check the purchase of upstream goods by online retailers. Under a similar vein, as a region gains better access to fulfillment facilities, it relaxes the burden of online retailer to the destination market, and is likely to be associated with reductions in wholesale trade flows.

3.3.2 Data

Products and Sellers on Amazon: The major data I used to test for the empirical implications and later quantitatively evaluate the model comes from Keepa

---

[1] The transport cost with upstream producers could also reduce in this period, due to general improvement in infrastructure and transportation and information technology. What's the key in driving the result, however, is the asymmetric changes in frictions, due to Amazon's more expansive presence on the downstream side.

(`www.Keepa.com`), an online marketing intelligence firm that serves both Amazon buyers and sellers by providing detailed information on products and sellers. Keepa started collecting data Amazon since 2011; once a product is searched by a consumer, Keepa will track it in its database. Therefore, Keepa's database includes any products that have ever been looked at by consumers, and is updated on a daily or weekly basis depending on the information. As of January 2023, Keepa's database includes more than 674 million products of 36 root categories sold on Amazon in the United States. For the purpose of my analysis, I took a 1 % random sample out of each category and restrict to the period 2016-2018, which is after Amazon pick-up of e-commerce's expansion. Online Appendix Table **??** details the number of products of each category included in the analysis of this paper.

The product data I collect from Keepa contains each product's root category and brand, as well as longitudinal information such as prices, sales rank, and ratings. Several studies in the marketing literature show that a Pareto distribution fits the sales rank and quantity relationship well over e-commerce platforms. Using a combination of a book publisher's data and authors' own experiment, Chevalier and Goolsbee (2003) found that the coefficient of a regression of log sales quantities on log rankings to be around -0.76 to -1.11, while using the online sales data of 734 products of a retailer, Brynjolfsson, Hu, and Simester (2011) found the coefficient to be -0.88. Therefore, I convert the sales rank into quantity sold by running a similar regression and adopt an coefficient of -0.9.[2] Together with price information I then obtain the total sales revenue of a product overtime.

---

[2] What will also be important for the imputation is the intercept of the regression, since different product categories might have different innate level of sales quantity, despite the Pareto distribution fits well the quantity-rank relationship. I adopt Brynjolfsson, Hu, and Simester (2011)'s estimated intercept of 8.13 since their data cover broader product categories

Moreover, I also obtain detailed seller information for the products sold on Amazon. Keepa starts to track sellers in 2016 and assign each seller with a unique identifier, which can then be linked to the seller profile on Amazon that contain information on the seller′s address, fulfillment method, and whether the seller ships products from China. I retain all sellers that are located within the United States and that do not directly ship from China. Since a product can be available from multiple sellers at each point in time, I assign the seller of product to be the one that appear in the "BuyBox" , which accounts for more than 80% of sales of a product.[3]

Amazon Facilities: I obtain information on Amazon's fulfillment and distribution facilities from the supply-chain consulting firm MWPVL (`www.mwpvl.com`). The provided data contains the specific year and location a facility is built, its square footage, and detailed description of its functionality. For the purpose of my analysis, I focus on relatively bigger fulfillment and distribution centers that handle the common-sized domestic orders of non-perishable goods in typical regions. These are the facilities that most likely will lead to a decrease in shipping cost and therefore, consumers' shopping patterns and sellers' locations decisions.[4]

Figure 3.2 illustrates the capacity changes of Amazon's fulfillment and distribution facilities from 2000-2020. Panel (1) and (2) illustrate that there is a huge increase in center sizes from 2010-2020, with the majority of facilities built in this

---

[3] BuyBox is the "Add to Cart" and "Buy Now" section of the product detail page. Winners of the BuyBox are determined by Amazon algorithm that takes into account the price, product rating, delivery method of the sellers.

[4] Amazon also runs other specific centers that deal with fresh food and orders placed through Prime Now or Whole Foods, as well as centers that deal with in-bound goods and located near the airports, or deal with small packages; these facilities are excluded from my analysis. Within the fulfillment and distribution category, I don't differentiate whether the center is serving more in terms of storage or sortation, as both reduces the shipping time and cost.

Figure 3.2: Expansion of Amazon Facilities

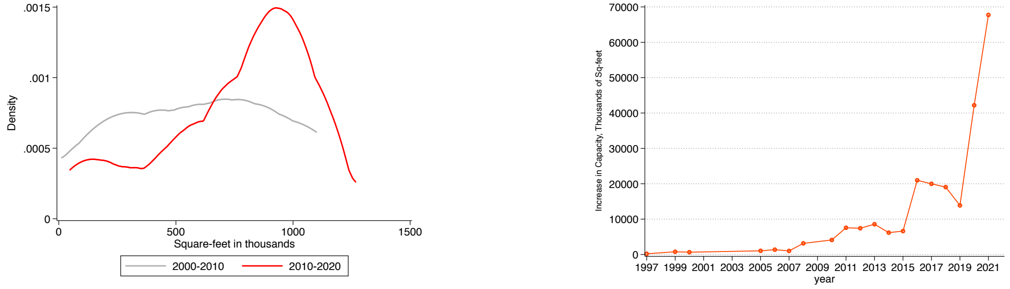

(1) Distribution of Facility Sizes

(2) Changes in Facility Capacity

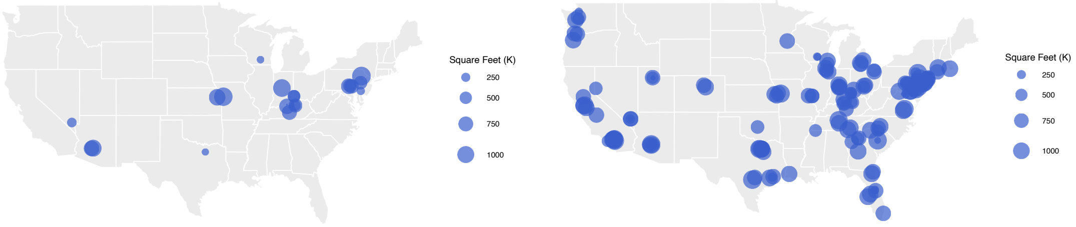

(3) Amazon Facilities Before 2010

(4) Amazon Facilities Before 2010-2020

period at around 1 million square feet. Since 2015, there is a huge upsurge of 14-20 million square feet per year, leading to the height of 42 square feet built in year 2020. Panel (3) and (4) maps the locations of the centers using geo-coded address. From 2000-2010, most centers are concentrated in 3 states: New York, Kentucky, and Arizona. The geographical distribution of centers spread substantially starting 2010, covering most U.S. states with concentrations in the east and west coast.

Trade Flow. For the data patterns regarding intra-regional trade flows, I rely on Commodity Flow Survey (CFS) that provides representative shipment level trade flows in value and quantity for all the 30 manufacturing and retail sectors across 50 U.S. states.

Figure 3.3: Spatial Concentration of Online vs. Overall Retail Sales and Sellers

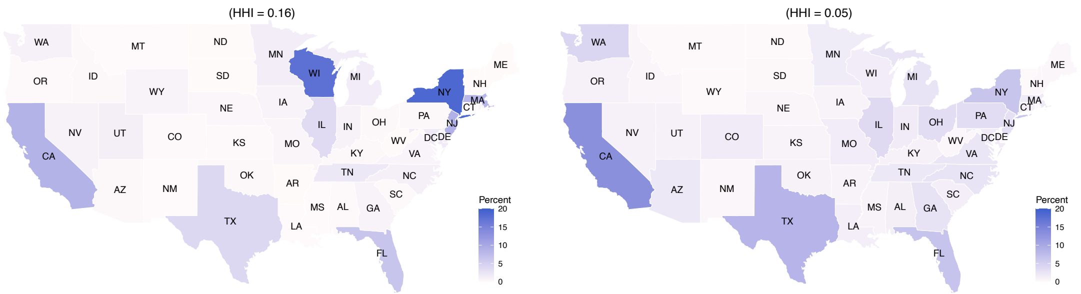

(1) Regional Share of E-commerce Sales

(2) Regional Share of Retail Sector Sales

(3) Regional Share of E-commerce Sales with FBA

(4) Regional Share of E-commerce Sales without FBA

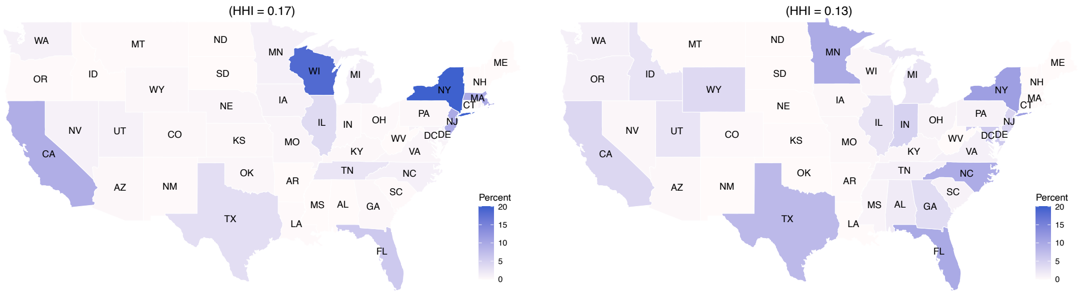

3.3.3 Data Patterns

In this section, I document four broad data patterns that point to the differential concentration of online sellers and implications for intra-regional trade flows.

Pattern 1a: *Online retail sales are more spatially concentrated than overall retail sector sales, particularly for those that are FBA.*

Figure 3.3 panel (1) and (2) depict different states' shares of total online retail sales on Amazon as well as of overall retail sector value-added, and clearly indicates that online retail sales are more spatially concentrated.[5] I assign online sellers' sales value to different states based on the sellers' addresses, and I obtain from BEA states' shares of retail sector value-added. The results indicates that

---

[5] States' shares of retail value added are good proxies for their shares of retail sales if the retail production function is Cobb-Douglas with constant factor shares across regions.

Table 3.1: HHI Index by Product Categories

Category name	**HHI Index**
Toys & Games	0.12
Patio, Lawn & Garden	0.12
Arts, Crafts & Sewing	0.07
Sports & Outdoors	0.14
Office Products	0.16
Grocery & Gourmet Food	0.08
Tools & Home Improvement	0.21
Movies & TV	0.08
Musical Instruments	0.10

two states—New York, and Wisconsin—have captured 36 percent of total online retail sales, while for the overall retail sector value-added, states' shares are more consistent with their population sizes. The value of Herfindahl-Hirschman Index (HHI) is 0.16 for online retail sales, and 0.05 for retail sector sales, confirming the greater concentration of the former.

The fulfillment services that Amazon provides ease online retailers' burden of shipping and could lead to greater agglomeration. Figure 3.3 panel (3) and (4) depict the states' shares of online retail sales that use Amazon's FBA service versus those don't use. Online retail sales through FBA is more spatially concentrated, and drives the overall concentration of online retail sales, with a higher HHI (0.17 versus 0.13).

Pattern 2: *Online retail sales of durable and standardized products are more concentrated than those of non-durable and non-standardized products.*

The agglomeration patterns of online sellers, as depicted in Figure 3.3, vary by product groups. Standardized and durable products, which benefit from economies of scale, predictable demand, and optimized transportation and storage costs, tend to have sellers concentrated in fewer locations. In contrast, prod-

Table 3.2: Online and Total Retail Sales with Population and Corporate Taxes

Dependent Variable (Log Sales)	Online Retail (Non-FBA)		Online Retail (FBA)		Overall Retail	
Corporate tax	-0.88	1.41***	0.29	0.92	-0.81	1.07***
	[0.55]	[0.52]	[0.83]	[0.81]	[1.90]	[0.28]
% Total population	3.05	-9.05	0.15***	15.63**	-0.02	0.96***
	[2.43]	[15.51]	[0.05]	[7.66]	[0.02]	[0.05]
Year, State FE		X		X		X
Observations	230	230	230	230	230	230
R-squared	0.11	0.36	0.13	0.55	0.99	1.00

Figure 3.4: Peak Period Congestion on the High-Volume Truck Routes in 2020

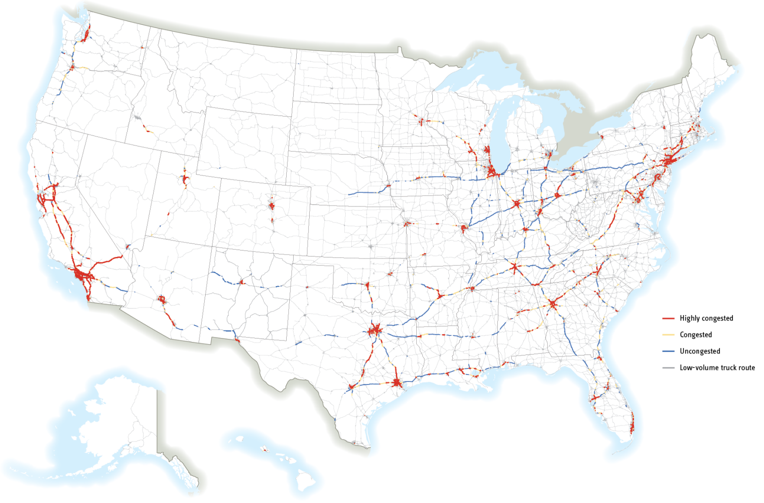

ucts that are non-durable and non-standardized typically have a more dispersed seller distribution. Table 3.1 displays the Herfindahl-Hirschman Index (HHI) of sales across regions for nine commonly purchased goods on Amazon. Notably, durable goods such as ”Tools & Home Improvement” and ”Office Products” exhibit high spatial concentration, with HHI indices of 0.21 and 0.16, respectively. In contrast, non-durable goods like ”Arts, Crafts & Sewing” and ”Grocery & Gourmet Food” show lower concentration, with HHI indices of 0.07 and 0.08, respectively.

Pattern 3: *Online retail sales are less explained by population or taxes than overall retail sector sales, and more explained by truck volumes.*

To obtain a clearer understanding of how online retail sales differ from the overall retail sector in terms of agglomeration patterns, Table 3.2 presents the results from a regression analysis of the relationship between states' shares of online retail sales and retail sector value-added with their percentages of population and corporate tax revenues over time. Generally, the population positively correlates with the regional sales shares of FBA online retailers and the entire retail sector, while corporate taxes have a positive association with the sales shares of non-FBA sellers and the entire retail sector. However, it is crucial to recognize that the R-squared values for the regressions involving both types of online retailers are significantly lower than those for the overall retail sector. This discrepancy suggests that population and taxes explain a larger variation in overall retail sales activities than in online retail sales.

Moreover, the regional sales shares in the overall online retail sector are more closely aligned with regional truck volumes. Figure 3.4 illustrates the peak period congestion on high-volume truck routes and highlights that states such as Wisconsin, Illinois, New York, Texas, Florida, and California experience the highest truck volume congestion. These same states also have the highest concentration of online retailers. This correlation may be attributed to the advanced transportation and logistics infrastructure and services available in these areas, contributing to the agglomeration of online sellers in these regions.

Pattern 4: *Destination markets with more online retailers import more wholesale trade goods, whereas origin markets with more online retailers export less wholesale trade*

Table 3.3: Wholesale Import/Export, Online Retailers, and Amazon Facility

Dependent Variable:	ln(Shipment)	
Share (%) of online sellers - destination	1.5*	
	[0.8]	
Share (%) of online sellers - origin	-3.7***	
	[1.0]	
Bilateral distance via Amazon facility		-0.20**
		[0.08]
Origin, Destination FE	✓	
Year, Industry FE	✓	
Observations	19,739	43,715
R-squared	0.2	0.4

goods.

Table 3.3 column (1) illustrates the association between online retailer agglomeration and the upstream wholesale trade flows. Since the Keepa data on online retailer location is available only after 2016, but CFS conducts survey every 5 years with the most recent one in 2017, I regress the changes in intra-regional trade flows between 2012-2017 on states' average share of online retailers between 2016-2017, controlling origin, destination, and industry fixed characteristics. Consistent with the predictions, a one percent increase in a state's share of online retailers is associated with a 1.5 percent increase in wholesale shipment if that state is the destination market, but is associated with a 3.7 percent drop in wholesale shipment if that state is the origin.

Pattern 5: *Regions near to Amazon's fulfillment facilities import and export less wholesale trade goods.*

As the fulfillment service that Amazon provides ease online retailers' burden of being closer to either the downstream consumer or upstream producer, the potential loss of online retailers will likely reduce the trade flow. Table 3.3 column

(2) shows a regression of log of shipment value of wholesale trade goods between an organ-destination pair on the distance between the pair through the nearest amazon fulfillment center that is likely to handle the shipment, controlling fixed origin, destination, industry characteristics, as well as time trend. To compute the distance to the nearest distance via an amazon facility, I follow Houde, Newberry, and Seim (2021), which shows more than 90 percent of ordered are handled by the 3 closest fulfillment centers to destination, and assign among the 3 closest centers, the one that is also closest to the origin to be the one handle the fulfillment. The result indicates that a one percent decrease in bilateral distance due to the expansion of amazon's fulfillment centers is associated with a 0.2 percent decrease in bilateral shipment of wholesale trade goods.

3.4 A Spatial Retail Trade Model

In this section, I build a multi-sector spatial retail trade model in which consumers search for retailers across regions, and online retailer choose their optimal location to import from the upstream sector and to ship to consumers. The model illustrates how online retailing, as embodied in the increase in match efficiency and reduction in shipping frictions affects different regions by altering the trade flows. It also shows that the location choices of online retailers plays an important role in determining these trade flows. The model provides a quantitative tool to evaluate the impact of e-commerce calibrated using data, which I will show in next section.

The general environment of the model contains N regions denoted by n or m, and J sectors denoted by j or k. For each region-sector pair, there is a representative manufacturer (M) and a brick-and-mortar store (B), as well as a flexible measure of online retailers (R). I show that consumers' search and shopping problem simplify to a CES demand with a demand shifter of online retailers representing the efficiency of matching. The production intermediate varieties follows multi-sector Eaton and Kortum (2002), for which regions with comparative advantages obtain higher share of the market demand subject to transport frictions. These varieties are then purchased by brick-and-mortar as well as online retailers. The distinguishing feature of the model is that online retailers choose the locations that will give them a cost advantage in terms of both purchasing and selling to multiple markets. Workers are heterogeneous in their productivities and optimally determine the sector of employment. In what follows I describe the spatial retail trade problem, the role played by e-commerce and then derive the comparative statistics.

3.4.1 Demand

In this section I show the sequential search problem of retail purchasing that consumer faces, sin which online platforms such as Amazon plays the role of match making gives rise naturally to an equilibrium demand of the CES form, for which the online matching efficiency is represented by a demand shifter.

Consumer Search: There is a continuum of consumers of region n, each consuming goods of different sectors with weights η_j. For retail sector goods, they purchase it from one retailer i among $O+1$ retailers, where $i = 0$ represents the local brick-and-mortar store, and there are O totally online retailers indexed by $i \geq 1$. Consumers have Cobb-Douglas utility over the sectoral goods $u_n = \sum_j \eta^j \ln c_{ni}^j$, and with income y_n, their optimal consumption from the chosen retailer becomes $c_i^{j*} = \eta^j y_n / p^j_{ni}$, where p^j_{ni} is the price that retailer i charges for sector j goods that includes the cost of obtaining the goods, such as commuting cost from the local brick-and-mortar store, or shipping cost from online retailers.

Consumers have imperfect knowledge about the goods, and to resolve the uncertainty, they need to search for the optimal goods. Specifically, a consumer's indirect utility of consuming sector j good from retailer i can be expressed as $v^j_{ni} = \ln \eta^j y_n - \ln p^j_{ni} + \epsilon^j_{ni}$. Here, ϵ^j_{ni} represents the idiosyncratic match value between a consumer and retailer pair and is assumed to be independently distributed according to $F(\epsilon)$ unknown to the pair. I normalize ϵ^j_{ni} such that consumers' match value is 0 with the local brick-and-mortar store ($\epsilon^j_{n0} = 0$) and the relative average match value ϵ^j_{ni} consumers have with online retailers is $\ln(\mu)$.[6] Consumers can either purchase from the local store as an outside option, or search sequentially for online retailers. If one chose the latter, at each step, the consumer decides whether to pay a cost s to observe ϵ^j_{ni} of a online retailer.

Optimal Demand: As the directed sequential search problem here represents that of Weitzman (1979), the optimal strategy for the consumers is to order their search of the retailers by $\epsilon^{j*}_{ni} - p^j_{ni}$, in which ϵ^{j*}_{ni} is the lowest match value that makes consumers indifferent between searching or not ($s = \int_{\epsilon^j_{ni}}^{\bar{\epsilon}^{j*}_{ni}} (1 - F(\epsilon)) d\epsilon$).[7] Several studies found that the search outcome can be simplified even further (Choi, Dai, and Kim 2018; Armstrong 2017; Armstrong and Vickers 2015). Let $\omega^j_{ni} \equiv min\{\epsilon^j_{ni}, \epsilon^{j*}_{ni}\}$, which stands for the "effective match value" of a retailer, the

---

[6] The relative match value $\ln(\mu)$ contrasts the shopping experience between the two modes. Taking log gives a cleaner representation and is without loss of generality. If $\mu > 1$, consumers obtain higher utility from shopping online, and vice versa for $\mu < 1$.

[7] The consumer will stop and purchase from either the local brick-and-mortar store or an online retailer i if $max\{v^j_{n0}, -max_{g \in \bar{O}} \ln p^j_{ni} + \epsilon^j_{ni}\} > max_{g \in \bar{O}} - \ln p^j_{ng} + \epsilon^j_{ng}$, where $\bar{O}$ stands for the retailers the consumer has checked so far.

consumer will buy from the retailer $i = argmax_i \; \omega^j_{ni} - p^j_{ni}$.[8]

$$D^j_{ni} = P(\omega^j_{ni} - \ln p^j_{ni} > max_g \; \omega^j_{ng} - \ln p^j_{ng}) = \int \Pi_{g \neq i} F_{\omega^j_{ng}}(\epsilon - \ln p^j_{ng}) f_{\omega^j_{ni}}(\epsilon - \ln p^j_{ni}) d\epsilon \tag{3.1}$$

The characterization of consumers' eventual purchase based on ω^j_{ni} allows a discrete choice formulation of the optimal demand (Anderson, Engers, and Savelle (2022)). Specifically, the representative consumer's demand of region n for retailer i of sector j goods can be expressed as in the equation of D^j_{ni} above, which is equivalent to a discrete choice model if $F_{\omega^j_{ni}} = F_{\epsilon^{j,DC}_{ni}}$, where $\epsilon^{j,DC}_{ni}$ is the random utility a consumer obtains from the retailer. The discrete choice formulation leads to a more frequently used, CES representation of consumer's demand, as CES can be considered as a special case of demand based on discrete choice. Specifically, since we know that the average ϵ^j_{ni} is 0 for brick-and-mortar stores, and $\ln(\mu)$ for online retailers, we can express $\epsilon^{j,DC}_{ni} = \ln(\mu) + \chi^j \tilde{\epsilon}^j_{ni}$ such that $\tilde{\epsilon}^j_{ni}$ has mean 0 and unit variance, and χ^j is the variance of the effective match value ω^j_{ni} that is assumed to differ across sectors but not regions.[9]

Theorem 1 gives the final CES demand of consumers. Under the assumption of extreme type I distribution of $\tilde{\epsilon}^j_{ni}$, region n's consumer demand for sector j goods from retail i becomes $D^j_{ni} = \frac{(p^j_{ni}/\mu)^{\frac{-1}{\chi^j}}}{\sum^N_{g=1}(p^j_{ng}/\mu)^{\frac{-1}{\chi^j}} + (p^j_{n0}/)^{\frac{-1}{\chi^j}}}$, a standard CES expenditure share. This expression clarifies the role of μ–as online shopping and matching becomes more efficient, μ increases, and consumers shift their demand more

---

[8] As Choi, Dai, and Kim (2018) shows, to guarantee the existence and uniqueness of the equilibrium, one needs the density and loss function sof ω^j_{ni} to be log-concave, and the density function to be unbounded above. These are taken as assumptions for this paper.

[9] This requires that the effective match value ω^j_{ni} will also have mean $\ln(\mu)$ and there is a large number online retailers relative to the local retail store. Since it is likely that the cost of searching for an additional retailer s is relatively small via online platforms, and since s is decreasing in ϵ^{j*}_{ni}, ω^j_{ni} is closer to ϵ^j_{ni} that has mean $\ln(\mu)$.

towards online retailers with online retailers. The variance of consumers' effective match value χ^j determines the elasticity of substitution among retailers $\sigma = \frac{1+\chi^j}{\chi^j}$: as there is less uncertainty about the value of the retail goods, retailers become more substitutable. Since each retailer sells its own variety, monopolistic competition implies that the mark-up retailers charge is $\tilde{\sigma}^j = \frac{\sigma^j}{\sigma^j - 1}$.

Theorem 1 *A representative consumer in region n with sectoral consumption weights* η^j *has nest Cobb-Douglas and CES demand as below under sequential ordered search if only if the effective match value* $\omega^j_{ni} = min\{\epsilon^j_{ni}, \epsilon^{j*}_{ni}\}$ *is distributed extreme type I*

$$C_n = \Pi^J_{j=1}(C^j_n)^{\eta^j}, \quad C^j_n = [(c_{n0})^{\frac{\sigma-1}{\sigma}} + \mu\sum_{i=1}^{N}(c_{ni})^{\frac{\sigma-1}{\sigma}}]^{\frac{\sigma_j}{\sigma_j - 1}} \textit{ for } j \geq 2 \tag{3.2}$$

Proof: See Appendix C.1.

3.4.2 Production

Production is multi-stage and vertical to capture the role of retailers. Both brick-and-mortar and online retailers first collect manufactured intermediate varieties across regions, and then turn them into final retail goods and charge a mark-up. Therefore, there are two layers of intra-regional trade in this framework: for each sector, trade happens in both the final and intermediate goods market. The difference between the two types of retailers is that brick-and-mortar can only serve the local consumers, whereas online retailers can sell to all the regions, and will choose the location optimally taking into account of trade costs. The location choices of online retailers then determine the intra-regional trade flows of retail goods.

Intermediate Varieties. There is a representative firm in each sector j of region n that produces a continuum of varieties $e^j \in [0, 1]$:

$$q_n^{j,M}(e^j) = a_n(e^j)\left[h_n(e^j)^{\beta_n} l_n(e^j)^{1-\beta_n}\right],$$

where $a_n(e^j)$ is the factor neutral productivity to produce variety e^j by the firm in region n, and $l_n(e^j)$ and $h_n(e^j)$ are labor and land or structures used. The production features labor and structures as complements, bundled together in a Cobb–Douglas function with their shares controlled by β_n. All firms across different regions have access to the this same technology, and with it being constant return to scale, no firm has any market power. The prices are set at the unit cost given by equation (3.3), where r_n^h and w_n^j are structure costs and wages respectively.

$$c_n^{j,M} = \left[(\frac{r_n^h}{\beta_n})^{\beta_n}(\frac{w_n^j}{1-\beta_n})^{1-\beta_n}\right] \tag{3.3}$$

The trade of intermediate varieties is subject to standard iceberg cost that requires κ_{ni}^M units of good for one unit of it to ship from i to n. Interpreting this cost as related to transportation expenses proportional to distance leads to the requirement that $\kappa_{ni}^M > 1$ for $i \neq n$ and $\kappa_{ni}^M = \kappa_{in}^M$. With the market structure of intermediate goods being perfectly competitive and consider a vector of draw of productivities across regions $a(e^j) = \{a_1(e^j), \ldots, a_N(e^j)\}$, the price of variety e^j in region n is the lowest of the effective unit cost multiplied by the iceberg cost:

$$p_{ni}^{j,M}(a(e^j)) = \min_i\{\kappa_{ni}^M \frac{c_n^{j,M}}{a_i(e^j)}\}.$$

By further parameterizing the probability structure of productivities as Eaton and Kortum (2002), one can obtain a gravity representation of trade across regions. Specifically, let $a_n(e^j)$ be a random draw from a Fréchet distribution with a shape and scale parameter given by θ^j and T_n^j respectively: $\phi_n^j(a_n(e^j)) = exp(-T_n^j z^{-\theta^j})$. The Fréchet shape θ^j determines the dispersion of productivities across regions and the within-sector specialization pattern, while T_n^j regulates regions' absolute advantages in production and across-sector specialization. Using properties of the Fréchet distribution, expenditure share of region n on i in sector j of intermediate goods $x_{ni}^{M,j} = X_{ni}^{M,j}/X_n^{M,j}$ can be expressed as a gravity formula:

$$x_{ni}^{j,M} = \frac{(\kappa_{ni}^M c_i^{j,M})^{-\theta^j} T_i^j}{\sum_{m=1}^N (\kappa_{nm}^M c_m^{j,M})^{-\theta^j} T_m^j}, \tag{3.4}$$

which represents the probability that for varieties in sector j, buyers in n purchase from low cost vendors from i. Note that this probability depends on both the unit cost of the product and iceberg cost between the two regions, therefore, θ^j can be interpreted as the trade elasticity with respect to cost. With a higher θ^j, the dispersion of productivities across regions is lower and import volumes become more responsie to cost changes.

Retail Sector. Both brick-and-mortar and online retailers of a given region and sector first collects different intermediate varieties $e^j \in [0, 1]$ from the lowest-cost producers and aggregate them into a retail bundle $q_n^{J,R/B}$ for the production of retail good, as shown in equation (3.5). Since the vector of draw of productivities for variety e^j across regions being $a(e^j) = \{a_1(e^j), \ldots, a_N(e^j)\}$, their joint distribution becomes $\phi^j(a^j(e^j)) = exp\{-\sum_{n=1}^N T_n^j(z)^{-\theta^j}\}$, while α_j controls the elasticity of substitution across varieties in sector j. This delivers the vertical production structure in this economy with upstream and downstream sectors.[10]

$$q_n^{j,R/B} = [\int_0^1 q_n^{j,M}(e^j)^{\frac{\alpha^j-1}{\alpha^j}} d\phi^j(a^n(e^j))]^{\frac{\alpha^j}{\alpha^j-1}} \tag{3.5}$$

---

[10] Note that this differs from input-output linkages in Costinot and Rodríguez-Clare (2014) and applied recently in quantitative trade models where the production of intermediate goods needs this aggregate as input. Here, intermediate production only needs primary factors, and the retail goods are purchased by consumers only, a more realistic reflection of the retail industry.

$$Q_n^{j,R/B} = z_n^{j,R/B}\left[(h_n^{j,R/B})^{\beta_n}(l_n^{j,R/B})^{1-\beta_n}\right]^{\gamma_n^j}\left[q_n^{j,R/B}\right]^{1-\gamma_n^j} \tag{3.6}$$

The retailers then combine the retail good aggregate with other inputs in a nested Cobb–Douglas production function to produce the final retail good, with share of value-added given by γ_n^j, as in equation (3.6). Both type of retailers uses labor and structure bundle with factor shares given by β_n. Given retail sector's production function, the unit cost of retail good is given by:

$$c_n^{J,R} = (\rho_n^{J,R}\omega_n^{J,R})^{\gamma_n^j}(p_n^{j,M})^{1-\gamma_n^j}, \tag{3.7}$$

$$\text{where } p_n^{j,M} \equiv \left(\Gamma\left(\frac{\theta^j + 1 - \alpha^j}{\theta^j}\right)\right)^{\frac{1}{1-\alpha^j}}\left(\sum_{m=1}^{N}(\kappa_{nm}^M c_m^{j,M})^{-\theta^j} T_m^j\right)^{\frac{1}{-\theta^j}}, \omega_n^{J,R} \equiv \left(\frac{r_n^h}{\beta_n}\right)^{\beta_n}\left(\frac{w_n^j}{1-\beta_n}\right)^{1-\beta_n}.$$

Here, $p_n^{j,M}$ is the price index of the intermediate varieties aggregate derived by applying the properties of the Fréchet distribution to the vector of productivities $\phi^j(a^j(e^j))$, where $\Gamma(.)$ is a gamma function and is evaluated at $\frac{\theta^j+1-\alpha^j}{\theta^j}$.[11] $\omega_n^{J,R}$ is the unit cost of the labor and structure bundle in the retail sector. Since the market structure of the retail sector is also competitive, the price of retail goods shipped from market i to n will be the product of the unit retail cost $c_n^{R,i}$ and the iceberg cost κ_{ni}^R between the two markets $p_{ni}^{J,R} = \kappa_{ni}^R c_i^{J,R}$.

Online Retailer Location: The distinguishing feature of online retailers is that while each brick-and-mortar store is characterized by the productivity in its own location $z_n^{j,B}$, the measure O of online retailer each draws a vector of productivity across different locations $(z_1^{j,R}, ..., z_N^{j,R})$. Online retailers can pay a fixed entry cost

---

[11] The parameter condition that $\theta^j + 1 - \alpha^j > 0$ is assumed to guarantee that the price index is well-defined.

in labor units f_m to locate in region m, from where they import intermediate varieties and ship the retail goods to consumers in different places. The optimal location choice based on profit maximization is then

$$m^* = argmin_m\{\tilde{\sigma} c_m^{j,R} \sum_n (\frac{\kappa_{nm}^R}{P_n^{j,R}})^{\sigma^j - 1} X_n\}$$

The online retailers will optimally locate in region m that minimizes the production cost times the weighted sum of normalized shipping cost to destinations, the weight being the total expenditure of the destination market X_n and the price index of retail goods $P_n^{j,R}$ serves as the normalizer. This expression clarifies the forces of agglomeration as well as dispersion in the model. Online retailers would want to locate in where the shipping cost is the lowest to the largest destination market (HME) or if the imported goods are the cheapest. Balancing these agglomeration forces, the increased concentration will lead to higher wage and land prices, pushing up the cost of production. An additional consideration for online retailer's choice is the entry cost of the location, and entrance only happens when the total revenue across destinations is greater than the cost: $\sum_n (\frac{p_{nm}^{j,R}/\mu}{P_n^{j,R}})^{1-\sigma^j} \eta^j X_n \geq \sigma^j w_m^{j,R} f_m$. Therefore, online retailers will only enter a region if the production cost is lower than the threshold:

$$\bar{c}_m^j = \frac{\mu}{\tilde{\sigma}^j} [\frac{\sigma^j}{\eta^j} \frac{w_m^{j,R} f_m}{\sum_n \kappa_{nm}^R P_n^{j,R} X_n}]^{\frac{1}{1-\sigma^j}} \quad (3.8)$$

The location of online retailers then determine the volume of trade flows of retail goods across regions. To gain tractability and derive closed form solution, I follow the multinational production literature (Arkolakis et al. (2018); Arkolakis, Rodrıguez-Clare, and Su (2017)) to assume that the productivity vectors of online retailers are randomly drawn from a multi-variate Pareto distribution $P(Z_1^j < z_1, ..., Z_N^j < z_N) = 1 - (\sum_{m=1}^N [A_{jm} z_m^{-\phi}]^{\frac{1}{1-\rho}})^{1-\rho}$, with support $z_m \geq (\sum_{g=1}^N A_{jg}^{\frac{1}{1-\rho}})^{1-\rho}$ and $\rho \in [0, 1)$. The scale parameter A_{jm} measures the absolute advantage of region

m in producing sector j goods, whereas θ controls the degree of heterogeneity across different vectors, and ρ controls the degree of heterogeneity within a single vector of different realizations. Define $\xi_m^j \equiv c_m^{j,R} \sum_n (\frac{\kappa_{nm}^R}{P_n^{j,R}})^{\sigma^j - 1} X_n$, so $m^* = argmin_m \{\frac{\tilde{\sigma}\xi_m^j}{z_m^j}$, the probability of a sector j retailer to locate in m can then be expressed as

$$\Psi_m^j = P(m = argmin_m\{\tilde{\sigma}\xi_m^j / z_m^j\} \cap c_m^{j,R} < \bar{c}_m^j) = \psi_m^j (\bar{c}_m^j)^{\phi}, \tag{3.9}$$

where $\psi_m^j = A_{jm}(\xi_m^j)^{\frac{-\phi}{1-\rho}} / \sum_{m=1}^N [A_{jm}(\xi_m^j)^{-\phi}]^{\frac{-\rho}{1-\rho}}$. Online retailers are more likely to locate in a region if it has lower weighted total cost of selling to destinations, or if it has higher productivity in producing retail goods, subject to the elasticity of substitution controlled by ϕ and ρ.

The location of online retailers plays an important role in the model: it determines the intra-regional aggregate trade flows. With a total of measure O of online retailers, the measure of online retailers in location m is $O_m = O\Psi_m^j$. Therefore, the total sales from region m to n is the product of sales per firm and the measure of firms: $(\frac{p_{nm}^{j,R}/\mu}{P_n^{j,R}})^{1-\sigma^j} \eta^j X_n O_m$. We can then obtain the bilateral online retail expenditure share $x_{nm}^{j,R}$ as in equation 3.10, which represents an extended gravity equation of Chaney (2008)'s version of the Melitz model. Unlike standard gravity equation of trade, the numerator ("bilateral resistance") depends not only on the retail production cost of the origin, but on the probability of online retailers locating in that region, as well as the online matching efficiency; the denominator ("multi-lateral resistance") includes both the sum of bilateral resistance as well as the cost of local brick-and-mortar store divided the measure of available online retailers.

$$x_{nm}^{j,R} = \frac{\psi_m^j (\bar{c}_m^j)^{\phi} (c_{nm}^{j,R}/\mu)^{1-\sigma}}{\sum_h \psi_h^j (\bar{c}_h^j)^{\phi} (c_{nh}^{j,R}/\mu)^{1-\sigma} + \frac{1}{O}(c_{n0}^{j,R})^{1-\sigma}} \tag{3.10}$$

3.4.3 Labor Supply

To characterize workers' sorting and heterogeneous labor supply across sectors, I adopt a Roy (1951) framework with probabilistic productivities (Lagakos and Waugh 2013; Hsieh et al. 2019; Galle, Rodríguez-Clare, and Yi 2022; Lee 2020). In each region, workers obtains a vector of region-sector specific productivities $z_n = \{z_n^0, z_n^S, z_n^{1,M}, z_n^{1,R}, z_n^{1,B}, ..., z_n^{j,M}, z_n^{j,R}, z_n^{j,B}\}$ for each unit of its labor provided, for which sector 0 is treated as non-employment as Dvorkin (2014) and Caliendo, Dvorkin, and Parro (2019).[12] The productivities are drawn independently from a Fréchet distribution $\psi_n^{j,K}(z_n^{j,K})$ with shape parameter ν_n and scale parameter $A_n^{j,K}$, $K = \{M, R, B, \emptyset\}$. The scale parameter $A_n^{j,K}$ gives the absolute advantage while the shape parameter ν_n regulates the comparative advantage of workers, jointly determining the sorting pattern on the labor market.

From properties of Fréchet distribution, the joint distribution of productivities draws follows another Fréchet distribution $\psi_n(z_n) = \sum_{j=0}^{J} \sum_{K=M,R} A_n^{j,K} z^{-\nu_n}$. Taking account the idiosyncratic productivity, workers' wage per unit of labor supply is $w_n^{j,K} z_n^{j,K}$, which workers seek to maximize by choosing sector (j, K) optimally. Define the optimum choice set for a sector (j, K) by $\Lambda_n^{j,K} \equiv \{z_n^{j,K} st. z_n^{j,K} > z_n^{H,k} \forall (H, k)\}$, then a worker will choose to be employed in (j, K) if the obtained vector draw of productivities is in this set. Applying the properties of the joint Fréchet distribution for the productivity draws $\psi_n(z_n)$, we can drive the probability of

---

[12] Non-employment is treated as a sector that workers can allocate their labor into, with a wage of w_n^0 per efficiency unit of labor that can be understood as the marginal return for home production, and households' consumption when non-employed depend on the labor units they withdraw from the employment sectors.

non-employment, as well as the employment in sector (j, K) as:

$$\pi_n^0 = \frac{A_n^0 (w_n^0)^{\nu_n}}{\Phi_n}, \ \pi_n^{j,K} = \frac{A_n^{j,K} (w_n^{j,K})^{\nu_n}}{\Phi_n}, \text{ where } \Phi_n = \sum_{j=1}^{J} \sum_{K=M,R} A_n^{j,K} (w_n^{j,K})^{\nu_n} + A_n^0 (w_n^0)^{\nu_n}. \tag{3.11}$$

The probability of being non-employed or employed in a certain sector is shown to be proportional to the return of home production or sectoral wage relative to the total returns of being employed and non-employed, scaled by the Fréchet parameter ν_n that plays the role of the elasticity of labor adjustment.[13] Therefore, as labor demand changes affect wages, they also alter households employment decisions. Another tractability gained from the Fréchet distribution is that the efficiency units of labor supply can be conveniently derived; specifically, for a sector (j, K):

$$l_n^{j,K} \equiv \Gamma(\frac{\nu_n - 1}{\nu_n}) \frac{\Phi_n^{1/\nu_n}}{w_n^{j,K}} \pi_n^{j,K} L_n \tag{3.12}$$

where $\Gamma(.)$ denotes a gamma function. Workers' income as well as firms' production depend on this efficiency units of labor provided, and the wage return for workers in sector $w_n^{j,K} l_n^{j,K}$ becomes $\Gamma(\frac{\nu_n - 1}{\nu_n}) \Phi_n^{1/\nu_n} \pi_n^{j,K} L_n$.

3.4.4 Market Clearing and Competitive Equilibrium

On the goods market there exist two types of expenditures: consumers purchase retail goods across retailers, and retailers acquire intermediate varieties from

---

[13] As discussed in Galle, Rodríguez-Clare, and Yi (2022), if $\nu_n \to \infty$, the households become homogeneous in employment choices and $\nu_n \to 1$ delivers the same comparative statics as sectoral specific labor supply.

different regions. In equilibrium, both of these markets need to be cleared:

$$X_n^{j,R} = \sum_{i=1}^{N} x_{in}^{j,R}(I_i L_i), \text{where } I_i L_i = \sum_{k=0}^{J} \sum_{K=M,R} (r_i^{h,k} h_i^{K,k} + w_i^k l_i^{K,k}) - \Omega_i, \tag{3.13}$$

$$X_n^{j,M} = \sum_{i=1}^{N} (1 - \gamma_i^j) x_{in}^{j,M} X_i^{j,R}. \tag{3.14}$$

The total expenditure or demand of sector j retail goods sold from region n denoted by $X_n^{j,R}$ has to equal to the product of the retail expenditure share on region n's retail goods $x_{in}^{j,R}$ and total income $I_i L_i$ across regions. In the benchmark model, households' total income comes from their wage earnings and ownership of land, minus a region's trade deficit denoted by Ω_i that is assumed to be exogenous.[14] On the other hand, the total demand for sector j intermediate goods from region n, denoted by $X_n^{j,M}$, equates the expenditure share on region n's intermediate goods $x_{in}^{j,M}$ times the portion of retail sector's spending on intermediate varieties $(1 - \gamma_i^j) X_i^{j,R}$ summed across regions. Accounting for regional trade deficits leads to the balance of trade equation:

$$\sum_{j=0}^{J} \sum_{i=1}^{N} (x_{ni}^{j,M} X_n^{j,M} + x_{ni}^{j,R} X_n^{j,R}) + \Omega_n = \sum_{j=0}^{J} \sum_{i=1}^{N} (x_{in}^{j,M} X_i^{j,M} + x_{in}^{j,R} X_i^{j,R}). \tag{3.15}$$

The clearing of the markets for primary factors including labor and structures follows the same manner that each of their return needs to equal to the portion of value-added. However, since these factors are used in the production of both intermediate and retail goods that are subject to different production functions, the market clearing conditions differ for intermediate and retail sectors. Specifically, for the labor market:

$$w_n^{j,M} l_n^{j,M} = w_n^{j,M} \int_0^1 h_n(e^j) d\phi_n^j(a_n(e^j)) = \beta_n X_n^{j,M}, \quad w_n^{j,R} l_n^{j,R} = \gamma_n^j \beta_n X_n^{j,R}, \tag{3.16}$$

---

[14] In Section V's discussion of policy interventions, households' total income will also depend on the "tariff" that a local region imposes on others, and an endogenous deficit that is affected by revenue reallocation.

$$r_n^h h_n^{j,M} = (1-\beta_n) X_n^{j,M}, \quad r_n^h h_n^{j,R} = \gamma_n^j (1-\beta_n) X_n^{j,R}. \tag{3.17}$$

Model Equilibrium and Comparative Statics. To characterize the competitive equilibrium for this interregional retail trade framework, we need to specify the economy's fundamentals and model parameters. The fundamentals of the model economy include the sector-region productivities in producing intermediate goods as well as retail goods $(T^M, T^R) = \{T_n^{j,M}, T_n^{j,R}\}_{n=1,j=1}^{N,J}$, workers' productivities in different sectors $A^K = \{A_n^{j,K}\}_{n=1,j=1}^{N,J}, K = \{M, R, B, \emptyset\}$, the demand shifters for retail goods across regions μ, the iceberg trade costs of manufacturing and retail goods $(\kappa^M, \kappa^R) = \{\kappa_{ni}^M, \kappa_{ni}^R\}_{n=1,i=1}^{N,N}$, the stock of structures across markets $(h^M, h^R) = \{h_n^{j,M}, h_n^{j,B}\}_{n=1,j=1}^{N,J}$, and the exogeneous trade deficits of different places $\Omega = \{\Omega_n\}_{n=1}^{N}$. For clarity, here I denote these fundamentals by $\Psi \equiv \{T^M, T^R, A^K, \mu, \kappa^M, \kappa^R, h^M, h^R, \Omega\}$.

The parameters of the model are related to the factor shares, elasticity of substitution of factors in production, as well as the Fréchet distribution parameters, all of which are assumed to be constant. The only endogenous variable of the economy is $\{L_n^{j,M}, L_n^{j,R}\}_{n=1,j=0}^{N,J}$ and all prices can be expressed with respecting wages. The equilibrium can then be defined as below.

Definition 2 (Competitive Equilibrium) *Given the fundamentals Ψ and labor supply L_n, a competitive equilibrium for this economy is a vector of wages $\mathbf{w} = \{w_n^j\}_{n=1,j=0}^{N,J}$ such that the optimality conditions are satisfied and all markets clear – equations (3.10), (3.11), (3.3), (3.4), as well as (3.13) to (3.17) hold.*

3.4.5 E-commerce and Equilibrium Outcomes

E-commerce Shock. Applying the theoretical framework, I intend to answer the question: what are the equilibrium implications of an e-commerce shock on the economy, particularly those related to the dispersion of economic outcomes across different regions? As shown in Definition (2), the model equilibrium is conditional on fundamentals, hence addressing this question requires to specify how the economic fundamentals might be affected the e-commerce. In light of the model in this paper, there are three channels through which e-commerce is likely to bear an impact. Firstly, as online shopping eases consumers' search frictions (Goldmanis et al. 2010; Dinerstein et al. 2018), it may alter the across-region demand shifter μ_{ni}^{j} such that online retailers seize a higher demand. Secondly, the rolling-out of fulfillment and distribution facilities of e-commerce giants such as Amazon significantly reduces the shipping costs of consumer goods (Houde, Newberry, and Seim 2021), lowering κ_{ni}^{R}.

Welfare Analysis. The general equilibrium effects of an e-commerce shock on welfare across different regions can also be conveniently analyzed in proportional changes. Define the welfare of a region by its real income per capita $W_n = \frac{Y_n/L_n}{P_n}$, where $Y_n = I_n L_n + \Omega_n$ is the total income in a region including trade deficit. Y_n can be further simplified into $Y_n = (\frac{1}{1-\beta_n})\Gamma(\frac{\nu_n-1}{\nu_n})\Phi_n^{1/\nu_n} L_n$. The changes in welfare can then be expressed as $\hat{W}_n = \hat{\Phi}_n^{1/\nu_n}\Pi_{j=1}^{J}(\hat{P}_n^{j,R})^{-\eta_j}$. Using labor market allocation, we can get $\hat{\Phi}_n^{1/\nu_n} = \hat{w}_n^0(\hat{\pi}_n^0)^{\frac{-1}{\nu_n}}$, while expression of retail trade share in equation (C.3) leads to that $\Pi_{j=1}^{J}(\hat{P}_n^{j,R})^{-\eta_j} = \Pi_{j=1}^{J}(\hat{x}_{nn}^{j,R})^{\frac{-\eta_j}{\sigma^j-1}}(\frac{\hat{c}_n^{j,R}}{\hat{\mu}})^{-\eta_j}$. Taken together, the counterfactual changes in welfare is:

$$\hat{w}_n^0(\hat{\pi}_n^0)^{\frac{-1}{\nu_n}}\Pi_{j=1}^J(\hat{x}_{nn}^{j,R})^{\frac{-\eta_j}{\sigma^j-1}}(\frac{\hat{c}_n^{j,R}}{\hat{\mu}})^{-\eta_j} \tag{3.18}$$

The above expression of welfare changes highlights several general equilibrium channels that e-commerce could affect an economy with interrelated regions and sectors as well as elastically supplied labor. The term $\Pi_{j=1}^J(\hat{x}_{nn}^{j,R})^{\frac{-\eta_j}{\sigma^j-1}}(\frac{\hat{c}_n^{j,R}}{\hat{\mu}_{nn}^j})^{-\eta_j}$ comes from the changes in consumer retail good price index aggregated across sectors $\Pi_{j=1}^J(\hat{P}_n^{j,R})^{-\eta_j}$ and captures the price effects of a shock. Such effects depend on the consumer expenditure share on a region's local goods $\hat{x}_{nn}^{j,R}$, and a negative power term that comprises the elasticity across retailers σ^j as well as consumers' expenditure shares η^j, both varying at the sector level. A region's expenditure share of its own good and the trade elasticity represent the sufficient statistics for welfare change in a wide variety of trade models, as discussed in Arkolakis, Costinot, and Rodríguez-Clare (2012). By shifting demand towards non-local retailers and reducing transportation friction, the rise of e-commerce will increase welfare through this price channel; adding to that, sectoral heterogeneity in trade elasticities and consumer's expenditure share also matters for welfare in this model.

Two additional terms appear in the composition of the price effects. First, the change in unit cost of local retail good production $\hat{c}_n^{j,R}$ affects the local retail price positively conditional on changes in trade share of a region's own goods. Hence, consumers benefit from reduction in the price of local retail goods if it doesn't alter the trade share of local goods, and note from equation (C.2) that such effect also depends on the input-output linkages. As the price of intermediate inputs decreases, the price of retail goods will also drop depending on the value-added share γ_n^j. Second, the increase in preference for local goods $\hat{\mu}_{nn}^j$ reduces local retail prices conditional its effects on expenditure share on local goods, but since tastes for local goods also affect the trade share $\hat{x}_{nn}^{j,R}$, the total effect on welfare depends on the magnitude of their changes. For both $\hat{c}_n^{j,R}$ and $\hat{\mu}_{nn}^{j}$, their effects on welfare changes and on $\hat{\mu}_{nn}^{j}$ are negatively correlated, hence counterbalances the local expenditure share in determining welfare variation.

With worker heterogeneity in labor supply and imperfect mobility across regions, employment rate across sectors will also affect households' welfare. The term $\hat{w}_n^0(\hat{\pi}_n^0)^{\frac{-1}{\nu_n}}$ represents the income effects on welfare conditional on price changes, and indicates that as non-employment rate decreases or wage return for non-employment rises, welfare will tend to increase. Since the change in total income can be is positively correlated with the change in wage and negatively correlated with the change employment of any sector $\hat{\Phi}_n^{1/\nu_n} = \hat{w}_n^{j,K}(\hat{\pi}_n^{j,K})^{\frac{-1}{\nu_n}}, \forall(j,K)$, as shown in Galle, Rodríguez-Clare, and Yi (2022), this implies that welfare will increase with the degree of specialization of workers. Therefore, regions with workers that have a comparative advantage in the sectors exporting more due to the e-commerce shock will see increases in welfare, while regions that loose jobs due to competition from elsewhere will see reductions in welfare. Taking stock, by explicitly capturing heterogeneous labor supply, and demand shift related to search transportation friction, the model delivers comparative statistics regarding welfare that are comprehensive of the general equilibrium mechanisms through which e-commerce affect different regional economies.

3.5 Model Quantification and Counterfactual Analysis

· In this section, I discuss the quantification of the model to evaluate the impact of e-commerce on regional economies. I first explain the data and measurement with respect to the general economic environment, specifically the fundamentals and parameters necessary to bring the model to the data. I then consider the rise of Amazon as a salient case of e-commerce shock, and discuss how to quantify its impacts on the fundamentals of the model. Counterfactual analysis on regional economic outcomes are presented afterwards.

3.5.1 General Environment

To study the impact of e-commerce on regional economies, I consider 2007 as the baseline economy since only after then the online sales of Amazon started to pick up, and I consider 2017 as the post-Amazon shock equilibrium economy. The model is fit to the data and variables on 50 U.S. states and 2 tradable good sectors (durable, non-durable), service sector, as well as a non-employment sector. For each of the tradable sector, there are three subsectors: manufacturing, online retail, and brick-and-mortar.[15] In the model, a labor market is a region and sector pair, which implies that there are 400 markets in the quantification. Table 3 lists for each model section, the parameters, fundamentals and shocks that need to be calibrated or estimated, and the sources of information, which I discuss below. Appendix C provides further details for the calibration of some of the parameters.

---

[15] In online Appendix Table 1, I show the allocation of 3-digit manufacturing sectors according to the North American Industry Classification System (NAICS) into durable and non-durable sectors. The breakdown by durability of online retail and brick-and-mortar sectors is discussed as below.

Table 3.4: Parameters, Fundamentals and Shocks for Model Quantification

Section	Param.	Description	Estimation/Caliberation
Consumer	η_n^j	Sector share of consumption	CFS 2007
	σ^j	Elasticity of subs. across retailers	Keepa + IV
Labor Supply	π_n^j	Share of empployment	CBP, ACS
	ν^n	Fréchet shape of worker product.	Galle, Rodríguez-Clare & Yi (2022)
Production	β_n^j	Share of structures	BEA + Greenwood et. al (1997)
	θ^j	Fréchet shape of sector product.	Caliendo and Parro (2015)
	γ_n^j	Value-added share of retail goods	BEA, CFS
Trade	$x_{ni}^{M,j}$	Interm. expenditure share	CFS 2007
Amazon Shock	$\hat{\kappa}_{nm}^R$	Iceberg cost change	Amazon data + CFS 2007 + IV
	μ	Matching efficiency	E-stats + CES
	Ψ_m^j	Online retailer location probability	Keepa
	O	Measure of online retailers	E-stats
	T_n^j	Fréchet scale of sectoral product.	Assume constant
	A_n^j	Fréchet scale of labor product.	Assume constant

Consumption and Labor Supply. I calibrate consumers' expenditure share for durable and non-durable sector goods η^j using Bureau of Economic Analysis (BEA)'s regional consumption data. A key parameter for consumption is the elasticity of substitution between different retailers (σ^j). From consumer's CES demand, one can drive that the expenditure on an online retailer i for a representative consumer from n is $X_{ni}^{j,R} = (\frac{p_{ni}^{j,R}}{\mu})^{1-\sigma^j}(P_n^j)^{\sigma^k-1}\eta^j E_n$. Aggregating consumer's expenditure on online retailer i across regions n and taking log, one can express the total sales of online retailer i as a function of prices in equation **??**. The identification challenge of using this equation to estimate the elasticity σ^j is that price is endogenous to other demand side factors that could also change quantity. To solve the endogeneity issue, I apply standard supply shifters to instrument for prices as detailed in Appendix C. The estimated elasticity of substitution between retailers is between 3.1 and 4.0, which is slightly higher than the elasticity between brick-and-motar retailers as in the IO literature, but lower than the elasticity between US commuting zones of 5.5 as in Gervais and Jensen (2019), and between United States and other countries as in Caliendo and Parro (2015).

From BEA, I also obtain value-added information for each region and sector that corresponds to the sectoral income $w_n^{j,K} L_n^{j,K}$ in the model. Since the BEA data doesn't breakdown the retail value added into online retailer and brick-and-motar, I use the regional sales shares computed from Keepa and total e-commerce sales from E-stats to infer regional online retail sales, which divided by the total regional retail sales gives the online retail value added share.[16]

On the worker side, the Census County Business Patterns (CBP) provides data on employment share by region and sector, denoted as π_n^j. To analyze the impact of e-commerce, it is essential to distinguish between employment in online retail and brick-and-mortar settings. Unfortunately, CBP does not separate the retail sector into these categories. Therefore, I utilized the regional online sales share data from Keepa combined with national e-commerce sales figures to estimate the employment shares for each retail type. Additionally, I calculated regional total retail sales by converting the Bureau of Economic Analysis's (BEA) regional retail value-added figures based on the regional value-added share into total sales numbers. However, it is important to note that BEA data does not specify whether retail sales are from durable or non-durable goods. I deduced this information by referencing BEA's regional consumption data for durable and

---

[16] It is important to differentiate between online retail versus brick-and-mortar employment to discuss the impact of e-commerce. Since CBP also doesn't provide the breakdown of retail sector, I adopt a similar imputation method, using the regional online sales share from Keepa and regional total retail sector sales to infer the employment share of the two different retailers. —— Specifically, regional online retail sales share times total e-commerce sales gives regional retail sales. The total retail sales of each region is computed as the regional retail value-added divided by regional value-added share. An additional caveat of the BEA data is that it also doesn't provide durable/non-durable information for the retail sector. I infer this information using BEA's regional consumption data on durable and non-durable goods.

non-durable goods. par For workers' labor supply elasticity, ν^n, I adapt the value estimated by Galle, Rodríguez-Clare, and Yi (2022), which presents a multi-sector Ricardian model with Roy (1951) type sorting of heterogeneous workers whose productivities similarly characterized by joint Fréchet distributions.[17] Bringing the model to data on U.S. commuting zones and other countries for 13 manufacturing and a nonmanufacturing sector and using a model implied Bartick type identification, they estimate the labor supply elasticity (analogous to ν^n) to range from 1.42 to 2.79, which are close to the across occupation elasticities estimated in Burstein, Morales, and Vogel (2019) and Hsieh et al. (2019) ranging from 1.2 to 3.44. Here I specify ν^n equal to 1.5, which is the value from their preferred specification.

Production and Trade. With regard to production, the share of structures in the structure-labor bundle β_n^j can be identified from the value-added share of labor over structure, which equals to $\frac{\beta_n^j}{1-\beta_n^j}$. BEA provides value-added and labor compensation, while Caliendo et al. (2018) derived value-added share of structures to be consistent with the share of capital estimates in Greenwood, Hercowitz, and Krusell (1997). I obtain the productivity dispersion parameter θ^j of different sectors directly from corresponding ones in Caliendo and Parro (2015), which used a multi-sector gravity equation to identify the values. For the value-added share of retail goods, γ_n^j, BEA provides the value-added for each sector, which divided by gross-output gives the share value.

---

[17] In their model, worker differ not only by region and sector, but also by groups that can be categorized by education level and demographics etc., leading to a more nuanced picture of welfare.

3.5.2 Calibrating the Amazon Shock

Shipping Cost Reduction. The drastic expansion of Amazon's centers both in their capacity and in geographic locations should lead to a reduction in shipping costs and as a result, the iceberg cost. To use the facility expansion to estimate shipping cost reduction requires imposing more micro-structures on the fulfillment order flows. I follow Houde, Newberry, and Seim (2021), which shows that more than 90 percent of the order are fulfilled by the 3 nearest centers to the destination, and further specify that the center nearest to the origin to be the one handling the order.[18] Table 3.5 illustrates the reduction in shipping distances due to the roll-out of Amazon fulfillment and distribution facilities. In 2007, an order of a region pair handled by the Amazon facilities on average travels 490 kilometers, while in 2017 it reduces to 288 kilometers. In terms of differences, the shipping distance reduced by 202 kilometers on average, and by 0.5 in log units.

Identification Strategy. Using the actual roll-out of Amazon's facilities to calibrate the shock is subject to key endogeneity issues, namely the location of that the new facilities expand to are correlated with GDP growth, population changes and other demand side factors that could directly affect the outcomes of interest. To overcome the endogeneity issue in Amazon's expansion, my current empirical strategy borrows insights from the transportation and infrastructure literature (Duflo and Pande 2007; Lipscomb, Mobarak, and Barham 2013) to build counterfactual distribution centers with simulated location choices based solely

---

[18] Houde, Newberry, and Seim (2021) applies a probit model of order assignment as $\tau_{ni,f} = \Phi(\alpha_1 d_{fn} + \alpha_2 d_{fi} + \alpha_3 k_f)$. The probability that a facility f processes an order from region i to n, $\tau_{ni,f}$, depends on three factors: the distance from the facility to i and n as well as the capacity of facility f. Therefore, for any order that originates in i and ends up in n, a vector of probabilities represents the chances that it is handled by each of Amazon's facilities. The parameters are then estimated by specifying the labor demand of facility and matching it to the data

Table 3.5: Transportation Cost Reduction via Amazon Facilities

	Mean	Std. Dev.	P25	P75	Corr.
Actual Amazon Facility					
2007	490.2	376.3	234.9	739.0	–
2017	287.9	225.6	124.7	409.0	–
Diff.	−202.2	295.6	−249.8	−12.5	–
Log Diff.	−0.5	0.6	−0.9	0.0	–
Counterfactual Amazon Facility					
2007	623.4	400.3	349.6	897.4	0.10
2017	335.2	278.4	143.9	412.1	0.58
Diff.	−288.2	361.8	−355.9	0.0	−0.22
Log Diff.	−0.7	0.8	−1.1	0.0	−0.02

on plausibly exogeneous geographic cost factors that are orthogonal to demand-side factors.[19] To implement such identification strategy, I obtain county-level geographic characteristics on land elevation and climate changes from Open Topography Global Datasets, as well as National Centers for Environmental Information (NCEI).

For my simulated instrument, a "budget" for Amazon of a particular year is determined based on the observed number of new facilities built in that year. Then each U.S. county is ranked based exclusively on topographic factors with respect to land elevation, as well as climatic factors that include temperature, precipitation and number of extreme weathers. The counties that rank the highest according to these factors will be assigned distribution facilities first depending on the budget for each year. Table 3 shows this cross-sectional probit regression of an indicator whether a county was assigned an Amazon facility on different

---

[19] Alternatively, I may leverage the changes in sales tax collection on Amazon, the so called "Amazon tax", or the nexus tax laws imposed by different states requiring sales tax collection where Amazon maintains a physical presence to identify the impact of e-commerce (Baugh, Ben-David, and Park 2018; Houde, Newberry, and Seim 2021). The major challenge for this kind of identification strategy is whether those places that passed these laws are plausibly comparable to those that didn't. I potentially can use this alternative strategy as a robustness check.

Table 3.6: Probability of Amazon Facilities on Geographic Cost Factors

Dependent: 1{AMZ Center}		
Temperature (Lag)	Mean	-0.011
	Minimum	-0.002
	Maximum	0.046^{***}
Precipitation (Lag)	Mean	-0.032
	Minimum	0.043
	Maximum	-0.015
Elevation	Mean	-0.001^{***}
	Minimum	0.000
	Maximum	0.001^{***}
Tornado	Magnitude	-0.051
	Injuries	-0.110
Year FE	X	
Observations	55,259	
Psudo R-squared	0.1663	

geographic factors. The observed facility expansion pattern prefers locations that are warmer and locate on a slightly lower-elevated land. Precipitation and number of tornados are negatively correlated with facility construction, though not significantly.

As a robustness check, the bottom of the table shows that the spearman rank correlation between the suitability index of distribution facility location and GDP growth is significantly negative, corroborating that exogeneity of the instrument with respect to demand-side factors related to economic growth. Figure 5 presents the counterfactual centers based solely on the geographic factors and Amazon's budget for each year. Overall, the simulated facilities approximate the pattern of actual facility locations reasonably well. For earlier years, these geographic factors aren't able to replicate Amazon's location decisions, over-predicting the centers built in California. Whereas for 2010-2020, the counterfactual locations match the observed locations quite well, though

Figure 3.5: Location of Counterfactual Fulfillment and Distribution Centers

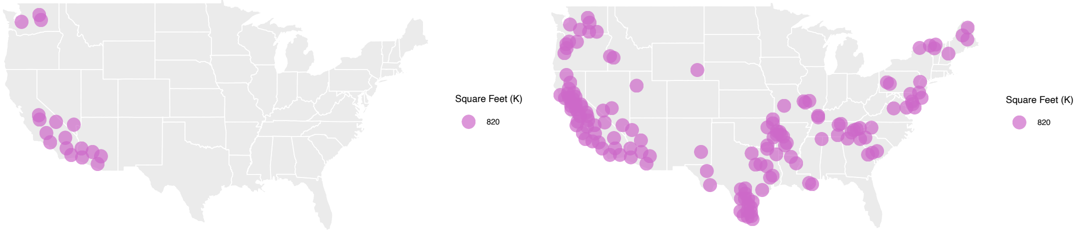
(1) Counterfactual Amazon Facilities Before 2010

(2) Counterfactual Amazon Facilities Before 2010-2020

California is still over-predicted and in Texas and Florida are under-predicted by the simulation model.[20]

The counterfactual transportation cost reduction to be used as instruments can be computed using these predicted location of Amazon's facilities using an order assignment parameters α estimated above. Table 3 presents that through the counterfactual facilities, an order shipped between a pair of regions travels 623 kilometers on average in 2007, while that distance reduces to 338 kilometers in 2017, a 288 kilometers reduction or 0.7 in log units. The greater shipping distance reduction via the counterfactual centers during the same period is due to the relatively densely distributed shipping centers in earlier year and more dispersed locations later. A major concern of using these geographic cost factors predicting center location is about their relevance.[21]

Demand Shift. Another key aspect that Amazon affects the economy is through increased online matching efficiency, hence turning the demand more towards

---

[20] To improve the precision of simulated centers particularly for earlier periods, I plan to include additional factors for Amazon's decision and experiment different specifications. Another factor that affects Amazon's location decision that is relatively exogeneous is that it doesn't want to build centers too adjacent to its existing ones. This will be added in the future analysis.

[21] Online Appendix Table 4 displays the first-stage results regressing actual shipping distances on predicted ones, showing a strong correlation between the two with a relatively large F-stats. I also examine the correlation between the counterfactual shipping distance and lag GDP and GDP growth, and correlation is either weak or negative for these demand indicators, corroborating the robustness of the instrument to demand side factors.

non-local retailers. As shown in the theoretical section, the match efficiency channel is reflected as a demand shifter μ in the nested Cobb-Douglas and CES consumption function. Taking the first order condition of the consumption function, log-linearize and take the difference between the initial value and its change due to Amazon, I obtain equation (3.19) that relate the changes in retail expenditure share $x_{ni}^{k,R}$ to the changes in cost of retail goods $c_i^{k,R}$ and in transportation cost as well as demand alteration, κ_{ni}^{R} and μ,

$$\Delta ln(x_{ni}^{Amz,k}) = (1-\sigma^j)\Delta ln(c_i^{Amz,k}) + (1-\sigma_k)[\Delta ln(\kappa_{ni}^{Amz}) + \Delta ln(\mu)] + \delta_n^k + \upsilon_{ni}^k. \tag{3.19}$$

This equation has an intuitive interpretation – it shows that conditional on cost of goods changes, the changes in consumer expenditure share within a region is either due to transportation cost variation or shifts of demands towards different retailers. Since we've estimated the transportation cost change from the last step, and the cost of good change can be directly obtained from CFS data, it is straight forward that the changes in match efficiency induced demand shift $\Delta ln(\mu)$ can be recovered as long as the changes in consumer expenditure shares due to Amazon $\Delta ln(x_{ni}^{k,R})$ is known. To measure the expenditure share change, the exogeneous changes in transportation costs through the constructed counterfactual centers appear to be useful. On the one hand, the transportation cost serves as an exposure measure of online sales and credible predictor of consumer expenditure shifts due to Amazon; on the other, the counterfactual variation in shipping distance is orthogonal to other demand factors that potentially affect other outcome variables.

$$x_{ni}^{Amz,k} = b_1 c_i^k + \delta\hat{\kappa}_{ni}^{Amz} + b_2\hat{\kappa}_{ni}^{Amz} Amz\hat{S}ale_n^k + Z_{ni}\Gamma' + \eta_t + \epsilon_{ni}^k. \tag{3.20}$$

Equation (3.19) displays the predictive regression where c_i^k is the csot of

Table 3.7: Estimates of Iceberg Cost Change and Demand Shift

δ^{dur}	δ^{nondur}	$\hat{\kappa}$	μ
1.5 [0.2]	2.1 [0.6]	0.97 [0.15]	1.27 [1.46]

sector k goods produced in i, $\hat{\kappa}_{ni}^{Amz}$ is the instrumented transportation reduction induced by Amazon's facility expansion, $Amz\hat{S}ale_n^k$ is the Amazon region-sector level sales, and Z_{ni} is the average demographics for pairs of regions. Since Amazon's regional sales could also be endogenous to other factors affecting the outcome, it is instrumented in a Hausemen approach. Then consumers' predicted expenditure share variation due to Amazon can then be recovered from the estimated coefficients as $\Delta x_{ni}^{Amz,k} = \delta\Delta\hat{\kappa}_{ni}^{Amz} + b_2\Delta(\hat{\kappa}_{ni}^{Amz} Amz\hat{S}ale_n^k)$. Prediction results show that Amazon's expansion in 2007-2017 predicts a 4.8 percent growth in consumers' expenditure share on non-local goods, which compared the actual expenditure share change of 16 percent on average, indicates that Amazon alone accounts for about 30 percent of the total increase in the purchase of non-local goods. Using $\Delta x_{ni}^{Amz,k}$ and $\Delta\hat{\kappa}_{ni}^{Amz}$, the estimate of demand shift is shown in Table 3.7. On average, consumers become about 27 percent more like to purchase from online retailer due to the growth of Amazon during 2007-2017.

3.6 The Impact of Amazon on Regional Economies

In this section, I evaluate the impact of Amazon's expansion on the aggregate and regional economies. The counterfactual question that I ask is that starting from the initial equilibrium, only the Amazon shock as embodied in iceberg cost change and match efficiency increase happen, keeping all other fundamentals constant, what are the impacts on aggregate and regional welfare and employment? To answer this question, I take the calibrated parameters and fundamentals as well as estimated Amazon shocks to the model to conduct counterfactual analysis. I also decompose different channels and compare which margin accounts more for the total effects.

Welfare: Starting with the changes in welfare induced by the Amazon shock. On average, states see a decline in total welfare of 1 percent. The slight decrease in total welfare on average is driven by the income effects, while the price effect has a positive impact on welfare. Leaving only to price effect, which is a result of price decline due to the Amazon expansion, total welfare would have increased by 40 percent. Mitigating the consumption benefit is the fact that Amazon's expansion also leads to the reallocation of economic activities, as well as of workers, changing the income level differentially across regions. The effect on welfare due to income changes would have decreased total welfare by 29 percent without the compensating price changes.

Underlying the overall welfare changes is a huge dispersion across regions. Figure 3.6 shows the state-level changes in total welfare and the decomposition into price and income effects. States in the north and middle west, which has a lower consumer expenditure on online retail goods enjoy a higher welfare due to the positive price effects. Wyoming and South Dakota have the highest regional price decline due to the expansion of Amazon, followed by Iowa and Montana. Meanwhile, states with a more diversified industrial composition, which tend to be larger states (such as North Dakota, California, and Washington), enjoy higher income effects since the reallocation of economic activities tends to tilt towards these regions, and workers in these locations also have better alternative options.

Figure 3.6: Welfare, Employment Changes and Decompositions

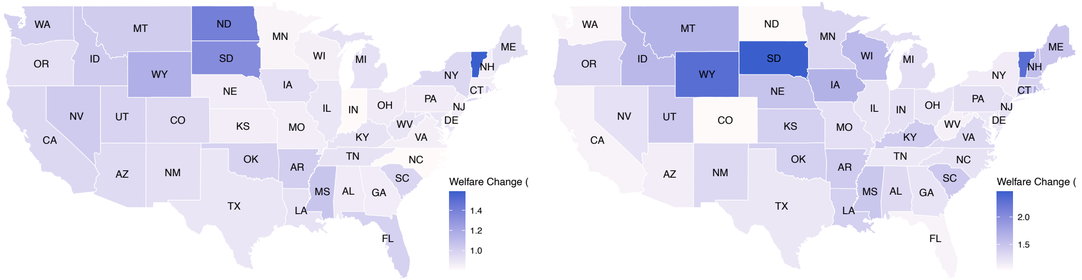

(1) Total Welfare Change

(2) Price Effect

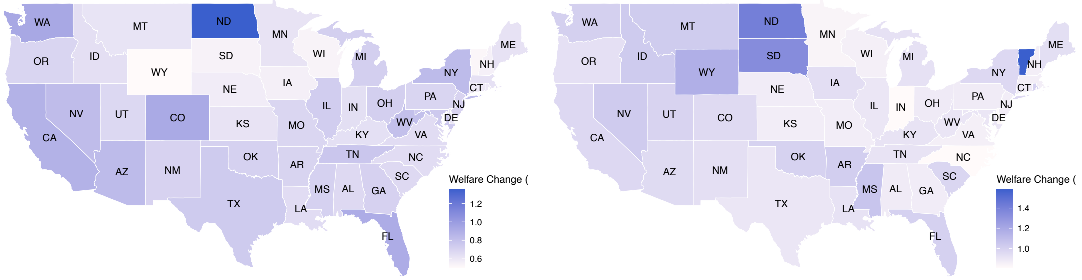

(2) Price Effect

(4) Employment Change

Table 3.8: Employment Changes by Sector and State Groups

	All States		Below 75th Percentile Online Sales Density		Below 50th Percentile Online Sales Density	
Sector	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Manufacturing	-4.3	(7.6)	-2.9	(6.0)	-1.8	(1.1)
Online Retail	109.8	(97.8)	87.6	(73.8)	63.3	(64.8)
Brick-and-Mortar	-11.1	(8.0)	-9.8	(6.2)	-8.6	(1.2)
Service	-1.6	(7.9)	-0.1	(6.2)	1.2	(1.2)
Non-Employment	-1.3	(8.1)	0.3	(6.3)	1.7	(0.8)

Employment: I now turn to discuss the employment changes implied by the Amazon shock and the model. Table 3.8 illustrates the average sectoral employment changes due to the Amazon shock, expressed in the ratio of the post-shock effective units of labor relative to the initial equilibrium. As can be seen from the table, the overall picture of employment changes due to Amazon is characterized by reallocation from the manufacturing sector to the retail sector and non-employment. Since the effect on non-employment is in terms of ratios, to convert it into levels, in 2007 the average non-employment rate was 38.5 percentage points, which implies that non-employment has increased by 2.3 percentage points (6 percent) due to the Amazon shock. Beneath this overall increase in non-employment is a huge regional dispersion. Regions with higher a less diversified industrial structure have higher increases in non-employment, resulting in higher dispersion of non-employment. The Gini index of non-employment has increased from 0.11 to 0.13, a 20 percent growth. This implies that the gap in employment opportunities has become wider due to Amazon.

3.7 Alternative Modeling of Online Retail Location

An alternative modeling approach is to follow Chaney (2008) to understand the dynamics of retailer entry and trade. We begin by considering the productivity distribution of retailers, represented by a Pareto distribution: $P\left(Z^j < z\right) = G^j(z) = 1 - z^{-\rho}$. Retailers decide to enter the market based on a profitability condition: the expected revenue must be greater than or equal to the costs of entry. This condition is given by $\sum_n \left(\frac{p_{nm}^{j,R}/\mu}{P_n^{R,j}}\right)^{1-\sigma} \eta^j Y_n \geq \sigma w_m^{j,R} f_m$. The threshold for entry denoted as $\bar{c}_m^j$ is then given by [22]

$$\bar{c}_m^j = \frac{\mu}{\widetilde{\sigma}} \left(\frac{\sigma}{\eta^j}\right)^{\frac{1}{1-\sigma}} \left[\frac{w_n^{j,R} f_m}{\sum_n \left(\kappa_{nm}^R / P_n^{R,j}\right)^{1-\sigma} Y_n}\right]^{\frac{1}{1-\sigma}}. \tag{3.21}$$

The trade flow equation, can then be derived to link to the relative productivity and cost structures of the trading regions. The bilateral export $X_{nm}^{j,R}$ from region m to n is a function of wage rates, productivity, and the relative costs of

---

[22] Since $\bar{c}_m^j = \frac{1}{z_m^j} \left(w_m^{j,R}\right)^{\gamma^j} \left(P_m^{j,M}\right)^{\left(1-\gamma^j\right)}$, we can also drive the threthold productivity $\bar{z}_m^j = \left(w_m^{j,R}\right)^{\nu_i^j} \left(P_m^{j,M}\right)^{\left(1-\gamma_i^j\right)} \frac{\tilde{\sigma}}{\tilde{\mu}} \left(\frac{\sigma}{\eta^j}\right)^{\frac{1}{\sigma-1}} \left[\frac{w_m^{j,R} f_m}{\sum_n \left(\kappa_{nm}^R P_n^{R,j}\right)^{1-\sigma} Y_n}\right]^{\frac{1}{\sigma-1}}$

retailing and manufacturing, as below.[23] This equation below suggests that an increase in the productivity or a decrease in the wage rate of the exporting region (region m) would lead to an increase in exports $X_{nm}^{j,R}$ to region n, all else being equal. Similarly, an improvement in the transportation technology (represented by κ_{nm}^{R}) would increase the trade flow.

$$X_{nm}^{j,R} = \lambda w_m^{j,B} l_m^{j,B} \left(\left(w_m^{j,R}\right)^{\gamma^j} \left(P_m^{j,M}\right)^{\left(1-\gamma^j\right)} \frac{\kappa_{nm}^{R}}{\mu} \right)^{-\rho} \times \left[\frac{w_m^{j,R} f_m}{\sum_n \left(\frac{\kappa_{nm}^{R}}{P_n^{R,j}} \right)^{1-\sigma} Y_n} \right]^{\frac{\sigma-\rho-1}{1-\sigma}} \eta^j Y_n \left(P_n^j\right)^{\sigma-1} . \tag{3.22}$$

Furthermore, local brick-and-mortar (BM) sales $X_{nn}^{j,B}$ in region n are also modeled, capturing the local market dynamics. This equation considers the local wage rates and productivity, and the price index P_n^j:

$$X_{nn}^{j,B} = \left(\left(\omega_n^{j,B}\right)^{\gamma^j} \left(P_n^{j,M}\right)^{\left(1-\gamma^j\right)} \right)^{1-\sigma} \eta^j Y_n \left(P_n^j\right)^{\sigma-1} \tag{3.23}$$

Finally, the expenditure share of region m in n, both for retail and BM sales, is derived from these trade flow equations, as well as the price index P_n^j. These shares reflects the distribution of consumption across different regions and sectors. For a more detailed mathematical exposition of the derivation, please refer to the Appendix.

The alternative modeling approach presented here offers a different view of the retail market dynamics. Rather than thinking of online retailing location as choice of onlinre retailers based on multivariate Pareto distribution, the

---

[23] In this equation $\lambda \equiv \tilde{\sigma}^{-\rho} \left(\frac{\sigma}{\eta^j \mu^{1-\sigma}} \right)^{\frac{\sigma-\rho-1}{1-\sigma}} \frac{-\rho}{\sigma-\rho-1}$

entry model of Chaney (2008) represents online retail entry based on regional conditions. As detailed in the Appendix, the comparative statics based on this alternative model is also different and requires additional calibration of the iceberg cost in the first period, price index of entire retail sector, and the change in online retail efficiency. Despite these differences, quantitative results based on this Alternative model are consistent with the main model.

3.8 Conclusion

The rapid expansion of e-commerce, as exemplified by Amazon's growth, has brought about significant shifts in regional economies and labor markets. This paper's findings show the noticeable effects of online retailing on spatial economic disparities. In particular, while e-commerce has led to a general decline in retail prices, benefiting consumers, there has been a noticeable worker reallocation away from manufacturing sectors, contributing to a 1 percent average decrease in welfare. While some regions have reaped the benefits of increased trade and economic activity, others have faced challenges. This results in an overall increase in regional inequality, and indicates important redistribution effects of e-commerce. This paper's results imply that the growth of e-commerce, despite its benefits in terms of efficiency and consumer choice, requires careful policy considerations to reduce adverse impacts on less advantaged regions and sectors.

APPENDIX A

APPENDIX FOR ”THE CHANGING SKILL MIXING IN THE US ECONOMY”

A.1 ADDITIONAL EMPIRICAL RESULTS

A.1.1 Data Construction

In this section, I give more details on data construction for the two primary datasets on job skill demand employed in Section 1.3 and 1.4, namely O*NET (Occupation Information Network) and Lightcast (previously known as ”Burning Glass”). Specifically, I discuss strategies for leveraging the longitudinal information in these datasets with higher precision. I also present an overview of their inherent characteristics, advantages and disadvantages, and how they are cross-walked with other datasets used in the analysis.

O*NET: Administered by the U.S. Department of Labor, O*NET is a replacement for the Dictionary of Occupational Titles (DOT). It is more comprehensive and more frequently updated and has been used widely to analyze occupation skill requirements and work settings (i.e., Acemoglu and Autor 2011; Yamaguchi 2012; Autor and Price 2013).

Nonetheless, to use the longitudinal variation from O*NET, the key challenge concerns partial updating – each new version of O*NET only updates an average of 110 targeted occupations among the 970 7-digit occupations. Online Appendix Table A.1 lists different versions of O*NET, the release year, and the year composition for 3 of the modules. Specifically, for each release of O*NET, I assign a "Considered Year" such that at least 55% to 60% of occupations are updated after that year.

Moreover, I use 4-year intervals. The last column of online Appendix Table A.1 shows the percent of occupations that are updated from the last considered year of data included in the analysis. On average, more than 50 percent of the occupations are updated across the succeeding years included in the analysis.

O*NET contains around 270 descriptors about occupations that are grouped into 9 modules: abilities, knowledge, skills, work context, work activities, experience/education requirement, job interest, work values, and work styles. For my main analysis, I only use descriptors from 3 modules: work context, work activities, and knowledge that are more interpretable as the skill requirements and are consistently evaluated by incumbent workers for each new release. These descriptors come as importance, level, extent, and relevance. To interpret the skill measures as gauging the intensity, I use the importance information, similar to i.e., Acemoglu and Autor (2011) and Guvenen et al. (2020), but the level and importance pieces of information are highly correlated and do not affect the qualitative patterns of skill mixing shown in the paper.

In Section 1.3, I show the longitudinal changes in skill mixing by combining O*NET and ACS datasets. O*NET uses SOC 2000 occupation classification for releases between 2000 and 2010 and SOC 2010 for years after 2010. To link O*NET and ACS, I first bridge SOC codes to the census' OCC 2000 and OCC 2010 codes respectively using crosswalks provided by the Analyst Resource Center and the Bureau of Labor Statistics. Then different years of OCC codes are homogenized using a balanced and consistent panel of occupation codes developed by Autor and Dorn (2013) and updated by Deming (2017). The same code is also used for combining all years of ACS and O*NET data.

Lightcast: Lightcast (formerly ”Burning Glass Technologies”) is an analytics software company that has developed a comprehensive and detailed dataset derived from online job postings, capturing real-time labor market information, and reflecting the current demand for skills and occupations. One of the key advantages of Lightcast data is its extensive coverage and high-frequency updates. By examining over 40000 online job boards and company websites, it provides a near universe of online posted vacancies; moreover, it provides a level of detail that is rarely matched by other sources of labor market data, such as job titles, employer information, and specific skill requirements. This allows for a very granular analysis of job skill requirements and labor market dynamics across different industries and regions.

The information that Lightcast collected is then parsed and deduplicated into a systematic list of thousands of codified skills. Similar to Hershbein and Kahn (2018) and Braxton and Taska (2023), the dataset that this study uses defines different skills if the codified skills from Lightcast contain relevant keywords. Specifically, the keywords used to capture analytical skill are: ”research”, ”analy”, ”decision”, ”solving”, ”math”, ”statistic”, and ”thinking”. The keywords used to capture interpersonal skills are ”communication”, ”teamwork”, ”collaboration”, ”negotiation”, and ”presentation”. For each occupation, the share of posted vacancies that require a particular skill is then the measure of skill for that occupation, capturing the extensive margin of firm skill demand.

However, like any data source, Lightcast data also has its limitations. For instance, it only covers online job postings, which may not represent the entire labor market, especially for low-skilled jobs or jobs in small firms that do not typically advertise online. It may also have a bias towards certain types of jobs or industries that use online job advertisements more frequently, and online vacancies by nature overrepresent growing firms (Davis, Faberman, and Haltiwanger 2013). One note of Lightcast data is that the measure of skill as introduced above focuses on the extensive margin – whether a job uses a skill or not – this is very different than the level and importance information that O*NET contains.

Table A.1: O*NET Versions and Corresponding Years

	Released Year	Division	Work Context	Work Activities	Knowledge	Skills	Abilities	Considered Year
O*NET 13.0	2008	Post 2005	73.79%	73.79%	73.79%	73.79%	73.79%	2005
		Before 2005	26.21%	26.21%	26.21%	26.21%	26.21%	
O*NET 18.0	2013	Post 2009	57.15%	57.21%	57.21%	99.89%	57.21%	2009
		Before 2009	42.85%	42.79%	42.79%	0.11%	42.79%	
O*NET 22.0	2017	Post 2013	57.84%	57.67%	57.67%	57.67%	57.67%	2013
		Before 2013	42.16%	42.33%	42.33%	42.33%	42.33%	
O*NET 25.0	2022	Post 2018	54.52%	54.52%	54.52%	54.52%	54.52%	2018
		Before 2018	45.48%	45.48%	45.48%	45.48%	45.48%	

*Notes: The table summarizes different versions of the O*NET (Occupational Information Network) database, along with their released year, year division for the 5 modules (work context, work activities, knowledge, skills, abilities), and the considered year for each version. The "Post" and "Before" rows indicate whether the data in each version was collected post or before a particular year. The "Considered Year" column represents the year considered to be corresponding to each release of O*NET based on the year division of data.*

A.1.2 Details of Skill Measures

In this section, I discuss the choice of skill measures used in the main analysis. Specifically, I show the composition of descriptors of each skill used in the main analysis. I also discuss the composite skill measures' validity and correlation with other measures used in the literature.

Table A.2 lists the O*NET descriptors for each of the constructed composite skill measures. The analytical measure corresponds to "non-routine cognitive analytic" and the interpersonal measure corresponds to "non-routine interpersonal" from Acemoglu and Autor (2011). I collapse Acemoglu and Autor (2011)'s "routine cognitive" (the first three items under Routine) and "routine manual" (the last three items under Routine) into a big routine skill, as occupations using these skills have been shown to have had similar labor market dynamics (Autor, Levy, and Murnane 2003; Acemoglu and Autor 2011). I didn't include the "non-routine manual" from Acemoglu and Autor (2011), since it includes descriptors from the "Abilities" module of O*NET that is evaluated solely by job analysts, and for consistency purposes I focus on occupation descriptors that are evaluated incumbents workers.

Further, I include two additional composite skills that are considered to be non-routine. First, I include a "leadership" composite skill that is comprised of descriptors of problem-solving, strategic thinking, teamwork, and communication. They all demand an ability to guide and manage teams, strategize and plan, solve problems, coordinate activities, and communicate effectively within a team or organizational context. Second, I include a "design" composite skill measure centering around technical proficiency and creativity. The composing descriptors entail a strong understanding of design principles, and the ability to draft and layout specifications for technical devices.

Table A.2: O*NET Skill Measures and Composing Descriptors

Non-routine Analytical

- Analyzing data/information
- Thinking creatively
- Interpreting information for others

Non-routine Interpersonal

- Establishing and maintaining personal relationships
- Guiding, directing and motivating subordinates
- Coaching/developing others

Computer

- Interacting With Computers
- Programming
- Computers and Electronics

Design

- Design
- Drafting, Laying Out, and Specifying Technical Devices, Parts, and Equipment

Routine

- Importance of repeating the same tasks
- Importance of being exact or accurate
- Structured v. Unstructured work (reverse)
- Pace determined by speed of equipment
- Controlling machines and processes
- Spend time making repetitive motions

Leadership

- Making Decisions and Solving Problems
- Developing Objectives and Strategies
- Organizing, Planning, and Prioritizing Work
- Coordinating the Work and Activities of Others
- Developing and Building Teams
- Guiding, Directing, and Motivating Subordinates
- Provide Consultation and Advice to Others

*Notes: This table shows the detailed O*NET descriptors for skill measures. The Non-routine Analytical and Non-routine Interpersonal skills align with Acemoglu and Autor (2011)'s "non-routine cognitive analytic" and "non-routine interpersonal" skills. A unified Routine skill measure combines Acemoglu and Autor (2011)'s "routine cognitive" and "routine manual" skills, reflecting their similar market trends. The study omits "non-routine manual" to maintain consistency with incumbent worker-evaluated descriptors. Two additional skills, 'leadership' and 'design', are included to capture managerial and creative competencies.*

Table A.3 shows the correlation among the chosen skills used in the main analysis, as well as math skill and social skill, which are constructed based on Deming (2017), and broader skill measures skills constructed using factor analysis as discussed in online Appendix A.1.6. It reveals the analytical skill (row 1), exhibits a strong positive correlation with computer skills (0.92) and a moderate correlation with math skills (0.50). This pattern suggests that positions requiring analytical skills frequently necessitate computer and mathematical proficiency. Interpersonal skills (row 3) indicate a moderate-to-strong positive correlation with social skills (0.61) and broader interpersonal skills (0.73). This correlation suggests that occupations demanding interpersonal skill also emphasize social abilities. These results validate the interpretation of the analytical and interpersonal skills with a strong positive correlation with math and social skills used in other studies.

On the other, a strong negative correlation exists between routine and interpersonal skills (-0.49) and between routine and interpersonal skills (-0.45), indicating that these skill sets rarely overlap in job requirements. The broader skill categories (rows 7 to 9) align well with their narrower counterparts, reinforcing the validity of these categorizations. In sum, there exist specific, identifiable skills in the labor market, some of which are more aligned with each other, but they tend not to overlap, reflecting distinct competencies.

Table A.3: Correlations Among Skill Measures

Skill Measures	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
(1) analytical	1.00								
(2) routine	-0.45	1.00							
(3) interpersonal	0.44	-0.49	1.00						
(4) computer	0.92	-0.27	0.25	1.00					
(5) math skill	0.50	-0.11	0.12	0.46	1.00				
(6) social skill	0.34	-0.54	0.61	0.24	0.09	1.00			
(7) analytical (broader)	0.84	-0.59	0.55	0.68	0.63	0.57	1.00		
(8) mechanical (broader)	-0.43	0.58	-0.24	-0.38	-0.11	-0.38	-0.49	1.00	
(9) interpersonal (broader)	0.10	-0.35	0.73	0.02	-0.09	0.70	0.28	-0.22	1.00

*Notes: This table reports the correlation among different skill measures constructed using O*NET data from 2000-2020. The first four skills measures in rows (1) to (4) are the ones used in the main text and are constructed using the O*NET descriptors shown in Table A.1. The next two measures in rows (5) to (6), math skill and social skill are constructed based on Deming (2017). Math skill is the average of 1) mathematical reasoning ability, 2) mathematics knowledge, and 3) mathematics skill. Social skill consists of the average of four variables, 1) social perceptiveness, 2) coordination, 3) persuasion, and 4) negotiation. Rows (7) to (9) contain the broader analytical, mechanical, and interpersonal skills constructed using factor analysis as discussed in online Appendix A.1.6 with their specific component variables.*

A.1.3 Alternative Non-parametric Examination of Trend

An intuitive alternative check of the changes in the degree of skill mixing across periods is to non-parametrically plot the density of skill intensities in different skill dimensions. Figure A.1 depicts the density of skill requirements of six skill pairs out of the four constructed skills in 2005 and 2018 respectively using O*NET data combined with ACS. As in previous studies of job attributes, I aggregate the ACS to sex-education-industry cells that implicitly control for changes in task inputs due to variations in skill and industry mixes in the U.S. economy. Employment weights are obtained as the total hours of work aggregated to each cell. The ACS then supplies the O*NET data with employment across worker types to present an overarching picture of skill intensities in the economy.

From the figure, there is a clear shift towards mixed skill requirements in panel (1) pertaining to analytical and computer skills where these skills are positively correlated. Two salient changes happened in this period: first, the entire distribution of skill intensities moves near the 45-degree line; second, there is a significant increase in density around the 45-degree line. Both of these changes will lead to an increased degree of skill mixing, according to how it is defined based on the position of skill vectors relative to the 45-degree line. Such a change is also salient for other non-routine skill combinations: in the analytical and interpersonal skills space (panel 2), as well as in the computer and interpersonal skills space (panel 4).

On the other hand, one can scarcely observe changes in mixing in the routine skill spaces, as shown in panels (3),(5), and (6). From these three plots, there is an increase in density towards the non-routine direction, losing density in routine skill, and the resulting change in relationship with the diagonal does not indicate a strengthening of mixing.

Figure A.1: Non-parametric Depiction of Skill Intensities, 2005 vs. 2018

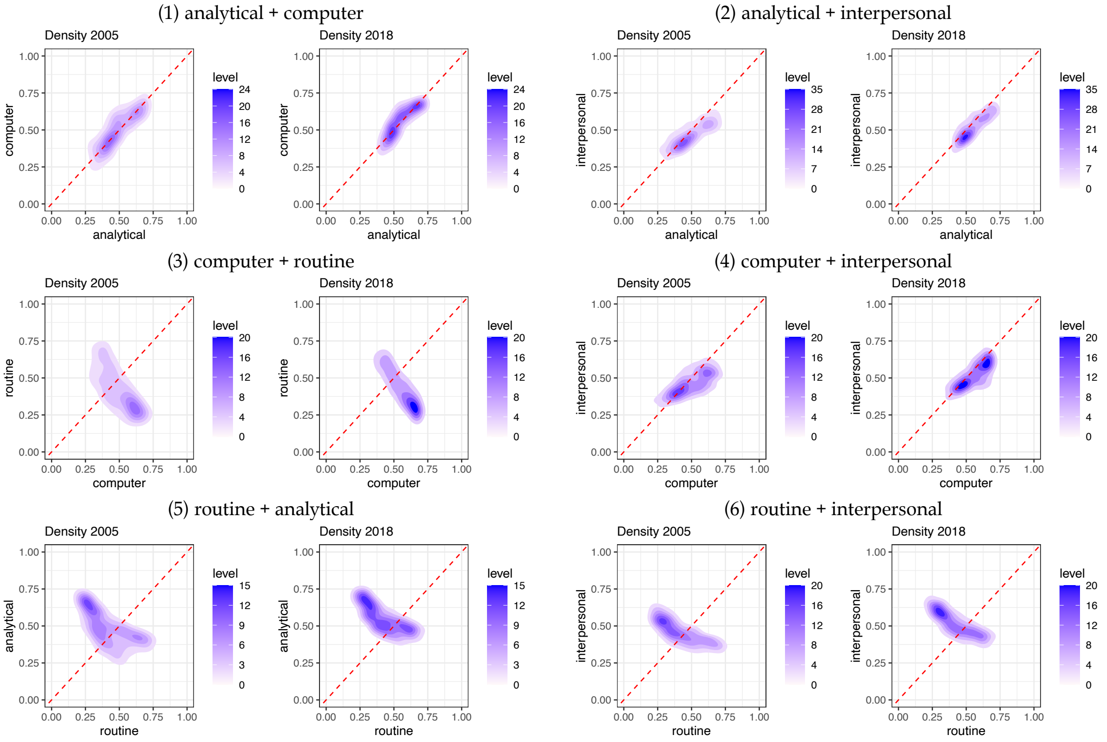

*Notes: These density plots show the intensity of occupation skill requirements across the U.S. economy in 2005 (column 1) and 2018 (column 2) in six two-dimensional skill spaces, as illustrated in the six panels. Darker colors indicate higher density and the 45-degree line is also plotted. O*NET and ACS data are combined for the construction of these plots. The two datasets are merged using consistent occupation codes constructed by Autor and Price (2013) and further developed by Deming (2017). Skill measures are constructed using the O*NET descriptors shown in Table A.1. All measures are normalized to [0,1].*

A.1.4 Robustness of Trend Results to Different Weights and Groupings

In this section, I discuss the robustness of the trend results in different weighting, granularities, and groupings. In particular, I show the density results using weighted skill mixing indexes instead of unweighted ones in the main analysis, as well as at a higher occupations level; the trend of skill mixing using indexes for different skill pairs, instead of high-dimensional indexes; the heterogeneity of skill mixing increases across occupations using indexes for different skill pairs; and the differential changes in skill mixing across industries.

One concern of the analysis of skill mixing shown in Figure 1.2 is that as it shows the changes in the density of skill mixing indexes without weighting, it might not accurately represent the overall picture of mixing in the whole economy. In Figure A.2 panel B, I weigh the skill mixing indexes using employment weight from the OEWS. The results show a similar message that there is a sizable increase in skill mixing particularly for non-routine skills. The only difference is that with employment weighting, the increase in the skill mixing of RNR skills is more discernable. This implies a relatively higher weight of occupations intensive in RNR skills that also increase skill mixing in these skills. In Figure A.2 panel B, I show the density results at a higher 4-digit occupation level, and a similar trend holds.

Next, I discuss the changes in skill mixing using indexes of different skill pairs instead of high-dimensional indexes. Figure A.6 panel (1) shows the results. The figure shows similar results as the main analysis: there is a stronger increase in skill mixing among non-routine skills. For the skill combinations involving routine skills, the change in skill mixing is negligible.

In Table A.5, I show the decomposition results of the changes in the skill mixing indexes for different skill pairs across different datasets. A similar pattern as the main analysis in 1.2, that is within-occupation variation surpassed across-occupation variation in accounting for the increase in skill mixing. This is particularly true using constantly updated occupations at 6-digit occupation level for non-routine skill pairs, and also quite apparent in the Lightcast data. The only slight difference is that for full O*NET data at the 7-digit level, across-occupation variation does contribute to a comparable amount to the change in skill mixing for skill pairs with routine skill.

Figure A.2: Density for Skill Mixing Indexes (Weighted Cosine Distances), 2005 vs. 2018

(1) 7-Digit Occupations

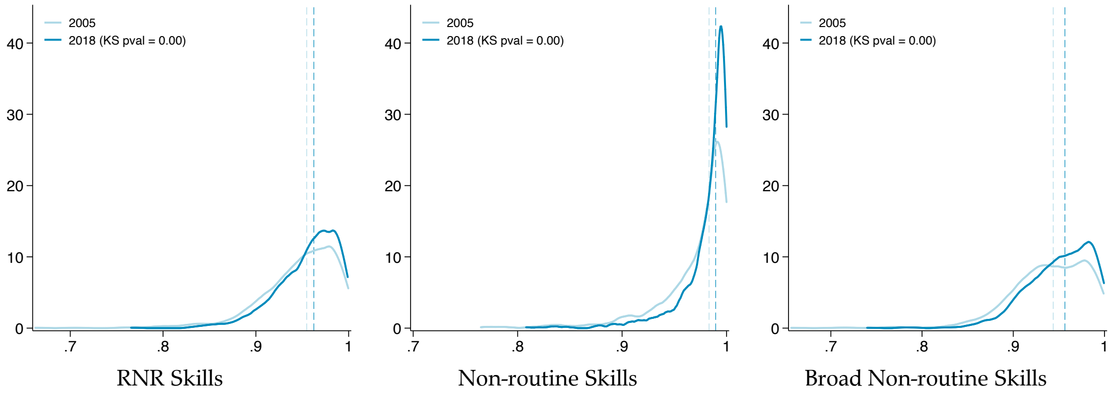

(2) 7-Digit Occupations Weighted by OEWS

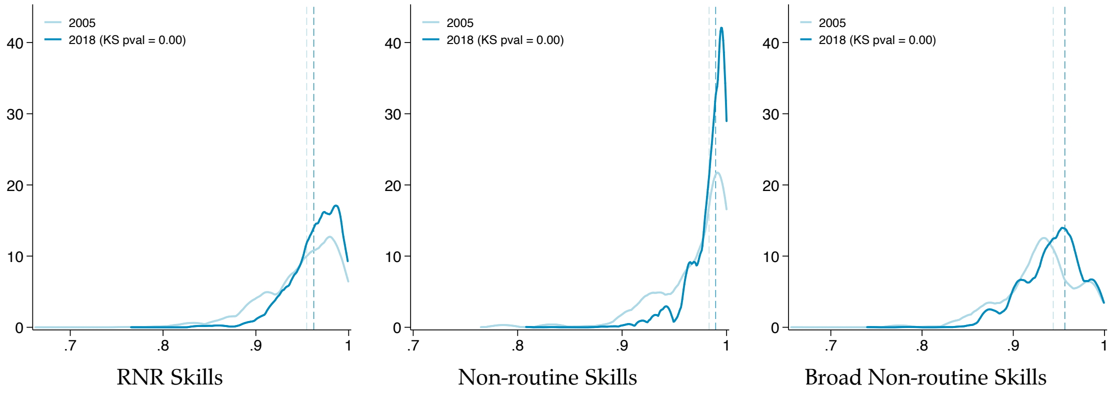

(3) Leave One Skill Out from Non-routine Skill Set

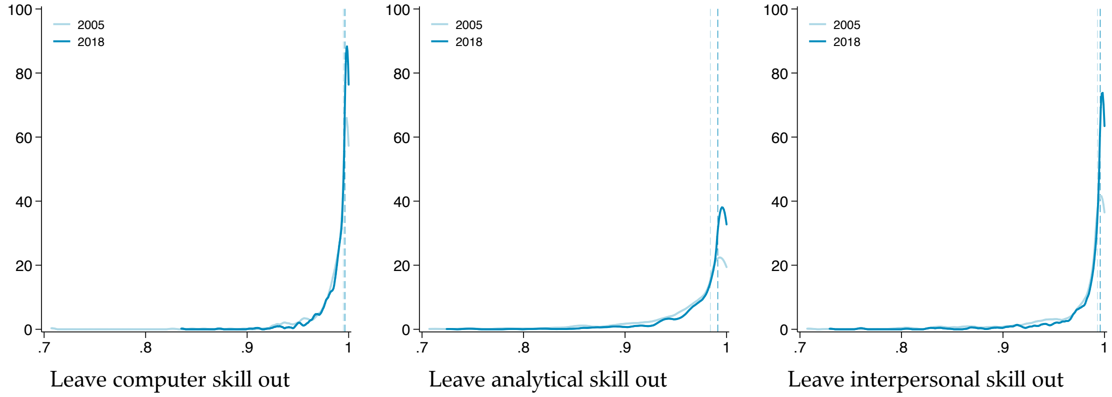

*Notes: These figures plot the PDF of different mixing indexes in 2005 (light blue line) and 2018 (dark blue line). The x-axis displays the value of mixing indexes with a maximum of 1 by construction. These plots are created using O*NET and ACS data merged with occupation codes constructed by Autor and Price (2013) and developed by Deming (2017).*

Figure A.3: Employment Share and Hourly Wage of Full and Updated O*NET

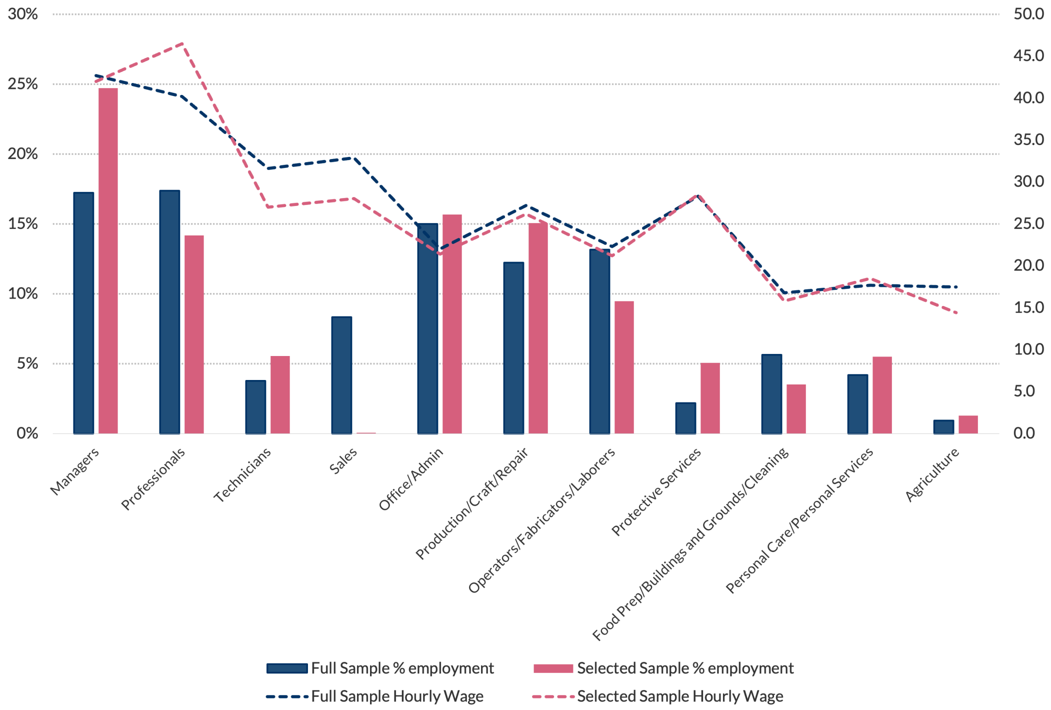

Additionally, in Figure A.3, I show employment percentages and hourly wages across various job categories in the full and the sample for constantly updated occupations. This information gives the occupational structure and returns for these two samples. It can be seen that while professionals make up a smaller percentage in the selected sample, they exhibit a higher average wage, suggesting a focus on higher-earning professionals in the selected sample. Conversely, the sales category shows a drastic reduction in the selected sample, indicating its limited representation. The hourly wage rates across the categories seem fairly consistent between the full and selected samples, with minor discrepancies.

Table A.4: R-Squared Values for Non-Routine Skills' Mixing Index by Polynomial Order

	Analytical	Computer	Interpersonal
1st Order Polynomial			
All occupations	0.11	0.12	0.15
High-wage	0.00	0.02	0.25
White-collar	0.17	0.00	0.38
Blue-collar	0.02	0.22	0.03
Service	0.22	0.37	0.18
3rd Order Polynomial			
All occupations	0.15	0.48	0.21
High-wage	0.03	0.45	0.55
White-collar	0.21	0.20	0.52
Blue-collar	0.05	0.56	0.15
Service	0.30	0.62	0.20
5th Order Polynomial			
All occupations	0.18	0.50	0.22
High-wage	0.04	0.46	0.55
White-collar	0.22	0.21	0.53
Blue-collar	0.07	0.57	0.16
Service	0.38	0.73	0.26

The table presents the R-Squared values from a polynomial regression analysis, assessing the relationship between non-routine skills' mixing index and each composing skill's polynomials up to order N over the period from 2000 to 2020. The regression formula used is $Mix(\mathbf{y})^{percentile}_{ijt} = \beta_1 y^1_{ijt} + \beta_2 y^2_{ijt} + ... + \beta_N y^N_{ijt}$, *where* $Mix(\mathbf{y})^{percentile}_{ijt}$ *indicates the percentile rank of an individual's* i *mixing index of non-routine skills in occupation* j *at time* t, *and* y_{ijt} *is the measure of a specific composing skill for the same individual and occupation at time* t. *The R-Squared values for polynomial orders* $N = 1, 3, 5$ *are provided, illustrating the degree to which each composing skill explains the variance in skill mixing.*

Table A.5: Decomposition of Mixing Indexes' Changes by Skill Pairs

	6-digit Occupations			3-digit Occupations		
Skill Groups	total	within	across	total	within	across
*Full O*NET*						
Analytical + Computer	10.52	6.40	4.12	8.13	6.71	1.42
Analytical + Interpersonal	5.36	2.90	2.46	6.42	4.21	2.21
Computer + Routine	4.38	2.41	1.97	2.65	3.03	-0.37
Computer + Interpersonal	7.23	3.60	3.63	10.28	7.67	2.60
Routine + Analytical	4.00	2.29	1.71	1.52	3.26	-1.75
Routine + Interpersonal	1.93	0.12	1.81	-1.25	1.13	-2.38
Constant Updates						
Analytical + Computer	5.59	6.03	-0.44	6.75	5.96	0.79
Analytical + Interpersonal	3.53	4.58	-1.05	4.24	3.15	1.09
Computer + Routine	2.88	3.69	-0.81	0.77	1.97	-1.20
Computer + Interpersonal	0.78	1.86	-1.09	7.24	6.06	1.17
Routine + Analytical	2.04	2.13	-0.09	1.72	3.69	-1.96
Routine + Interpersonal	0.81	0.82	-0.01	-0.08	1.53	-1.61
Lightcast						
Analytical + Computer		—		13.20	11.74	0.90
Analytical + Interpersonal		—		2.73	2.20	0.31
Computer + Interpersonal		—		-3.90	-3.79	-0.39

Notes: This table shows the shift-share decomposition of changes in the average level of different mixing indexes between 2005-2018 in percentile units. Specifically, for a change in the percentile of a mixing index h over two periods t and τ*, its change* $\Delta T_{h\tau} = T_\tau - T_t$ *which can be decomposed to* $\Delta T_h = \sum_j \left(\Delta E_{j\tau}\alpha_{jh}\right) + \sum_j \left(E_j \Delta\alpha_{jh\tau}\right) = \Delta T_h^a + \Delta T_h^w$ *where* $E_{j\tau}$ *is employment weight in occupation* j *in year* τ*, and* $\alpha_{jh\tau}$ *is the level of mixing index* h *in occupation* j *in year* τ*,* $E_j = \frac{1}{2}(E_{jt} + E_{j\tau})$ *and* $\alpha_{jh} = \frac{1}{2}(\alpha_{jht} + \alpha_{jh\tau})$*.* ΔT_h^a *and* ΔT_h^w *then represent across-occupation and within-occupation change.*

Table A.6: Relationship between Robotics, IT Capital, and Skill Mixing Shifts

	Non-routine Skills		RNR Skills	
	(1)	(2)	(3)	(4)
*A. Skill Mixing Index, 2005-2018 (O*NET)*				
IT capital stock	-0.00	0.09***	0.12**	-0.09**
	[0.02]	[0.03]	[0.04]	[0.03]
Observations	821,030	821,030	821,030	821,030
R-squared	0.07	0.19	0.11	0.24
B. Skill Mixing Index, 2007-2017 (Lightcast)				
IT capital stock	0.02	0.06		
	[0.07]	[0.07]		
Observations	518,520	518,520		
R-squared	0.09	0.26		
*C. Change in Skill Mixing Index, 2005-2010 and 2010-2015 (O*NET)*				
Δ industrial robots	-1.38***	-0.41	-1.77**	-1.24***
	[0.36]	[0.38]	[0.54]	[0.31]
Observations	97,650	97,650	97,650	97,650
R-squared	0.02	0.12	0.04	0.14
Year FE	X	X	X	X
Experience and education controls	X	X	X	X
Gender × education FE	X		X	
Gender × education × industry FE		X		X

*Notes: This table provides regression results on the changes in skill mixing indexes for non-routine and RNR skills between 2005-2018, measured in percentile units based on their distributions in 2005. The analysis integrates data from O*NET and Lightcast to derive skill intensities for calculating skill mixing, which are then merged with ACS data. The data on IT capital stock is sourced from the Bureau of Labor Statistics Total Multifactor Productivity tables. It reflects the productive capital stock for "Total information processing equipment" in billions of 2017 dollars, which is then converted into logarithmic values. The information on the number of industrial robots per thousand workers is sourced from the International Federation of Robotics (IFR), which covers seven industries including Manufacturing, Agriculture, Mining, Utilities, Construction, Education, and Services, and covers the periods 2004-2010 and 2010-2014. Following the methodology in Acemoglu and Restrepo (2020), I use IFR data from Denmark, Finland, France, Italy, and Sweden to assess the influence of global technological progress. Robust standard errors are reported in brackets. *** p¡0.01, ** p¡0.05, and * p¡0.1.*

A.1.5 Additional Results on Trend Heterogeneity

Now, I turn to discuss the robustness of the occupation heterogeneity in skill mixing changes. Figure A.5 provides a detailed view of the changes in the skill mixing of different skill pairs across various occupations. Overall, the increase in the degree of mixing of non-routine skill pairs is higher than the increase in the mixing of skill pairs that include routine skills. Service and blue-collar occupations experience the highest increases in skill mixing of different skills, surpassing white-collar and high-wage occupations. When it comes to routine skills, blue-collar jobs lead other occupations in terms of increase in mixing.

Figure A.5 also provides a detailed view of the changes in the skill mixing of RNR and non-routine skills across various industries. The main patterns indicate that the private service sector, followed by retail trade and construction, leads others in the growth of skill mixing, while public, education, social, and professional services experience the least increases in skill mixing. There is also noticeable heterogeneity across industries in terms of the skills that are mixed. For instance, in finance, real estate, and professional services, there is much higher mixing in non-routine skills relative to RNR skills; conversely, in industries like mining, transportation, public utilities, and construction, RNR skills are mixed in a higher degree.

Figure A.4: Mixing Index Change by Occupation Groups, 2005-2018

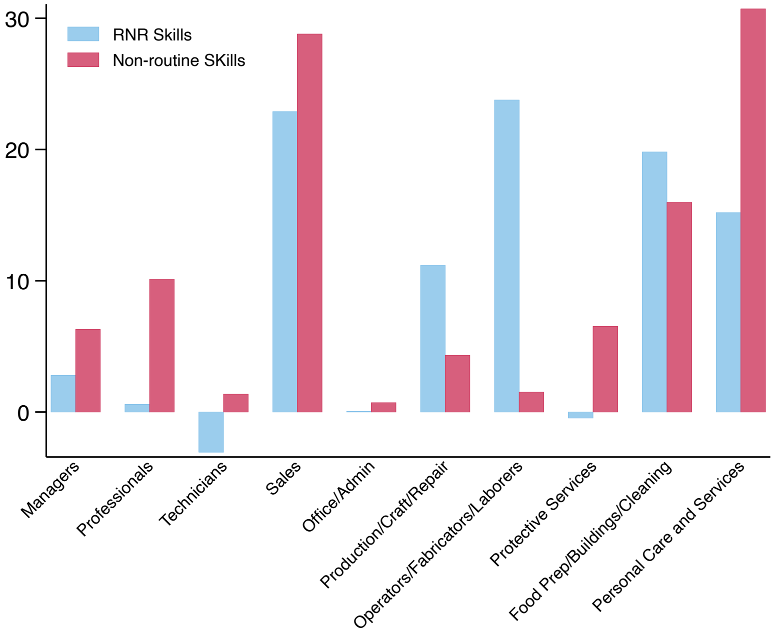

*Notes: This figure plots the changes in skill mixing indexes across different occupation groups. The units of the index changes are percentiles of their distributions in 2005. Workers are categorized into 10 1-digit occupational groups that cover the entirety of US non-agricultural employment following Acemoglu and Autor (2011). O*NET and ACS data are combined for these figures with consistent occupation codes from Autor and Price (2013) and developed by Deming (2017).*

Figure A.5: Mixing Index Change by Industry and Occupation Groups, 2005-2018

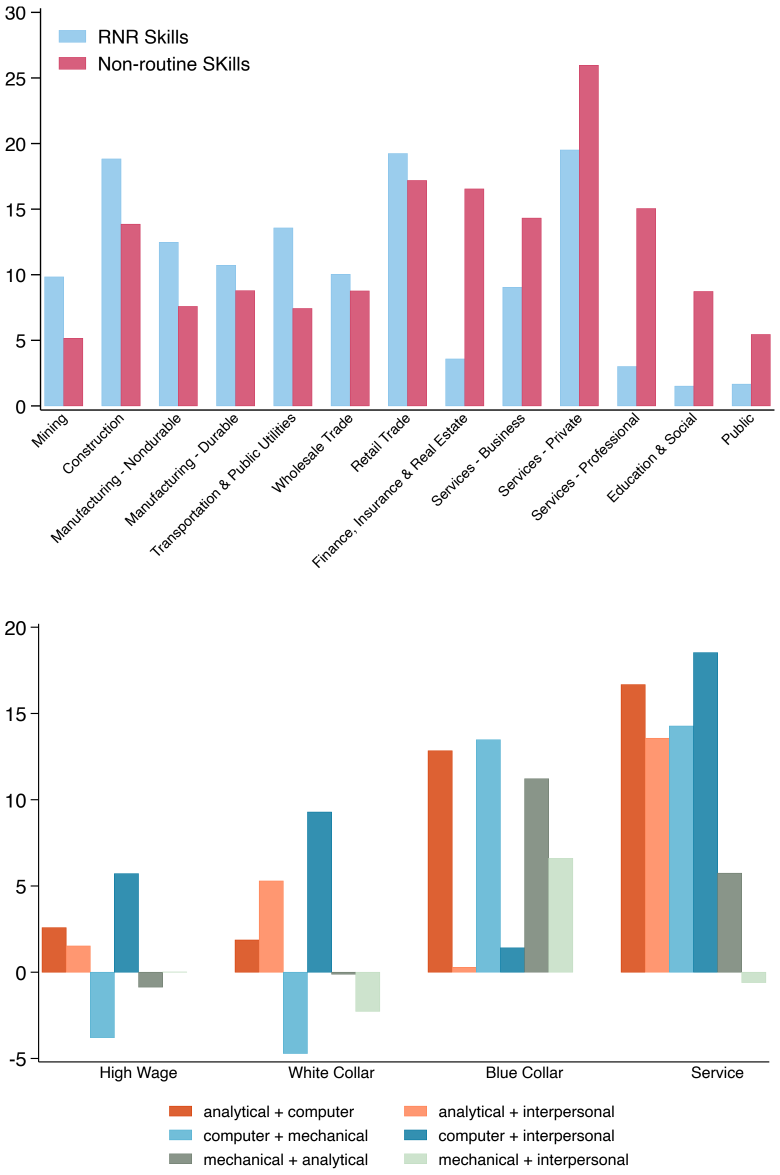

*Notes: These two figures plot the changes in mixing indexes across different industry and occupation groups. The units of the index changes are percentiles of their distributions in 2005. The industry grouping is based on the industrial classification from the 1990 census. The occupation groups (High-wage, White-collar, Blue-collar, Service) follow Acemoglu and Autor (2011). O*NET and ACS data are combined for these figures with consistent occupation codes from Autor and Price (2013) and developed by Deming (2017).*

A.1.6 Robustness of Trend Results to Measures of Skills

In this section, I discuss the robustness of the trend results to using alternative measures of skills. Specifically, I present alternative trend results using different ways of processing skill descriptors from O*NET, such as not using PCA, and standardizing rather than rescaling. I also show the robustness using broader skill measures than those applied in the main analysis.

Alternative Construction of Skills: Since O*NET contains a large number of descriptors, many of which capture the same dimensions of skill requirements, it becomes standard practice to first abstract useful information from the descriptors to construct lower-dimensional measures of skills. The first approach, as in Autor, Levy, and Murnane (2003), Acemoglu and Autor (2011) and Deming (2017), takes the average of a subset of variables and assumes that such average represents a particular broader skill intensity and not others. The other approach, as in Lise and Postel-Vinay (2020), applies PCA to the entire set of variables, which assumes that each of the variables contains information about underlying components that are orthogonally distributed. Both approaches impose different assumptions, with the first one giving more easily interpretable skill groups while the second being more data-driven. A third approach, as in Yamaguchi (2012), first picks descriptors that are ex-ante most easily interpretable with respect to each skill dimension, and then conducts PCA on those descriptors to abstract the most relevant variation. The main body of the paper adopts the third approach; here I show robustness checks using alternative skill measures.

Online Appendix Figure A.6 presents the trend results using skill measures constructed by taking an average of the descriptors without imposing PCA (panel 2) and using skill measures normalized by standard deviation rather than linearly scaled to [0, 1] (panel 3). Normalizing by standard deviation necessarily creates negative values for the skills; since the mixing index is defined based on positive real values, having these negative values invalidates the mixing index in measuring skill mixing. One solution is to add a positive number to the skill measures. As any number chosen is essentially arbitrary, here I added the negative of the smallest value such that the re-scaled measure lies exactly above 0. For both of these robustness exercises, the main message is similar to the main text: there is a significant increase in mixing for non-routine skills, and less so for RNR skills.

Skill Measures: Another concern is that by using skill measures from Acemoglu and Autor (2011), each of which is constructed from a few descriptors, the resulting skill measures could be relatively "narrow" and do not provide a comprehensive depiction of the skill spaces. To alleviate this concern, I construct skill measures using a broader set of descriptors, similar to that of Lise and Postel-Vinay (2020). I first select descriptors from abilities, knowledge, skills, and work activities files that are more relevant for job skill demand, leaving me with around 163 descriptors. I then combine each year's O*NET data with ACS and conduct PCA on the merged data from the years 2005-2018.

The result from this approach supports the choice of analytical, routine, and interpersonal skills in the main text The first three factors out of PCA explain around 60% of the variation across all the descriptors for years. The first factor has a strong positive association with reason and math skills, such as "Deductive Reasoning", "Inductive Reasoning" and "Mathematics", while the second factor relates more to motor coordination and mechanical work, such as "Multi-limb Coordination", "Mechanical" and "Equipment Maintenance". The third factor is clearly more associated with interacting with other people, such as ”Selling or Influencing Others” and ”Resolving Conflicts and Negotiating with Others”. I interpret the second factor as ”mechanical” rather than routine for the broader skill measures.

After conducting PCA, one could directly extract the factors imposing the assumptions that these factors are orthogonal to each other. While this is obviously quite convenient, it nevertheless creates the challenge of interpretability, since each of the factors has been constructed such that it is positively or negatively correlated with all of the 163 descriptors, and the assumption of orthogonality appears strong if the underlying skills are complementary in production across occupations. To take a fine balance between comprehensibility and interpretability, I adopt an approach similar to the measurement validation literature (Costello and Osborne 2005; Thompson and Daniel 1996), where I first conduct PCA/factor analysis to reveal the underlying dimensionality and structure of the measure (as has been done in the previous step). Guided by the factor loadings, I then hand-pick the skill descriptors into three broad groups “analytical”, “mechanical” and “interpersonal” without imposing the orthogonality assumption.

Online Appendix Table A.7 illustrates the selected descriptors for each of the composite skill measures. These descriptors are broadly in line with Acemoglu and Autor (2011) but have several distinctions. First, the descriptors coming from factor analysis lean more toward reasoning, comprehension, and expression. Second, the mechanical skill used in the main text is the average of two ASVAB test scores that are constructed by the weighted average of 26 O*NET descriptors. The ASVAB “Mechanical Comprehension” tests contestants’ “understanding of the principles of mechanical devices, structural support, and properties of materials" and the ASVAB Electronics Information tests contestants' "understanding of electrical current, circuits, devices, and systems", both stressing one's knowledge basis. On the other hand, the descriptors chosen by conducting PCA relate more to physical control, coordination, and machine operation aspects rather than mental perception. Third, the descriptor choices for interpersonal skill from factor analysis also emphasize interactions with others as in Acemoglu and Autor (2011) but are more comprehensive.

Online Appendix Figure A.6 panel (4) illustrates the trend results using these broader skill measures. The message on the growth of skill mixing remains the same as the main text, that is there is strong growth of skill mixing for non-routine skills. Nonetheless, for RNR skills, the degree of skill mixing has decreased using the broader measures.

Figure A.6: Trend of Skill Mixing with Alternative Skill Measures

(1) Skill Pairs

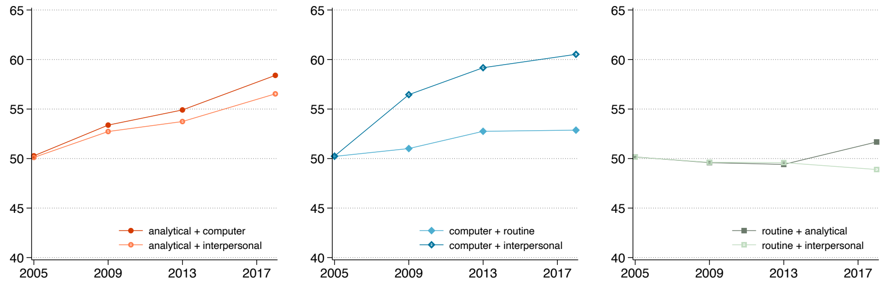

(2) Without PCA (3) Standardized Skill Measures (4) Broader Skill Measures

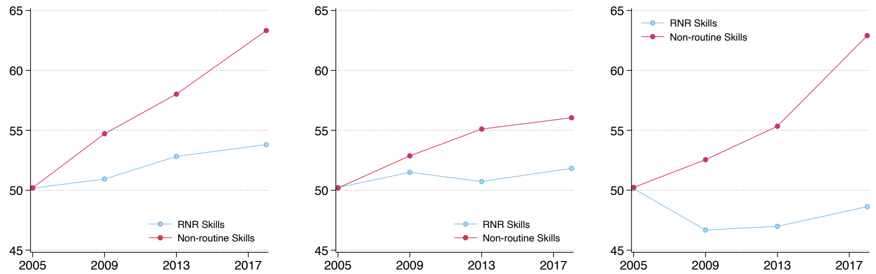

*Notes: These three panels plot the employment-weighted mixing indexes of different skills in the U.S. economy from 2005-2018 using O*NET and ACS data. Panel (1) shows the changes in skill mixing indexes of 6 distinct skill pairs of the 4 skills. In panel (2), skill mixing indexes are calculated using skill measures without using PCA, and in panel (3), skill measures are normalized to have mean 0 and standard deviation 1. Panel (4) shows the changes in mixing indexes using broader skill measures as described in online Appendix A.1.6.*

Table A.7: Components of Broader Skill Measures

Analytical

- Deductive Reasoning
- Inductive Reasoning
- Mathematical Reasoning
- Number Facility
- Mathematics
- Economics and Accounting
- Reading Comprehension
- Writing
- Speaking
- Oral Comprehension
- Written Comprehension
- Oral Expression
- Written Expression

Mechanical

- Multilimb Coordination
- Speed of Limb Movement
- Mechanical
- Performing General Physical Activities
- Handling and Moving Objects
- Controlling Machines and Processes
- Operate Vehicles, Mechanized Devices or Equipmnt
- Repairing and Maintaining Mechanical Equipment
- Repairing and Maintaining Electronic Equipment
- Installation
- Equipment Maintenance
- Repairing
- Production and Processing

Interpersonal

- Assisting and Caring for Others
- Selling or Influencing Others
- Resolving Conflicts and Negotiating
- Coaching and Developing Others
- Staffing Organizational Units
- Service Orientation
- Administration and Management
- Customer and Personal Service

*Notes: This table lists the O*NET descriptor components for each of the constructed (broader) composite skill groups as discussed in online Appendix A.1.6.*

A.1.7 Robustness of Trend Results to Measures of Skill Mixing

I introduce two additional measures and show the robustness of the trend results using these alternative mixing measures.

A first commonly used measure for concentration or specialization based on the share of a total quantity is the Herfindahl–Hirschman Index (HHI).[1] Equation (A.1) shows how to use inverse HHI to measure skill mixing for an occupation represented by $(\alpha_{ja}, \alpha_{js})$. Observe that this index is maximized when $\alpha_{ja} = \alpha_{js}$, exactly corresponding to the case when the skill vector lies on the unit vector and becomes most mixed. If one skill's intensity is greater than the other, the occupation becomes less mixed and this index becomes smaller. Similar to an angle-based mixing index, this measure is insensitive to the length of a skill vector, since each skill is normalized by the total quantity of skills in that occupation.

$$\Big[\Big(\frac{\alpha_{ja}}{\alpha_{ja}+\alpha_{js}}\Big)^2 + \Big(\frac{\alpha_{js}}{\alpha_{ja}+\alpha_{js}}\Big)^2\Big]^{-1} \tag{A.1}$$

$$-\frac{|\alpha_{ja}-\alpha_{js}|}{\alpha_{ja}+\alpha_{js}} \tag{A.2}$$

Under a similar vein, the degree of skill mixing could also be measured by normalizing the absolute distance between skill intensities for a skill vector: as this distance decreases, the overall skill portfolio becomes more balanced; normalization then eliminates the effect of the length of the skill vector. Equation

---

[1] For applications in the labor literature, Ransom and Phipps (2017) and Jin (2017) use the inverse of HHI as the "variety index" to examine how diverse the jobs held for students who graduated from a certain major. Similar logic can be applied to the measurement of skill mixing: in the context of 2-dimensional skill space, the more "varied" skills an occupation uses essentially means that skills are more mixed.

(A.2) gives a particular specification of such a measure. As can be seen from this construction, as the absolute distance between skill intensities decreases and the degree of mixing increases, this measure also increases, though from the direction of $(-\infty, 0)$.

In Online Appendix Figure A.7, I show the robustness of the trend results using these alternative measures in panels (1) and (2). Both measures deliver the same message as the cosine mixing index in the main text, that is, there is a sizable increase in skill mixing, particularly for non-routine skills. The only difference is that for the HHI skill index, there is also a comparable increase in skill mixing for RNR skills.

Figure A.7: Trend of Skill Mixing with Alternative Indexes

(1) Inverse Herfindahl

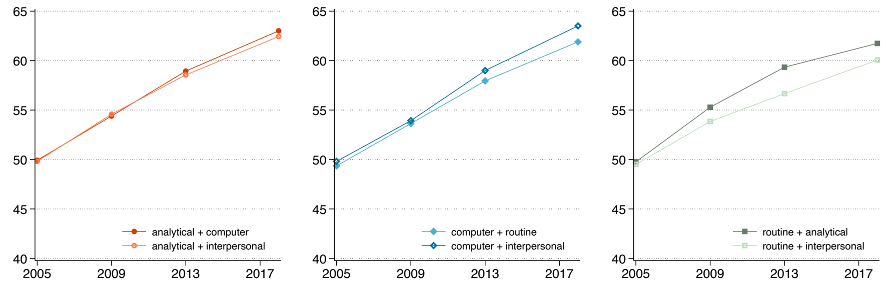

(2) Absolute Distance

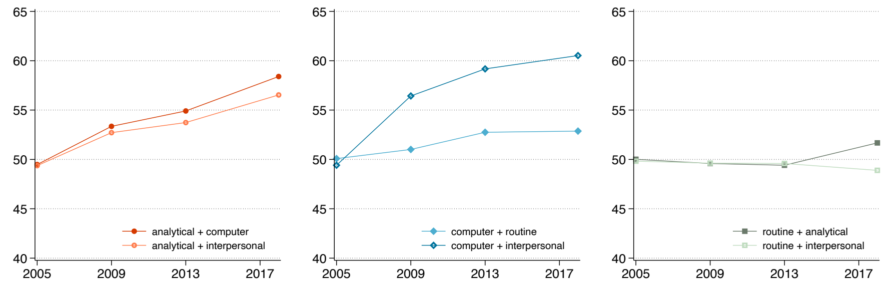

*Notes: These three panels plot the employment-weighted mixing indexes of different skills in the U.S. economy from 2005-2018 using O*NET and ACS data. In panels (1) and (2), mixing indexes are calculated using the Inverse Herfindahl index and Absolute Distance as discussed in online Appendix A.1.7.*

A.1.8 Additional Results on Wage and Employment Returns

In this section, I provide more detailed results on wage returns, and results relating employment to occupation skill mixing. I also provide robustness results to the analysis of the returns to skill mixing in the main paper.

Detailed Results on Wage and Employment Returns: I first check the returns to individual skills and how they interact with the returns to skill mixing. Table A.8 column (1) shows that in a cross-sectional regression analytical and computer skills both have significant positive returns. Workers employed in occupations requiring a higher degree of these two skills earn more. Nonetheless, workers in occupations that require a higher level of interpersonal skills have a wage reduction.

Column (2) of Table A.8 shows that by restricting to within-occupation variation and including skill mixing measures an important pattern appears: the coefficients for most individual skills become slightly more negative (except for routine skill),[2] while the skill mixing indexes of analytical paired with interpersonal skills, as well as routine and interpersonal skills show significant positive returns. Such a pattern persists in columns (3) and (4) including worker skills and fixed effects, only that the skill mixing of analytical and computer skills is more precisely estimated to have a positive return. This indicates that the mixing of skills earns separate and additional rewards beyond those predicted by individual skills.

---

[2] The insignificant or even negative return to analytical skill over time also finds support from the literature. Lise and Postel-Vinay (2020) shows a strong negative 14.4 percent return on cognitive skill using NLSY data with 3-digit occupation fixed effects. Deming (2017) found that the return to cognitive skills has declined across the NLSY79 and NLSY97 cohorts, similar to Castex and Kogan Dechter (2014).

Turning to employment, there is also a positive employment premium for workers with a more mixed skill set. Column (6) of Table A.8 shows that workers with a more mixed level of computer and interpersonal routine skills, or computer and interpersonal skills, or routine and interpersonal skills, are more likely to move from unemployment to employment. Workers with a more mixed level of analytical and computer, or analytical and interpersonal skills, are also more likely to exit unemployment, but the results are not precisely estimated. On the other hand, workers with a more mixed level of routine and interpersonal skills are less likely to find employment.

Robustness Checks: Table A.9 shows the robustness checks to the results in Table 1.4. Specifically, Columns (1) and (2) utilize the Absolute Distance and Inverse Herfindahl measures to formulate mixing indexes (refer to online Appendix A.1.7 for details), while Columns (3) and (4) employ standardized and broader measures of skills (refer to online Appendix A.1.6 for details).

The findings presented in Table A.9 clearly indicate a consistent trend: workers experience a positive return when they are employed in occupations that are more mixed with analytical with computer skills, analytical with interpersonal skills, and routine with interpersonal skills. Specifically, a notable increase in wage is observed with workers in occupations more mixed of analytical and computer skills, especially when applying standardized and broad skill measures; similarly, occupations becoming more mixed of analytical and interpersonal skills, when assessed using the Absolute Distance and Inverse Herfindahl measures, also show a significant positive return. The mixing of routine and interpersonal skills exhibits a positive return as well across the different measures.

On the other hand, the mixing of computer and routine skills, computer and interpersonal skills, and routine and analytical skills all exhibit significant negative wage returns at the occupational level. These negative coefficients may indicate that the combination of these particular skills is less beneficial or leads to inefficiency

Table A.8: Return to Skill Mixing Full Table

Dependent Variable:	ln(hourly wage)				Employed
	(1)	(2)	(3)	(4)	(5)
Occupation Skills					
Analytical	-0.023**	-0.023**	-0.022**	-0.015*	
	[0.009]	[0.010]	[0.010]	[0.008]	
Computer	-0.008	-0.014	-0.015	-0.009	
	[0.010]	[0.011]	[0.011]	[0.009]	
Interpersonal	-0.009	-0.014	-0.015*	-0.013*	
	[0.009]	[0.009]	[0.009]	[0.008]	
Mechanical	0.021**	0.029***	0.028***	0.019**	
	[0.010]	[0.011]	[0.011]	[0.009]	
Mix (analytical + computer + interpersonal)	**0.017*** **	**0.015*** **	**0.001**	**0.014*** **	
	[0.005]	**[0.005]**	**[0.006]**	**[0.005]**	
Mix (routine + computer)	-0.035***	-0.045***	-0.044***	-0.037***	
	[0.008]	[0.008]	[0.008]	[0.007]	
Mix (routine + analytical)	-0.041***	-0.045***	-0.042***	-0.039***	
	[0.007]	[0.008]	[0.008]	[0.007]	
Mix (routine + interpersonal)	0.029***	0.035***	0.033***	0.025***	
	[0.009]	[0.009]	[0.009]	[0.008]	
Worker Skills					
AFQT (analytical)		0.074***	0.073***		-0.009**
		[0.011]	[0.011]		[0.004]
Computer		0.045***	0.044***		0.056***
		[0.006]	[0.006]		[0.002]
Social (interpersonal)		0.016***	0.015***		-0.001
		[0.005]	[0.005]		[0.002]
ASVAB mechanical (routine)		-0.015	-0.014		-0.002
		[0.015]	[0.015]		[0.005]
Mix (AFQT + computer + social)		**0.065*** **	**0.070*** **		**0.135*** **
		[0.017]	**[0.017]**		**[0.009]**
Mix (ASVAB mechanical + computer)		0.029*	0.024		0.038***
		[0.017]	[0.017]		[0.010]
Mix (ASVAB mechanical + AFQT)		0.006	0.007		0.000
		[0.008]	[0.008]		[0.004]
Mix (ASVAB mechanical + social)		-0.039***	-0.038***		0.030***
		[0.008]	[0.008]		[0.004]
Interaction			**0.032*** **		
			[0.008]		
Ethnicity*Gender, Age/Year, Region, Edu FE	X	X	X	X	X
Occupation FE	X	X	X	X	X
Worker FE				X	
Observations	88,391	79,343	79,343	88,391	94,062
R-squared	0.416	0.430	0.431	0.756	0.136

Notes: See Table 1.4 notes.

Table A.9: Robustness Checks of Return to Skill Mixing

Dependent Variable:	ln(hourly wage)		
	(1)	(2)	(3)
Occupation Skills			
Analytical	-0.014*	-0.008	-0.013
	[0.008]	[0.033]	[0.008]
Computer	-0.002	0.069**	-0.038***
	[0.009]	[0.027]	[0.010]
Interpersonal	-0.019**	-0.118***	-0.014*
	[0.008]	[0.030]	[0.008]
Routine	0.026***	0.091***	0.010
	[0.009]	[0.017]	[0.008]
Mix (analytical + computer)	0.007	-0.040	0.020***
	[0.005]	[0.036]	[0.007]
Mix (analytical + interpersonal)	0.010**	0.156***	0.025***
	[0.004]	[0.042]	[0.005]
Mix (computer + routine)	-0.028***	-0.045***	-0.087***
	[0.007]	[0.015]	[0.013]
Mix (computer + interpersonal)	-0.011**	-0.019	-0.021***
	[0.005]	[0.033]	[0.008]
Mix (routine + analytical)	-0.033***	-0.080***	-0.041**
	[0.007]	[0.015]	[0.018]
Mix (routine + interpersonal)	0.010	0.033**	0.026**
	[0.007]	[0.016]	[0.012]
Ethnicity*Gender, Age/Year, Region, Edu FE	X	X	X
Occupation FE	X	X	X
Worker FE	X	X	X
Observations	87,655	87,655	87,655
R-squared	0.757	0.757	0.758

*Notes: This table reports the robustness checks to the results in Table 1.4. Columns (1) and (2) use Absolute Distance and Inverse Herfindahl measures to construct mixing indexes (see online Appendix A.1.7 for details) and Columns (3) and (4) use standardized and broad measures of skills (see online Appendix A.1.6 for details). Log hourly wages are the outcome variables and person-year is the unit of observation. The occupational skill and skill mixing measures come directly from O*NET and are merged to NLSY79&97 based on census occupation codes. All measures of skill and skill mixing are normalized to have mean 0 and standard deviation 1. Ethnicity-by-gender, age, year, census region, urbanicity, and a 5-category (no high-school, high-school graduate, some college, college graduate, post-college) education fixed effects are included for all regressions, with additional fixed effects as indicated in the table. Standard errors are clustered at the individual level. *** p¡0.01, ** p¡0.05, * p¡0.10.*

A.1.9 Additional Results on College Major's Skill Mixing

In online Appendix Table A.10, I list the top majors both in terms of the levels and changes in the degree of skill mixing for different skill pairs. Architecture and Environmental Design stands out as the highest major in mixing the three non-routine skills, followed by Computer and Information Sciences, and Communications. Two other majors: Social Sciences and Agriculture and Natural Resources are among the top majors in mixing routine and non-routine skills.

In Table A.8 column (4), I represent a worker's human capital by the skill content of a worker's accumulated education experience.[3] Such a designation necessarily restricts the analysis to those who have entered college and brings up selection concerns; however, controlling for worker fixed effects and fixed and time-varying occupation attributes, the estimates show whether it is rewarding to studying a more skill-mixed major conditional on one's job choices. The result in column (4) shows a positive return of around 3 percent studying a college major that is associated with a standard deviation higher mixing of non-routine skills. Interestingly, when taking into account the skill mixing of a worker's college major, the wage premium to occupational skill mixing becomes insignificant. This is due to the correlation between the skill mixing of college majors and subsequent occupational choices, and shows that the former plays a more significant role in driving the wage returns.

---

[3] This is determined using rolling averages of skill and mixing measures from the worker's entire educational history, since workers may have studied multiple majors.

Table A.10: Top College Majors in Skill Mixing

Hybrid Index – Level	**Hybrid Index – Change**
analytical + computer + interpersonal	
Physical Sciences	Architecture and Environmental Design
Engineering	Computer and Information Sciences
Letters	Communications
analytical + computer	
Physical Sciences	Interdisciplinary Studies
Engineering	Area Studies
Letters	Computer and Information Sciences
analytical + interpersonal	
Public Affairs and Services	Architecture and Environmental Design
Business and Management	Computer and Information Sciences
Social Sciences	Communications
computer + interpersonal	
Social Sciences	Architecture and Environmental Design
None, General Studies	Computer and Information Sciences
Public Affairs and Services	Engineering
routine + computer	
Transportation	Social Sciences
Fine and Applied Arts	Agriculture and Natural Resources
Engineering	Foreign Languages
routine + analytical	
Transportation	Agriculture and Natural Resources
Health Professions	Social Sciences
Computer and Information Sciences	Foreign Languages
routine + interpersonal	
Transportation	Agriculture and Natural Resources
Health Professions	Architecture and Environmental Design
Military Sciences	Social Sciences

*Notes: This table lists the top 3 college majors for each mixing index both in terms of levels and in terms of changes from 2005 to 2019. To calculate the degree of skill mixing for college majors, I first map the occupation level degree of skill mixing contained in the O*NET data to NLSY, and then calculate for each college major's students, the employment weighted average of skill intensities and mixing indexes of their occupations. I use both NLSY79&97 to get the employment weight on occupations.*

Table A.11: Major Crosswalk

NLSY97 Code (before 2010)	Major Field of Study	NLSY79 Code
0	None, no major yet (didn't/don't) have to declare yet;	.
1	Agriculture/Natural resources	1
2	Anthropology	22
3	Archaeology	22
4	Architecture/Environmental design	2
5	Area studies	3
6	Biological sciences	4
7	Business management	5
8	Communications	6
9	Computer/Information science	7
10	Criminology	22
11	Economics	22
12	Education	8
13	Engineering	9
14	English	15
15	Ethnic studies	3
16	Fine and applied arts	10
17	Foreign languages	11
18	History	22
19	Home economics	13
20	Interdisciplinary studies	49
21	Mathematics	17
22	Nursing	12
23	Other health professions	12
24	Philosophy	15
25	Physical sciences	19
26	Political science and government	21
27	Pre-dental	4
28	Pre-law	14
29	Pre-med	4
30	Pre-vet	4
31	Psychology	20
32	Sociology	22
33	Theology/religious studies	23
36	Nutrition/Dietetics	4
37	Hotel/Hospitality management	5
38	Other - Recoded to Liberal Arts and Sciences	49
39	Other - Recoded to Automobile/Automotive Mechanics Technology/Technician	24
40	Other - Recoded to Human Services, General	21
41	Other - Recoded to Social Work	21
42	Other - Recoded to Electrical/Electronics Maintenance and Repair Technology	24
43	Other - Recoded to Geography	22
44	Other - Recoded to International Relations & Affairs	21
45	Other - Recoded to transportation & materials moving	25
46	Other - Recoded to security and protective services	21
47	Other - Recoded to legal support services	14
48	Other - Recoded to other sciences/applied sciences	49
99	UNCODABLE	99

NLSY97 Code (CM10)	Major Field of Study	NLSY79 Code
1	Agriculture, agriculture operations, & related sciences	1
3	Natural resources and conservation	1
4	Architecture and related services	2
5	Area, ethnic, cultural, gender, and group studies	3
9	Communications, journalism, and related programs	6
10	Communications technologies/technicians & support services	6
11	Computer & information sciences & support services	7
12	Personal and culinary services	49
13	Education	8
14	Engineering	9
15	Engineering technologies & engineering-related fields	9
16	Foreign languages, literatures, and linguistics	11
19	Family and consumer sciences/human sciences	13
22	Legal professions and studies	14
23	English language and literature/letters	15
24	Liberal arts and sciences, general studies & humanities	49
25	Library science	16
26	Biological and biomedical sciences	4
27	Mathematics and statistics	17
28	Military science, leadership, and operational art	18
29	Military technologies and applied sciences	18
30	Multi/interdisciplinary studies	49
31	Parks, recreation, leisure, and fitness studies	21
32	Basic skills development/remedial education	8
33	Citizenship activities	21
34	Health-related knowledge and skills	12
35	Interpersonal and social skills	6
36	Leisure and recreational activities	49
37	Personal awareness and self-improvement	8
38	Philosophy and religious studies	15
39	Theology and religious vocations	23
40	Physical sciences	19
41	Science technologies/technicians	24
42	Psychology	20
43	Homeland security, law enforcement, firefighting, and related protective services	18
44	Public administration and social service professions	21
45	Social sciences	22
46	Construction trades	24
47	Mechanic and repair technologies/technicians	24
48	Precision production	24
49	Transportation and materials moving	25
50	Visual and performing arts	10
51	Health professions and related programs	12
52	Business, management, marketing, & related support services	5
53	High school/secondary programs and certificates	8
54	History	22
60	Residency programs	12
999	Uncodable	99

NLSY79 Code	Major Field of Study
0	None, General Studies
1	Agriculture and Natural Resources
2	Architecture and Environmental Design
3	Area Studies
4	Biological Sciences
5	Business and Management
6	Communications
7	Computer and Information Sciences
8	Education
9	Engineering
10	Fine and Applied Arts
11	Foreign Languages
12	Health Professions
13	Home Economics
14	Law
15	Letters
16	Library Science
17	Mathematics
18	Military Sciences
19	Physical Sciences
20	Psychology
21	Public Affairs and Services
22	Social Sciences
23	Theology
24	Mechanics
25	Transportation
49	Interdisciplinary Studies
99	Other

APPENDIX B

APPENDIX FOR "A MULTI-DIMENSIONAL SKILL DIRECTED SEARCH MODEL WITH OCCUPATION DESIGN"

B.1 THOERY AND QUANTITATIVE

B.1.1 Propositions and Proofs

Lemma 1 *An occupation* $\mathbf{y}^j = \{y_1^j, ..., y_k^j, \ldots, y_K^j\} \in S \subset \mathbb{R}^{K+}$ *within a closed skill space* S *of dimension* K *is more mixed in skills based on Definition 1 if for any pair of skills* (h, k)*, the ratio of* $\frac{y_h}{y_k}$ *becomes closer to 1.*

Proof of Lemma 1: For the occupation $\mathbf{y}$ we want to establish how the degree of skill mixing changes if the skill dimensions for j and k are to vary. The lemma can be simply proved by considering the skill mixing index for this occupation. Let $y_k = ry_h$ and denote y_h by y, the mixing index for $\mathbf{y}$ is:

$$\frac{y + ry + A}{\sqrt{K}\sqrt{y^2 + r^2y^2 + B}},$$

where A and B are two constants that do not depend on y_k and y_h. The above equation is maximized at $r = 1$. Therefore, for any y_h, the occupation is more skill-mixed if the ratio r is close to 1. This completes the proof. *Q.E.D.*

Proposition 1 (Changes in Skill Mixing) *Consider an occupation* $\mathbf{y}^j = \{y_1^j, ..., y_k^j, \ldots, y_K^j\} \in S \subset \mathbb{R}^{K+}$ *within a closed skill space* S *of dimension* K*. Assume that firms operate the occupation with a production technology as described by equation (2.1) and under an* *occupation operation cost defined by equation (2.5). Under these conditions, occupation* $\mathbf{y}^j$ *will show an increased degree of skill mixing given the following conditions:*

(i) The skills within the vector $\mathbf{y}^j$ *demonstrate a rise in complementarity in production (a decrease in* σ*), provided that* σ *does not undergo a change in sign.*

(ii) The skills within the vector $\mathbf{y}^j$ *exhibit an higher increasing marginal cost (an increase in* ρ*), under the condition that* $\rho > \sigma^j$*.*

(iii) Additionally, occupation $\mathbf{y}^j$ *will exhibit a increased degree of skill mixing in the* (y^j_k, y^j_h) *dimension if the ratio between* (x_k, x_h) *approaches unity.*

Proof of Proposition 1: Lemma 1 posits that an occupation $\mathbf{y}^j$ exhibits greater skill mixing if the ratio across all skill dimensions approximates 1. Therefore, establishing the influence of the ratio on the degree of skill mixing suffices. The initial step concentrates on any two skills within the vector (y^j_k, y^j_h). I subsume occupation superscript j in the proof below.

The firm value function indicates that the firm re-optimizes the choice of $\mathbf{y}$ in each period. Consequently, within a given submarket at a particular time instance $(\mathbf{x}, \mathbf{y}, \omega)$, the firms' choices of $\mathbf{y}$ remain uninfluenced by the continuation value, rendering it a static problem. Time subscript is subsumed in the subsequent proof.

By deriving the first-order condition of firms' optimization problems in the submarket $(\mathbf{x}, \mathbf{y})$ and taking ratios, one obtains the following condition: $\frac{y_h}{y_k} = (\frac{x_h}{x_k})^{\frac{\sigma}{\rho-\sigma}}(\frac{\alpha_h}{\alpha_k})^{\frac{\sigma}{\rho-\sigma}}$. Therefore, the ratio of firms' optimal skill requirement choices for any two skills (y_h, y_k) is influenced by four variables: the elasticity parameter of substitution in production σ, the degree of increasing marginal occupation operation cost ρ, the ratio of worker skills in the submarket (x_h, x_k), and the ratio of skill efficiencies (α_h, α_k).

From the equation, it is evident that as σ decreases, indicating an increase in skill complementarity in production, $\frac{y_h}{y_k}$ will converge to 1 for any two skills (y_h, y_k), under the assumption that σ does not change sign. Similarly, as ρ increases, $\frac{y_h}{y_k}$ will approximate 1 for any two skills (y_h, y_k), given that $\rho - \sigma$ does not change sign.

The influence of worker skill bias on the degree of skill mixing of $\mathbf{y}$ presents a more complex scenario, as a change in the ratio $\frac{x_h}{x_k}$ does not directly imply a change in the ratio of other skill pairs. Consequently, to gauge its impact on the degree of skill mixing, the focus must remain on the (y_h, y_k) dimension. For this specific dimension, if $\frac{x_h}{x_k}$ converges to 1, then $\frac{y_h}{y_k}$ also approaches 1. *Q.E.D.*

Proposition 2 (Changes in Wage and Job Finding) *Consider an occupation* $\mathbf{y} \in S \subset \mathbb{R}^K$ *within a closed skill space* S *of dimension* K*. Assume that firms operate the occupation with a production technology as described by equation (2.1) and under an occupation operation cost defined by equation (2.5). Also, these firms offer an output share* ω *to workers and have value functions described by equation (2.3). Further, let worker value be described by equation (2.2). Under these conditions, workers in occupation* $\mathbf{y}$ *will earn a higher wage, and unemployed workers will have a higher job-finding probability under conditions (i) and (ii) of Proposition 1*

Proof of Proposition 2:

Wages: To establish the change in wages, one needs to show that the output of the worker-firm match increases as the elasticity parameter σ decreases and approaches 0 from 1, or if σ decreases in the negative range, consistent with skills becoming more complementary in production. At a particular output share rate ω, such value changes of σ will lead to higher wages.

Now, let us obtain the first derivative of σ for the production function 2.1. WLOG, let's consider the case of two skills, and express y_1x_1 and y_2x_2 as m and n. The output of a worker-firm match can be expressed as $q = (m^\sigma + n^\sigma)^{1/\sigma}$. We can take log of the production function $\ln(q) = \frac{1}{\sigma}\ln(m^\sigma + n^\sigma)$ and then take logarithmic differentiation that gives the following:

$$\frac{1}{q}\frac{\partial q}{\partial \sigma} = -\frac{1}{\sigma^2}\ln(m^\sigma + n^\sigma) + \frac{1}{\sigma}\frac{1}{m^\sigma + n^\sigma}(m^\sigma \ln(m) + n^\sigma \ln(n))$$

Solving for $\frac{\partial q}{\partial \sigma}$ gives:

$$\frac{\partial q}{\partial \sigma} = q[-\frac{1}{\sigma^2}\ln(m^\sigma + n^\sigma) + \frac{1}{\sigma}\frac{1}{m^\sigma + n^\sigma}(m^\sigma \ln(m) + n^\sigma \ln(n))]$$

$$\frac{\partial q}{\partial \sigma} = q[-\frac{1}{\sigma}\ln(q) + \frac{1}{\sigma}q^{-\sigma}(m^\sigma \ln(m) + n^\sigma \ln(n))]$$

In the case of the calibration of the model, since m, n, and y are all in the range of $[0, 1]$, one can show that the above derivative is negative when $0 < \sigma < 1$ or when $\sigma < 0$.

With respect to (ii) of Proposition 1, it is easy to see that since for the analysis of this paper, both $(\mathbf{x}, \mathbf{y})$ are in the range $[0, 1]$, therefore the occupation operation cost is decreasing in ρ, so wage should increase as marginal cost increases.

Employment: For job finding probability, it suffices to show that $p(\theta_t(\mathbf{x}, \mathbf{y}, \omega))$ is increasing in σ and ρ. This becomes simpler since the above proof establishes that worker-firm output is increasing in both σ and ρ, and so does the firm's value $J(\mathbf{x}, \mathbf{y}, \omega)$. By the free entry condition in equation (2.4), at a fixed vacancy posting cost, an increase in $J(\mathbf{x}, \mathbf{y}, \omega)$ implies a decrease in $q(\theta_t(\mathbf{x}, \mathbf{y}, \omega))$ and therefore implies an increase in $p(\theta_t(\mathbf{x}, \mathbf{y}, \omega))$ under constant return to scale matching technology. *Q.E.D.*

B.1.2 Equilibrium Definition and Block Recursivity

In this section, I define a block-recursive equilibrium (BRE) for the economy following Menzio and Shi (2011). I further show that the equilibrium of the economy is unique and is block-recursive.

Definition 3 (Block-recursive Equilibrium) *Let $\psi \in \Psi$ be the aggregate state of the economy, which is a distribution of agents across employment status $e = U, W$, skill profiles $\mathbf{x}$, occupational skill requirements $\mathbf{y}$, and output shares ω.*

A block-recursive equilibrium for this economy consists of value functions for both unemployed and employed workers $U(\mathbf{x}) : S \to \mathbb{R}$, $W(\mathbf{x}, \mathbf{y}, \omega) : S \times S \times [0,1] \to \mathbb{R}$, and their respective policy functions $y'_U(\mathbf{x}) : S \to S \times [0,1]$, $y'_W(\mathbf{x}, \mathbf{y}, \omega) : S \times S \times [0,1] \to S \times S \times [0,1]$; firms' policy function $J(\mathbf{x}, \mathbf{y}, \omega) : S \times S \times [0,1] \to \mathbb{R}$ and corresponding policy function $y'_J(\mathbf{x}, \mathbf{y}, \omega) : S \times S \times [0,1] \to S \times S \times [0,1]$; labor market tightness $\theta(\mathbf{x}, \mathbf{y}, \omega) : S \times S \times [0,1] \to \mathbb{R}_+$; and aggregate state $\psi \in \Psi$ such that:

1. *The worker's value functions $U(\mathbf{x})$ and $W(\mathbf{x}, \mathbf{y}$ satisfy (2.2) for all states $\psi \in \Psi$ and $y'_U(\mathbf{x})$, $y'_W(\mathbf{x}, \mathbf{y}, \omega)$ are the associated policy functions respectively*
2. *Firms' value function $J(\mathbf{x}, \mathbf{y}, \omega)$ satisfy (2.3) for all states $\psi \in \Psi$ and $y'_J(\mathbf{x}, \mathbf{y}, \omega)$ is the associated policy function*
3. *The labor market tightness $\theta(\mathbf{x}, \mathbf{y}, \omega)$ in each submarket $(\mathbf{x}, \mathbf{y}, \omega)$ for all states $\psi \in \Psi$ is consistent with free-entry condition in equation (2.4)*

From the above definition of block-recursive equilibrium agents' value functions and policy functions, as well as the market tightness are independent of the aggregate state, only requiring that they are consistent with the aggregate state distribution of agents. Such an equilibrium is easier to characterize analytically and solve numerically. Note a key difference between the above definite of BRE and the one defined in Menzio and Shi (2011). In the economy studied in this paper, because I use the model to study the steady-state equilibrium, the value functions, policy functions, and market tightness are entirely independent of the aggregate state. Whereas Menzio and Shi (2011) studies out-of-steady-state dynamics, the value functions, policy functions, and market tightness still depend on the aggregate productivity shocks but are independent of the distribution of agents across employment status and match-specific shocks.

Now, I show that a block-recursive equilibrium exists and is unique.

Proposition 3 (Existence and Uniqueness of BRE) *Under the model specification of linear utility and invertible and constant returns to scale matching function, also assume that the support for worker and occupation skill profiles S has bounded, then: i) all equilibria are block recurve as defined in definition 3; ii) there exists a unique block-recursive equilibrium.*

Proof of Proposition 3:

The proof first establishes the uniqueness of value functions (U, W, J), as well as policy functions and market tightness $(y'_U, \omega'_U, y'_W, \omega'_W, y'_J, \theta)$; then, the proof establishes their independence from the aggregate state.

Uniqueness: I first show that the value functions for workers and firms as defined in equation (2.2) and (2.3) are contractions. Let $\Theta = S \times S \times [0, 1]$, which is bounded based on the assumption that S is bounded. Let $B(\Theta)$ the space of bounded functions $V : \Theta \to \mathbb{R}$ and the operator associated with the worker or firm value functions denoted by $T : B(\Theta) \rightarrow B(\Theta)$. It is straightforward to verify that T satisfies monotonicity and discounting properties:

1. (monotonicity) For $V, V' \in B(\Theta)$, $V \leq V'$ implies $T(V) \leq T(V')$
2. (discounting) For $V \in B(\Theta)$ and $\epsilon > 0$, $T(V + \epsilon) =$

The above conditions establish that the operator T associated with either firm or worker values functions is a contraction bunder Blackwell's sufficient conditions. Therefore, the optimal values workers and firms obtain through dynamic optimization problems are unique.

Next, I show that the policy functions and market tightness are also unique. Since the optimal values firms and workers obtain for their dynamic optimization problems (2.2) and (2.3) is unique, the associated policy functions $(y'_U, \omega'_U, y'_W, \omega'_W, y'_J)$ are also unique due to concavity of the production function defined in equation (2.1) and workers have linear utility over consumption. To show the uniqueness of market tightness, first note that since it is assumed that the matching function is invertible, one may directly obtain market tightness through the market clearing condition (2.4) with $\theta > 0$. The uniqueness of θ then follows from the uniqueness of firms' value function.

Independence of Aggregate State: In the model economy, workers with different skill profiles $\mathbf{x}$ search in their own market, and firms with different skill requirements $\mathbf{y}$ post jobs in these separated markets, therefore, one can establish that the value functions of firms, workers and the market tightness are all independent of the aggregate state ψ. I establish this argument more rigorously through a backward induction argument as in Braxton and Taska (2021). For this purpose, I introduce back time subscript in the notation.

At the terminal period $t = T$, for an employed worker, the continuation value is zero for $T + 1$ onward, so the worker's dynamic programming problem does not depend on the aggregate distribution across states, and is equal to the worker's share of output $W_T(\mathbf{x}, \mathbf{y}, \omega) = \omega f(\mathbf{x}, \mathbf{y})$.

Similarly, the firm's value function also remains independent of the aggregate distribution $J_T(\mathbf{x}, \mathbf{y}, \omega) = (1 - \omega) f(\mathbf{x}, \mathbf{y})$. As a result, through the free entry condition in equation (2.4), the market tightness $\theta_T(\mathbf{x}, \mathbf{y}, \omega)$ is also independent of the aggregate distribution.

Firms at $T - 1$ make occupation design choices $\mathbf{y}$ to solve the firm dynamic programming problem in equation (2.3); workers at $T - 1$ make labor market search choices over occupations $\mathbf{y}$ to solve the worker dynamic programming problem in equation (2.2); As long as $\mathbf{y}$ is within a bounded interval, the extreme value theorem assures at least one solution to this problem. This process is repeated stepping back from $t = T - 1, ..., 1$, which completes the proof. *Q.E.D.*

B.1.3 Identification of Parameters

I begin by estimating the elasticity parameters in production and occupation operation cost, denoted by σ and ρ. As highlighted by Caselli and Coleman (2006), the challenge arises when allowing for the endogenous choice of the efficiency of inputs under constraints, as the elasticity parameters cannot be separately identified. To overcome this challenge, I estimate σ using the *relative wage within occupation* instead of relying on absolute wage levels.

Specifically, based on the model, the wage that workers receive per period is given by the share ω of the output of the worker-firm match, reduced by the occupation design cost, formulated as $w(\mathbf{x},\mathbf{y}) = \omega f(\mathbf{x},\mathbf{y}) - C(\mathbf{y})$. Consequently, within each occupation, the difference in wage relative to a base worker type $\Delta w(\mathbf{x},\mathbf{y})$ can be articulated as follows:

$$\Delta w(\mathbf{x},\mathbf{y}) = \omega \left[\sum_{k=1}^{K} (x^k y^k)^{\sigma} \right]^{\frac{1}{\sigma}} - A, \tag{B.1}$$

where A is occupation-specific and does not depend on the cost parameter τ or ρ. This formulation enables the identification of σ independent of the cost parameters. To carry out the estimation equation (B.1), I first adjust the wage for occupation fixed effects in order to account for A and ω. Next, I compute the within-occupation difference of the adjusted wage relative to the lowest skill type worker.[1] Last, I target the correlation between this adjusted within-occupation relative wage and worker abilities $\mathbf{x}$.[2]

I now turn to the identification of the cost parameters ρ and τ. To begin with, note that the first-order condition of firms' optimization problems in the

---

[1] Refer to Section 2.4 for an in-depth discussion on how worker skill types are calibrated.

[2] According to equation (B.1), σ can be identified from the correlation of any skill with the adjusted wage, which is what I use as the target.

submarket $(\mathbf{x}, \mathbf{y})$ can be simplified in ratios to $\frac{y_h}{y_k} = \left(\frac{x_h}{x_k}\right)^{\frac{\sigma}{\rho-\sigma}}$, a relationship that exclusively depends on the parameters σ and ρ. With σ already estimated, I then target the skill ratio y_j/y_k, which aligns with the moment of the degree of hybridization of occupations. Further, for employed workers, the distribution of employment across various occupations is governed by wages $w(\mathbf{x}, \mathbf{y})$. Given the parameters described above, this functional relationship allows the estimation of τ.

Lastly, given the calibrated unemployment benefits b, the parameters of the matching, production and cost functions, equation (2.2) reveals that the probability of exiting unemployment only depends on the vacancy posting cost. By targeting unemployment level, c is identified.

B.1.4 Calibration of Skill Supply

I carry out the calibration of two key aspects of skill supply variation: the Markov probability of worker skill adjustment in a steady state equilibrium and the variation in worker skill supply spanning two data periods that the model aims to align with two steady states. I will first delve into the details of the skill variation between data periods and then explore the skill evolution within a model period as guided by the Markov process, following the approach of Lise and Postel-Vinay (2020).

Across-period Skill Supply Variation: Considering the potential influence of skill supply variation on skill mixing, I calibrate the model to reflect workers' choices in occupation, college major (if attended), and employment status, in line with the approach of Lise and Postel-Vinay (2020). This calibration introduces variation in worker skill supply across two periods. Worker skills are adjusted based on the requirements of an occupation or a college major; they increase if the requirements exceed the original skills and decrease if the requirements are lower or if the worker is unemployed. The speed of this adjustment is asymmetric and skill-specific.

Specifically, following the estimates from Lise and Postel-Vinay (2020), as presented in online Appendix Table B.1, a worker's skills accumulate at a rate of γ_j times the gap between the worker's skill j and the occupation's requirement for that skill each year. The value of γ_j depends on whether it relates to learning or depreciation (upward or downward accumulation). Additionally, workers can lose skills when not employed, with unemployment treated as requiring a zero level for all skills. However, I specify such that a worker's skill level cannot fall below their initial endowments. For changes in skills while in school, I specify that workers spend an average of three years learning the skills of their majors.

I incorporate two modifications into this framework. First, since Lise and Postel-Vinay (2020)'s estimates are based on weekly data, I adjust them by multiplying by the number of working weeks, set at 47. Second, I align Lise and Postel-Vinay (2020)'s estimates of cognitive, interpersonal, and manual skills with my analysis's categories of analytical, interpersonal, and routine skills.[3] Since Lise and Postel-Vinay (2020)'s estimates do not include computer skills, I use their cognitive skill estimates as a proxy.

In calculating the skill adjustment, I first standardize both worker skills and

---

[3] Their exclusion restriction imposes that (i) the ASVAB mathematics knowledge score only reflects cognitive skills; (2) the ASVAB automotive and shop information score only reflects manual skills; (3) the Rosenberg self-esteem score only reflects interpersonal skills.

occupation skill requirements. Then, for example, if a worker is employed in an occupation that requires a standard deviation higher in analytical skill compared to the worker's analytical skill, the worker will accumulate 0.36 standard deviations of analytical skill in a year due to learning on the job. Conversely, if a worker's interpersonal skill is higher than required, it will decrease by only 3×10^4 standard deviations, almost remaining unchanged, as interpersonal skills are estimated to be very hard to lose.

Markov Skill Supply Adjustment: I now discuss the Markov process of skill adjustment. Specifically, considering each skill j in the worker's skill profile $\mathbf{x}$ as an element of the finite set S, the evolution of this skill follows a Markov process $\pi(x_j'|x_j, y_j)$, conditional on the worker's current skill level and employed occupation. If a worker is matched with an occupation that requires a skill level exceeding his or her own ($x_j < y_j$), the worker's skill j will adjust upward in the next period: $x_j' > x_j$, and the inverse applies for a worker whose skill is lower than the requirements of their current occupation.

The calibration of the Markov adjustment probability is conducted in a similar manner to that of the across-period skill supply variation. The annual adjustment rates for different skills *gamma* represent the rate at which worker skills approach occupation skill requirements, and it is regarded as the probability that a worker's skill j will adjust to the corresponding value.

The key challenge in this calibration process arises when quantifying the model: both worker skill and occupation skill requirements are discretized as grid values. To accommodate this discretization, the probability that a worker moves up or down a grid for skill j based on the occupation is scaled as below.

The Markov probability of upward adjustment is determined by:

$$\frac{x_j^{up} - x_j}{y_j - x_j} \mathbf{1}(x_j^{up} < y_j) \times \gamma_j^{up}$$

Similarly, the Markov probability of downward adjustment is given by:

$$\frac{x_j^{down} - x_j}{y_j - x_j} \mathbf{1}(y_j < x_j^{down}) \times \gamma_j^{down}$$

Here, x_j represents the current grid value of worker skill j, while x_j^{up} or x_j^{down} denotes the value of worker skill j up or down a grid, respectively. The indicator variables $\mathbf{1}(y_j < x_j^{down})$ or $\mathbf{1}(x_j^{up} < y_j)$ evaluates whether the skill j grid value of the worker's current employed occupation is greater or smaller than the value of the worker's skill j grid. This means that a worker will only adjust up or down a grid if the occupation's skill is larger or smaller than the corresponding up or down grid value for the worker's skill. This process specifies the interplay between skill adjustment and occupation requirements and allows for a precise calibration within the model's framework.

Table B.1: Annual Skill Learning and Depreciation Rate

O*NET Measure	NLSY Measure	γ_{school}^{learn}	γ_j^{up}	γ_j^{down}
analytical	AFQT score	0.33	0.36	0.10
interpersonal	Deming (2017) social skill	0.33	0.05	0.00003
routine	ASVAB	0.33	1	0.36
computer	OCC/Major's 2005 Value	0.33	0.36	0.10

*Notes: This table illustrates for each O*NET skill measure, its corresponding skill measure using NLSY79&97 data, and the learning and depreciation rate for these different skills. The AFQT is the same as the one used by Altonji, Bharadwaj, and Lange (2012) followed by Deming (2017), which controls for age-at-test, test format, and other idiosyncrasies. Deming (2017)'s social skill measure consists of sociability in childhood and sociability in adulthood in NLSY79, and two questions from the Big 5 inventory gauging the extraversion in NLSY97. The average of workers' ASVAB mechanical orientation and electronics test scores are used for mechanical skill. Since ASVAB scores are not available for the NLSY97 survey, they are imputed based on predictive regression using the NLSY79 survey. Workers' occupations' or college majors' O*NET computer skill scores in the year 2000 are used as their endowed computer skill. The skill accumulation/depreciation rate is directly from Lise and Postel-Vinay (2020)'s estimates based on monthly data converted to annual values. Skill learning/depreciating while attending college is specified to be 33% per year.*

B.1.5 Algorithm and Solution Method

The quantitative method used for estimation is SMM. Given the parameters in the model that are internally estimated $\Theta = \{\sigma, \rho, \tau, c, \alpha_k\}$, each iteration of SMM first solves the steady state firm and worker policy function, after which a panel of worker is simulated to obtain the equilibrium distribution of labor market outcomes.

Specifically, to find the steady state policy of agents, I use value function iteration:

1. Fix the number of periods T
2. Starting from the terminal period T, solve the firm problem as in equation (2.3)

3. Use the free entry condition in equation (2.4) to obtain the market tightness $\theta_T(\mathbf{x}, \mathbf{y}, \omega)$
4. With the market tightness, solve the worker dynamic programming problem in equation (2.2)
5. Repeated stepping back from $t = T - 1, ..., 1$
6. Check if the difference in worker value $U_{t+1} - U_t$, $W_{t+1} - W_t$ and the firm value $J_{t+1} - J_t$ is less than a predetermined tolerance level. If yes stop, if not increase T and go back to the first step

Next, I simulate 10,000 workers to obtain a distribution of labor market outcomes across different occupations and worker types. Finally, the SMM procedure minimizes the Euclidean distance between the model-implied moments and the data moments.

B.1.6 Additional Counterfactual Results

Figure B.1: Model Counterfactual

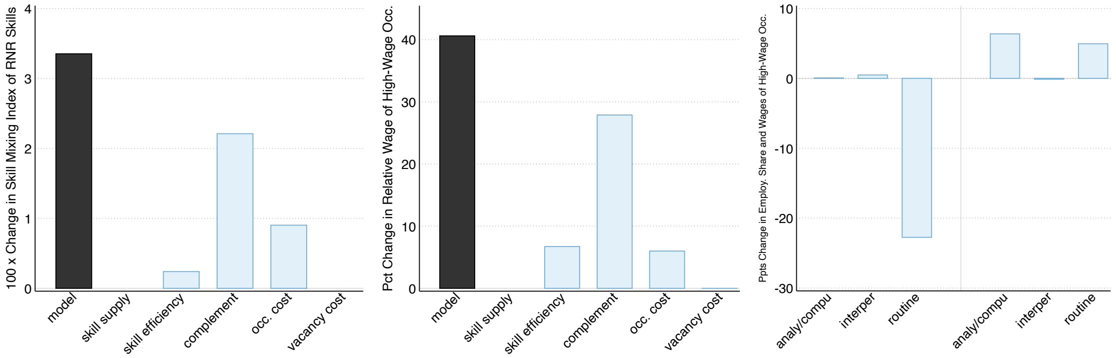

(1) Skill mixing

(2) Wages – unweighted

(3) Wages (left) and employment (right) by individual skills

Notes: These figures plot the model generated changes in skill mixing in high-skill occupations (column 1), changes in wages unweighted by employment (column 2), and changes in wages and employment from individual skills (column 3). Different model channels are shut down individually by eliminating the relative calibrated values to highlight the contribution of each channel. The full model has all the model features. The values of skill complementarity in production, cost of skills in occupation operation, efficiency differential, and vacancy posting cost across the two periods are shown in Table 2.2. Worker skill supply distribution variation across the periods are calibrated according to Table B.1. Panel (3) and (4) depict the model generated changes in skill mixing in low-skill occupation and the relative wage of high-skill occupations by shutting down the skill efficiency differential for analytical/computer, interpersonal, and routine skills individually; also by shutting down τ and ϕ individually.

APPENDIX C

APPENDIX FOR "E-COMMERCE AND REGIONAL INEQUALITY: A TRADE FRAMEWORK AND EVIDENCE FROM AMAZON'S EXPANSION"

C.1 Derivation of Demand Function

Proof of Theorem 1: In a sequential ordered search model, consumers in region n optimally choose or purchase a good from sector j at retailer i where $\omega^j_{ni} - p^j_{ni}$ is maximized. Denoting this demand as D^j_{ni}, it can be expressed as $D^j_{ni} = P(\omega^j_{ni} - \ln p^j_{ni} > \max_g \ \omega^j_{ng} - \ln p^j_{ng}) = \int \Pi_{g \neq i} F_{\omega^j_{ng}}(\epsilon - \ln p^j_{ng}) f_{\omega^j_{ni}}(\epsilon - \ln p^j_{ni}) d\epsilon$. This demand D^j_{ni} equates to a discrete choice model with indirect utility $v^j_{ni} = -\ln p^j_{ni} + \epsilon^{j,DC}_{ni}$ if $F_{\omega^j_{ni}} = F_{\epsilon^{j,DC}_{ni}}$, where $\epsilon^{j,DC}_{ni}$ is the random utility a consumer derives from the retailer.

To transition from a discrete choice model to CES demand, we note that the average ϵ^j_{ni} is zero for brick-and-mortar stores and $\ln(\mu)$ for online retailers. Therefore, we can express $\epsilon^{j,DC}_{ni}$ as $\ln(\mu) + \chi^j \tilde{\epsilon}^j_{ni}$ where $\tilde{\epsilon}^j_{ni}$ has mean zero and unit variance, and χ^j is the variance of the effective match value ω^j_{ni}, assumed to vary across sectors but not regions. The demand then becomes $D^j_{ni} = \int \Pi_{g \neq i} F_{\epsilon^{j,DC}_{ng}}(\epsilon - \ln p^j_{ng}) f_{\epsilon^{j,DC}_{ni}}(\epsilon - \ln p^j_{ni}) d\epsilon$.

Assuming $F_{\omega^j_{ni}} = F_{\epsilon^{j,DC}_{ni}}$ follows an extreme type I distribution, the demand for retailer i if i is an online retailer becomes

$$D^j_{ni} = \frac{(p^j_{ni}/\mu)^{\frac{-1}{\chi^j}}}{\sum_{g=1}^{N} (p^j_{ng}/\mu)^{\frac{-1}{\chi^j}} + (p^j_{n0})^{\frac{-1}{\chi^j}}}.$$

If i is brick-and-mortar, then

$$D_{ni}^{j} = \frac{p_{n0}^{j})^{\frac{-1}{\chi^{j}}}}{\sum_{g=1}^{N}(p_{ng}^{j}/\mu)^{\frac{-1}{\chi^{j}}} + (p_{n0}^{j})^{\frac{-1}{\chi^{j}}}}.$$

Denote the elasticity of substitution among retailers by σ_j, then $\sigma_j = \frac{1+\chi^j}{\chi^j}$. This demand function leads to sector j's demand as $C_n^j = \left[(c_{n0})^{\frac{\sigma-1}{\sigma}} + \mu \sum_{i=1}^{N}(c_{ni})^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma_j}{\sigma_j - 1}}$. Given that the consumer's expenditure share is controlled by η^j in a Cobb-Douglas manner, the final demand function is $C_n = \Pi_{j=1}^{J}(C_n^j)^{\eta^j}$.

C.2 Comparative Statics in Hat Algebra

Comparative Statics. Computing the equilibrium outcomes out of the model requires solving a system of nonlinear equations (3.10), (3.11), (3.3), (3.4), and (3.13) to (3.17), which requires pinning down the levels of a large number of fundamentals and parameters. To ease the comparative statics analysis, I adopt the "exact hat algebra" method (Dekle, Eaton, and Kortum 2008) to characterize the equilibrium variables and solve for the economy in proportional changes, which greatly reduces the number of fundamentals and parameters to identify. Specifically, define $\hat{x} \equiv x'/x$ the relative change of any variable from its original to counterfactual equilibrium values, x and x' respectively. Since e-commerce shocks function in three channels relating to search and transportation frictions and capital capacity, proportional changes in these fundamentals can be expressed as $\hat{\mu}_{ni}^{j}$, $\hat{\kappa}_{ni}^{R}$, and $\hat{\rho}_{n}^{j}$. The equilibrium in relative changes under the e-commerce shock can be characterized by the following equations.

The share of labor in different sectors is given by:

$$\hat{\pi}_n^0 = \frac{\hat{A}_n^0(\hat{w}_n^0)^{\nu_n}}{\hat{\Phi}_n}, \ \hat{\pi}_n^{j,K} = \frac{\hat{A}_n^{j,K}(\hat{w}_n^{j,K})^{\nu_n}}{\hat{\Phi}_n}, \text{ where } \hat{\Phi}_n = \sum_{h=0}^{J}\sum_{K=M,R}\pi_n^{K,h}\hat{A}_n^{K,h}(\hat{w}_n^{K,h})^{\nu_n}. \tag{C.1}$$

The input costs are given by:

$$\hat{c}_n^{j,M} = \hat{\omega}_n^{j,M}, \quad \hat{c}_n^{j,R} = (\hat{\rho}_n^{j,R}\hat{\omega}_n^{j,R})^{\gamma_n^j}(\hat{P}_n^{j,M})^{1-\gamma_n^j}, \tag{C.2}$$

$$\text{where } \hat{\omega}_n^{j,K} = \hat{w}_n^{j,K}(\hat{l}_n^{j,K})^{\beta_n} = (\hat{w}_n^{j,K})^{1+\beta_n}(\hat{\pi}_n^{j,K})^{\frac{(\nu_n-1)\beta_n}{\nu_n}},$$

$$\text{and } \hat{P}_n^{j,M} = \left(\sum_{i=1}^{N} x_{ni}^{j,M}(\hat{\kappa}_{ni}^{M}\hat{c}_i^{j,M})^{-\theta^j}\hat{T}_i^j\right)^{\frac{-1}{\theta^j}}$$

The trade shares are given by:

$$x_{ni}^{\prime\, j,M} = x_{ni}^{j,M}\left(\frac{\hat{\kappa}_{ni}^{M}\hat{c}_i^{j,M}}{\hat{P}_n^{j,R}}\right)^{-\theta_j}\hat{T}_i^j, \quad x_{ni}^{\prime\, j,R} = x_{ni}^{j,R}\left(\frac{\hat{\kappa}_{ni}^{R}\hat{c}_i^{j,R}}{\hat{\mu}_{ni}^{j}\hat{P}_n^{j,R}}\right)^{1-\sigma^j}, \tag{C.3}$$

$$\text{where } \hat{P}_n^{j,R} = \left(\sum_{i=1}^{N} x_{ni}^{j,R}\left(\frac{\hat{\kappa}_{ni}^{R}\hat{c}_i^{j,R}}{\hat{\mu}_{ni}^{j}}\right)\right)^{\frac{1}{1-\sigma^j}}.$$

Market clearing conditions now become:

$$X_n^{\prime\, j,R} = \sum_{i=1}^{N} x_{in}^{\prime\, j,R}\eta^j\left[\sum_{k=0}^{J}\sum_{K=M,R}\left(\frac{1}{1-\beta_i}\right)\hat{\rho}_i^{K,k}\hat{w}_i^{K,k}\hat{l}_i^{K,k}\rho_i^{K,k}w_i^{K,k}L_i^{K,k} - \Omega_i\right], \tag{C.4}$$

$$X_n^{\prime\, j,M} = \sum_{i=1}^{N}(1-\gamma_i^j)x_{ni}^{\prime\, j,M}X_n^{\prime\, j,R}, \tag{C.5}$$

$$\hat{w}_n^{j,M}\hat{l}_n^{j,M}w_n^{j,M}L_n^{j,M} = \beta_n\hat{X}_n^{j,M}, \quad \hat{w}_n^{j,R}\hat{l}_n^{j,R}w_n^{j,R}L_n^{j,R} = \frac{1}{\hat{\rho}_i^{j,R}}\gamma_n^j\beta_n\hat{X}_n^{j,R} \tag{C.6}$$

Equations (C.1)-(C.4) from above illustrate that given the e-commerce shock $(\hat{\mu}_{ni}^j, \hat{\kappa}_{ni}^R, \hat{\rho}_n^j)$, solving for the equilibrium in proportional changes only requires information on initial allocations $(x_{ni}^{j,K}, X_{ni}^{j,K}, K = \{M, R\})$, value-added and capital capacities $(w_n^{j,K}, L_n^{j,K}, \rho_n^{j,K}, K = \{M, R\})$, exogenous trade deficits (Ω_n), as well as parameters with respect to value-added shares (β_n and γ_n^j), consumption shares (η_n^j), and trade elasticities (σ^j and θ^j). All other equilibrium variables, economic fundamentals, and parameters turn out to be irrelevant for computing real wage changes – this significantly reduces the estimation burden of conducting counterfactual analysis of the e-commerce shock.

C.3 Alternative Modeling Details

The price index P_n^j for sector j in region n is a function of the aggregated price levels of imports from other regions and the local price level for brick-and-mortar (BM) stores. It integrates over all possible productivity levels z above a certain threshold $\bar{z}_m^j$, weighted by the productivity distribution $G(z)$, and sums up contributions from all other regions m to region n. The equation is expressed as:

$$P_n^j = \left[\sum_{m=1}^{N} Y_m \int_{\bar{z}_m^j} \left(\tilde{\sigma} \frac{\left(w_m^{j,R}\right)^{\gamma^j} \left(P_m^{j,M}\right)^{\left(1-\gamma^j\right)} k_{nm}^R}{\mu z_m^j} \right)^{1-\sigma} dG(z) + \left(\left(\omega_n^{j,B}\right)^{\gamma^j} \left(P_n^{j,M}\right)^{\left(1-\gamma^j\right)} \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}$$

$$= \left[\sum_{m=1}^{N} Y_m \left(\tilde{\sigma} \left(w_m^{j,R}\right)^{\gamma^j} \left(P_m^{j,M}\right)^{\left(1-\gamma^j\right)} \frac{\kappa_{nm}^R}{\mu} \right)^{1-\sigma} \frac{-\rho}{\sigma-\rho-1} \bar{z}_m^{j\ \sigma-\rho-1} + \left(\left(\omega_n^{j,B}\right)^{\gamma^j} \left(P_n^{j,M}\right)^{\left(1-\gamma^j\right)} \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}$$

$$= \left[\tilde{\sigma}^{1-\sigma} \frac{-\rho}{\sigma-\rho-1} \sum_{m=1}^{N} Y_m \left(\left(w_m^{j,R}\right)^{\gamma^j} \left(P_m^{j,M}\right)^{\left(1-\gamma^j\right)} \frac{\kappa_n^R}{\mu} \right)^{1-\sigma} \left[\left(w_m^{j,R}\right)^{\gamma^j} \left(P_m^{j,M}\right)^{\left(1-\gamma^j\right)} \frac{\tilde{\sigma}}{\mu} \left(\frac{\sigma}{\eta^j}\right)^{\frac{1}{\sigma-1}} \left[\frac{w_m^{j,R} f_m}{\sum_n \left(\frac{R_{nm}^R}{P_n^{R j}\big)^{1-\sigma}} Y_n \right)} \right]^{\frac{1}{\sigma-1}} \right]^{\sigma-\rho-1} \right. + \Big(\big(\omega_n^{j,B}\big)^{\gamma^j} \big(P_n^{j,$$

$$= \left[\lambda \sum_{m=1}^{N} Y_m \left(\left(w_m^{j,R}\right)^{\gamma^j} \left(P_m^{j,M}\right)^{\left(1-\gamma^j\right)} \frac{\left(\kappa_{nm}^R\right)^{\frac{\sigma-1}{\rho}}}{\mu} \right)^{-\rho} \left[\frac{w_m^{j,R} f_m}{\sum_n \left(\frac{\kappa_{nm}^R}{P_n^{R,j}} \right)^{1-\sigma} Y_n} \right]^{\frac{\sigma-\rho-1}{\sigma-1}} + \left(\left(\omega_n^{j,B}\right)^{\gamma^j} \left(P_n^{j,M}\right)^{\left(1-\gamma^j\right)} \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}$$

$$= \left[\lambda \sum_{m=1}^{N} Y_m \left(\left(w_m^{j,R}\right)^{\gamma^j} \left(P_m^{j,M}\right)^{\left(1-\gamma^j\right)} \frac{\left(k_{nm}^R\right)^{\frac{\sigma-1}{\rho}}}{\mu} \right)^{-\rho} \left[\frac{w_m^{j,R} f_m}{\theta_m^j} \right]^{\frac{\sigma-\rho-1}{\sigma-1}} + \left(\left(\omega_n^{j,B}\right)^{\gamma^j} \left(P_n^{j,M}\right)^{\left(1-\gamma^j\right)} \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}$$

The second part of the model deals with the total exports from region m to n, denoted as $X_{nm}^{j,R}$. This equation calculates the aggregate value of goods from sector j that are exported from region m to region n. The exports are determined by the productivity threshold, wage rates, prices, and sectoral income levels in both the exporting and importing regions:

$$X_{nm}^{j,R} = \int_{\bar{z}_m^j} w_m^{j,B} l_m^{j,B} X_{nm}(\phi) dG(\phi) = \int_{\bar{z}_m^j} Y_m \left(\frac{p_{nm}^j(\phi)}{P_n^j} \right)^{1-\sigma} \eta^j Y_n dG(\phi)$$

$$= \int_{\bar{z}_m^j} Y_m \left(\tilde{\sigma} \frac{\left(w_m^{j,R}\right)^{\gamma^j} \left(P_m^{j,M}\right)^{\left(1-\gamma^j\right)} \kappa_{nm}^R}{z_m^j \mu P_n^j} \right)^{1-\sigma} \eta^j Y_n dG(\phi)$$

$$= Y_m \left(\tilde{\sigma} \frac{\left(w_m^{j,R}\right)^{\gamma^j} \left(P_m^{j,M}\right)^{\left(1-\gamma^j\right)} \kappa_{nm}^R}{\mu P_n^j} \right)^{1-\sigma} \eta^j Y_n \frac{-\rho}{\sigma - \rho - 1} \bar{z}_m^{j\ \sigma-\gamma-1}$$

$$= Y_m \left(\tilde{\sigma} \frac{\left(w_m^{j,R}\right)^{\gamma^j} \left(P_m^{j,M}\right)^{\left(1-\gamma^j\right)} \kappa_{nm}^R}{\mu P_n^j} \right)^{1-\sigma} \eta^j Y_n \frac{-\rho}{\sigma - \rho - 1} \left[\left(w_m^{j,R}\right)^{\gamma^j} \left(P_m^{j,M}\right)^{\left(1-\gamma^j\right)} \frac{\tilde{\sigma}}{\mu} \left(\frac{\sigma}{\eta^j} \right)^{\frac{1}{\sigma-1}} \left[\frac{w_m^{j,R} f_m}{\sum_n \left(\frac{\kappa_{nm}^R m}{P_n^{R,j}} \right)^{1-\sigma} Y_n} \right]^{\frac{1}{\sigma-1}} \right]^{\sigma-\rho-1}$$

$$= \tilde{\sigma}^{-\rho} \left(\frac{\sigma}{\eta^j} \right)^{\frac{\sigma-\rho-1}{1-\sigma}} \frac{-\rho}{\sigma - \rho - 1} Y_m \left(\left(w_m^{j,R}\right)^{\gamma^j} \left(P_m^{j,M}\right)^{\left(1-\gamma^j\right)} \frac{\left(\kappa_{nm}^R\right)^{\frac{\sigma-1}{\rho}}}{\mu} \right)^{-\rho} \left[\frac{w_m^{j,R} f_m}{\sum_n \left(\frac{\kappa_{nm}^R m}{P_n^{R,j}} \right)^{1-\sigma} Y_n} \right]^{\frac{\sigma-\rho-1}{\sigma-1}} \eta^j Y_n \left(P_n^j\right)^{\sigma-1}$$

$$= \lambda_2 Y_m \left(\left(w_m^{j,R}\right)^{\gamma^j} \left(P_m^{j,M}\right)^{\left(1-\gamma^j\right)} \frac{\left(\kappa_{nm}^R\right)^{\frac{\sigma-1}{\rho}}}{\mu} \right)^{-\rho} \left[\frac{w_m^{j,R} f_m}{\sum_n \left(\frac{\kappa_{nm}^R m}{P_n^{R,j}} \right)^{1-\sigma} Y_n} \right]^{\frac{\sigma-\rho-1}{\sigma-1}} \eta^j Y_n \left(P_n^j\right)^{\sigma-1}$$

The total BM sales in region n, $X_{nn}^{j,B}$ can then be expressed as:

$$X_{nn}^{j,B} = \left(\frac{p_{nn}^{j,B}}{P_n^j} \right)^{1-\sigma} \eta^j Y_n = \left(\left(\omega_n^{j,B}\right)^{\gamma^j} \left(P_n^{j,M}\right)^{\left(1-\gamma^j\right)} \right)^{1-\sigma} \eta^j Y_n \left(P_n^j\right)^{\sigma-1}$$

Finally, the model considers the expenditure share of region m in region n and how it changes over time that reflects how shifts in variables like wages, prices, and productivity can impact the flow of goods and services between regions:

$$x_{nm}^{j,R} = \frac{\lambda Y_m \left(\left(w_m^{j,R} \right)^{\gamma^j} \left(P_m^{j,M} \right)^{\left(1-\gamma^j\right)} \frac{\left(\kappa_{nm}^R \right)^{\frac{\sigma-1}{\rho}}}{\mu} \right)^{-\rho} \left[\frac{w_m^{j,R} f_m}{\sum_n \left(\frac{\kappa_{nm}^R}{P_n^{R,j}} \right)^{1-\sigma} Y_n} \right]^{\frac{\sigma-\rho-1}{\sigma-1}}}{\sum_h \lambda Y_h \left(\left(w_h^{j,R} \right)^{\gamma^j} \left(P_h^{j,M} \right)^{\left(1-\gamma^j\right)} \frac{\left(\kappa_{nm}^R \right)^{\frac{\sigma-1}{\rho}}}{\mu} \right)^{-\rho} \left[\frac{w_h^{j,R} f_h}{\sum_n \left(\frac{\kappa_{nh}^R}{P_n^{R,j}} \right)^{1-\sigma} Y_n} \right]^{\frac{\sigma-\rho-1}{\sigma-1}} + \left(\left(\omega_n^{j,B} \right)^{\gamma^j} \left(P_n^{j,M} \right)^{\left(1-\gamma^j\right)} \right)^{1-\sigma}}$$

$$x_{nn}^{j,B} = \frac{\left(\left(\omega_n^{j,B} \right)^{\gamma^j} \left(P_n^{j,M} \right)^{\left(1-\gamma^j\right)} \right)^{1-\sigma}}{\sum_h \lambda Y_h \left(\left(w_h^{j,R} \right)^{\gamma^j} \left(P_h^{j,M} \right)^{\left(1-\gamma^j\right)} \frac{\left(\kappa_{nm}^R \right)^{\frac{\sigma-1}{\rho}}}{\mu} \right)^{-\rho} \left[\frac{w_h^{j,R} f_h}{\sum_n \left(\frac{\kappa_{nh}^R}{P_n^{R,j}} \right)^{1-\sigma} Y_n} \right]^{\frac{\sigma-\rho-1}{\sigma-1}} + \left(\left(\omega_n^{j,B} \right)^{\gamma^j} \left(P_n^{j,M} \right)^{\left(1-\gamma^j\right)} \right)^{1-\sigma}}$$

BIBLIOGRAPHY

Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, 67(2):251–333.

Acemoglu, D. (1999). Changes in unemployment and wage inequality: An alternative theory and some evidence. *American economic review*, 89(5):1259–1278.

Acemoglu, D. (2007). Equilibrium bias of technology. *Econometrica*, 75(5):1371–1409.

Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of labor economics*, volume 4, pages 1043–1171. Elsevier.

Acemoglu, D. and Restrepo, P. (2020). Robots and jobs: Evidence from us labor markets. *Journal of political economy*, 128(6):2188–2244.

Acemoglu, D. and Zilibotti, F. (2001). Productivity differences. *The Quarterly Journal of Economics*, 116(2):563–606.

Adao, R., Arkolakis, C., and Esposito, F. (2019). General equilibrium effects in space: Theory and measurement. Technical report, National Bureau of Economic Research.

Altonji, J. G., Bharadwaj, P., and Lange, F. (2012). Changes in the characteristics of american youth: Implications for adult outcomes. *Journal of Labor Economics*, 30(4):783–828.

Amior, M. and Manning, A. (2018). The persistence of local joblessness. *American Economic Review*, 108(7):1942–70.

Anderson, S., Engers, M., and Savelle, D. (2022). An equilibrium analysis of ordered search. *Working Paper*.

Antràs, P., Fort, T. C., Gutiérrez, A., and Tintelnot, F. (2022). Trade policy and global sourcing: An efficiency rationale for tariff escalation. Technical report, National Bureau of Economic Research.

Arkolakis, C., Costinot, A., and Rodríguez-Clare, A. (2012). New trade models, same old gains? *American Economic Review*, 102(1):94–130.

Arkolakis, C., Ramondo, N., Rodríguez-Clare, A., and Yeaple, S. (2018). Innovation and production in the global economy. *American Economic Review*, 108(8):2128–2173.

Arkolakis, C., Rodrıguez-Clare, A., and Su, J.-H. (2017). A multivariate distribution with pareto tails and pareto maxima. *Working paper*.

Armstrong, M. (2017). Ordered consumer search. *Journal of the European Economic Association*, 15(5):989–1024.

Armstrong, M. and Vickers, J. (2015). Which demand systems can be generated by discrete choice? *Journal of Economic Theory*, 158:293–307.

Atalay, E., Phongthiengtham, P., Sotelo, S., and Tannenbaum, D. (2020). The evolution of work in the united states. *American Economic Journal: Applied Economics*, 12(2):1–34.

Atkinson, A. B. and Stiglitz, J. E. (1969). A new view of technological change. *The Economic Journal*, 79(315):573–578.

Austin, B., Glaeser, E., and Summers, L. (2018). Jobs for the heartland: Place-based policies in 21st-century america. *Brookings Papers on Economic Activity*, page 151.

Autor, D. H. and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the us labor market. *American economic review*, 103(5):1553–97.

Autor, D. H. and Handel, M. J. (2013). Putting tasks to the test: Human capital, job tasks, and wages. *Journal of labor Economics*, 31(S1):S59–S96.

Autor, D. H., Katz, L. F., and Krueger, A. B. (1998a). Computing inequality: have computers changed the labor market? *The Quarterly journal of economics*, 113(4):1169–1213.

Autor, D. H., Katz, L. F., and Krueger, A. B. (1998b). Computing inequality: have computers changed the labor market? *The Quarterly journal of economics*, 113(4):1169–1213.

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

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

Autor, D. H. and Price, B. (2013). The changing task composition of the us labor market: An update of autor, levy, and murnane (2003). *Working Paper*.

Baley, I., Figueiredo, A., and Ulbricht, R. (2022). Mismatch cycles. *Journal of Political Economy*, 130(11):2943–2984.

Basu, S. and Weil, D. N. (1998). Appropriate technology and growth. *The Quarterly Journal of Economics*, 113(4):1025–1054.

Baugh, B., Ben-David, I., and Park, H. (2018). Can taxes shape an industry? evidence from the implementation of the “amazon tax”. *The Journal of Finance*, 73(4):1819–1855.

Braxton, J. C., Herkenhoff, K. F., and Phillips, G. M. (2020). Can the unemployed borrow? implications for public insurance. Technical report, National Bureau of Economic Research.

Braxton, J. C. and Taska, B. (2021). Technological change and the consequences of job loss.

Braxton, J. C. and Taska, B. (2023). Technological change and the consequences of job loss. *American Economic Review*, 113(2):279–316.

Brynjolfsson, E., Hu, Y., and Simester, D. (2011). Goodbye pareto principle, hello long tail: The effect of search costs on the concentration of product sales. *Management science*, 57(8):1373–1386.

Burstein, A., Morales, E., and Vogel, J. (2019). Changes in between-group inequality: computers, occupations, and international trade. *American Economic Journal: Macroeconomics*, 11(2):348–400.

Caliendo, L., Dvorkin, M., and Parro, F. (2019). Trade and labor market dynamics: General equilibrium analysis of the china trade shock. *Econometrica*, 87(3):741–835.

Caliendo, L. and Parro, F. (2015). Estimates of the trade and welfare effects of nafta. *The Review of Economic Studies*, 82(1):1–44.

Caliendo, L., Parro, F., Rossi-Hansberg, E., and Sarte, P.-D. (2018). The impact of regional and sectoral productivity changes on the us economy. *The Review of economic studies*, 85(4):2042–2096.

Caselli, F. and Coleman, Wilbur John, I. (2006). The world technology frontier. *American Economic Review*, 96(3):499–522.

Castex, G. and Kogan Dechter, E. (2014). The changing roles of education and ability in wage determination. *Journal of Labor Economics*, 32(4):685–710.

Chaney, T. (2008). Distorted gravity: the intensive and extensive margins of international trade. *American Economic Review*, 98(4):1707–1721.

Chetty, R. and Hendren, N. (2018). The impacts of neighborhoods on intergenerational mobility ii: County-level estimates. *The Quarterly Journal of Economics*, 133(3):1163–1228.

Chetty, R., Hendren, N., and Katz, L. F. (2016). The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment. *American Economic Review*, 106(4):855–902.

Chevalier, J. and Goolsbee, A. (2003). Measuring prices and price competition online: Amazon.com and barnesandnoble.com. *Quantitative marketing and economics*, 203.

Choi, M., Dai, A. Y., and Kim, K. (2018). Consumer search and price competition. *Econometrica*, 86(4):1257–1281.

Choné, P. and Kramarz, F. (2021). Matching workers' skills and firms' technologies: From bundling to unbundling. *Working Paper*.

Cortes, G. M., Jaimovich, N., and Siu, H. E. (2021). The growing importance of social tasks in high-paying occupations: implications for sorting. *Journal of Human Resources*, pages 0121–11455R1.

Costello, A. B. and Osborne, J. (2005). Best practices in exploratory factor analysis: Four recommendations for getting the most from your analysis. *Practical assessment, research, and evaluation*, 10(1):7.

Costinot, A., Donaldson, D., Vogel, J., and Werning, I. (2015). Comparative advantage and optimal trade policy. *The Quarterly Journal of Economics*, 130(2):659–702.

Costinot, A. and Rodríguez-Clare, A. (2014). Trade theory with numbers: Quantifying the consequences of globalization. In *Handbook of international economics*, volume 4, pages 197–261. Elsevier.

Costinot, A. and Vogel, J. (2010). Matching and inequality in the world economy. *Journal of Political Economy*, 118(4):747–786.

Davis, S. J., Faberman, R. J., and Haltiwanger, J. C. (2013). The establishment-level behavior of vacancies and hiring. *The Quarterly Journal of Economics*, 128(2):581–622.

Dekle, R., Eaton, J., and Kortum, S. (2008). Global rebalancing with gravity: Measuring the burden of adjustment. *IMF Staff Papers*, 55(3):511–540.

Deming, D. and Kahn, L. B. (2018). Skill requirements across firms and labor markets: Evidence from job postings for professionals. *Journal of Labor Economics*, 36(S1):S337–S369.

Deming, D. J. (2017). The growing importance of social skills in the labor market. *The Quarterly Journal of Economics*, 132(4):1593–1640.

Deming, D. J. and Noray, K. (2020). Earnings dynamics, changing job skills, and stem careers. *The Quarterly Journal of Economics*, 135(4):1965–2005.

Dinerstein, M., Einav, L., Levin, J., and Sundaresan, N. (2018). Consumer price search and platform design in internet commerce. *American Economic Review*, 108(7):1820–59.

Dodini, S., Lovenheim, M., and Willen, A. (2022). The Changing Skill Content of Private Sector Union Coverag.

Dolfen, P., Einav, L., Klenow, P. J., Klopack, B., Levin, J. D., Levin, L., and Best, W. (2019). Assessing the gains from e-commerce. Technical report, National Bureau of Economic Research.

Duflo, E. and Pande, R. (2007). Dams. *The Quarterly Journal of Economics*, 122(2):601–646.

Dvorkin, M. (2014). Sectoral shocks, reallocation and unemployment in competitive labor markets. Technical report, Yale University.

Eaton, J. and Kortum, S. (2002). Technology, geography, and trade. *Econometrica*, 70(5):1741–1779.

Edmond, C. and Mongey, S. (2021). Unbundling labor. *Working Paper*.

Eeckhout, J. and Kircher, P. (2018). Assortative matching with large firms. *Econometrica*, 86(1):85–132.

Ellison, G. and Ellison, S. F. (2018). Match quality, search, and the internet market for used books. Technical report, National Bureau of Economic Research.

Faber, B. (2014). Trade integration, market size, and industrialization: evidence from china's national trunk highway system. *Review of Economic Studies*, 81(3):1046–1070.

Fan, J., Tang, L., Zhu, W., and Zou, B. (2018). The alibaba effect: Spatial consumption inequality and the welfare gains from e-commerce. *Journal of International Economics*, 114:203–220.

Freeman, R. B., Ganguli, I., and Handel, M. J. (2020). Within-occupation changes dominate changes in what workers do: A shift-share decomposition, 2005–2015. In *AEA Papers and Proceedings*, volume 110, pages 394–99.

Galle, S., Rodríguez-Clare, A., and Yi, M. (2022). Slicing the Pie: Quantifying the Aggregate and Distributional Effects of Trade. *The Review of Economic Studies*.

Gervais, A. and Jensen, J. B. (2019). The tradability of services: Geographic concentration and trade costs. *Journal of International Economics*, 118:331–350.

Gibbons, R., Katz, L. F., Lemieux, T., and Parent, D. (2005). Comparative advantage, learning, and sectoral wage determination. *Journal of labor economics*, 23(4):681–724.

Goldin, C. and Katz, L. F. (2010). *The race between education and technology*. harvard university press.

Goldmanis, M., Hortaçsu, A., Syverson, C., and Emre, Ö. (2010). E-commerce and the market structure of retail industries. *The Economic Journal*, 120(545):651–682.

Goos, M., Manning, A., and Salomons, A. (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American economic review*, 104(8):2509–2526.

Greenwood, J., Hercowitz, Z., and Krusell, P. (1997). Long-run implications of investment-specific technological change. *The American economic review*, pages 342–362.

Guvenen, F., Kuruscu, B., Tanaka, S., and Wiczer, D. (2020). Multidimensional skill mismatch. *American Economic Journal: Macroeconomics*, 12(1):210–44.

Hagedorn, M., Law, T. H., and Manovskii, I. (2017). Identifying equilibrium models of labor market sorting. *Econometrica*, 85(1):29–65.

Heckman, J. J., Lochner, L., and Taber, C. (1998). Explaining rising wage inequality: Explorations with a dynamic general equilibrium model of labor earnings with heterogeneous agents. *Review of economic dynamics*, 1(1):1–58.

Heckman, J. J. and Sedlacek, G. (1985). Heterogeneity, aggregation, and market wage functions: an empirical model of self-selection in the labor market. *Journal of political Economy*, 93(6):1077–1125.

Hershbein, B. and Kahn, L. B. (2018). Do recessions accelerate routine-biased technological change? evidence from vacancy postings. *American Economic Review*, 108(7):1737–1772.

Hopenhayn, H. A. (1992). Entry, exit, and firm dynamics in long run equilibrium. *Econometrica: Journal of the Econometric Society*, pages 1127–1150.

Houde, J.-F., Newberry, P., and Seim, K. (2021). Economies of density in e-commerce: A study of amazon's fulfillment center network. Technical report, National Bureau of Economic Research.

Hsieh, C.-T., Hurst, E., Jones, C. I., and Klenow, P. J. (2019). The allocation of talent and us economic growth. *Econometrica*, 87(5):1439–1474.

Jin, X. (2017). The returns to specialization: Evidence from education-occupation match in the us from 1993 to 2017. *Working Paper*.

Johnson, G. E. (1997). Changes in earnings inequality: the role of demand shifts. *Journal of economic perspectives*, 11(2):41–54.

Jones, C. I. (2005). The shape of production functions and the direction of technical change. *The Quarterly Journal of Economics*, 120(2):517–549.

Kaas, L. and Kircher, P. (2015). Efficient firm dynamics in a frictional labor market. *American Economic Review*, 105(10):3030–3060.

Katz, L. F. and Murphy, K. M. (1992). Changes in relative wages, 1963–1987: supply and demand factors. *The quarterly journal of economics*, 107(1):35–78.

Kline, P. and Moretti, E. (2013). Place based policies with unemployment. *American Economic Review*, 103(3):238–43.

Krugman, P. (1991). Increasing returns and economic geography. *Journal of political economy*, 99(3):483–499.

Krugman, P. and Venables, A. J. (1995). Globalization and the inequality of nations. *The quarterly journal of economics*, 110(4):857–880.

Krusell, P., Ohanian, L. E., Ríos-Rull, J.-V., and Violante, G. L. (2000). Capital-skill complementarity and inequality: A macroeconomic analysis. *Econometrica*, 68(5):1029–1053.

Lagakos, D. and Waugh, M. E. (2013). Selection, agriculture, and cross-country productivity differences. *American Economic Review*, 103(2):948–80.

Lazear, E. P. (2009). Firm-specific human capital: A skill-weights approach. *Journal of political economy*, 117(5):914–940.

Lee, E. (2020). Trade, inequality, and the endogenous sorting of heterogeneous workers. *Journal of International Economics*, 125:103310.

León-Ledesma, M. A. and Satchi, M. (2019). Appropriate technology and balanced growth. *The Review of Economic Studies*, 86(2):807–835.

Lindenlaub, I. (2017). Sorting multidimensional types: Theory and application. *The Review of Economic Studies*, 84(2):718–789.

Lipscomb, M., Mobarak, A. M., and Barham, T. (2013). Development effects of electrification: Evidence from the topographic placement of hydropower plants in brazil. *American Economic Journal: Applied Economics*, 5(2):200–231.

Lise, J. and Postel-Vinay, F. (2020). Multidimensional skills, sorting, and human capital accumulation. *American Economic Review*, 110(8):2328–76.

Menzio, G. and Shi, S. (2010). Block recursive equilibria for stochastic models of search on the job. *Journal of Economic Theory*, 145(4):1453–1494.

Menzio, G. and Shi, S. (2011). Efficient search on the job and the business cycle. *Journal of Political Economy*, 119(3):468–510.

Mercan, Y. and Schoefer, B. (2020). Jobs and matches: Quits, replacement hiring, and vacancy chains and vacancy chains. *American Economic Review: Insights*, 2(1):101–124.

Murphy, K. (1986). Specilization and human capital. *PhD dissertation, University of Chicago - Department of Economics*.

Neal, D. (1999). The complexity of job mobility among young men. *Journal of labor Economics*, 17(2):237–261.

Ocampo, S. (2022). A task-based theory of occupations with multidimensional heterogeneity. *Working Paper*.

Osberg, L. (1993). Fishing in different pools: job-search strategies and job-finding success in canada in the early 1980s. *Journal of labor economics*, 11(2):348–386.

Pozzi, A. (2013). The effect of internet distribution on brick-and-mortar sales. *The RAND Journal of Economics*, 44(3):569–583.

Puga, D. (1999). The rise and fall of regional inequalities. *European economic review*, 43(2):303–334.

Ransom, M. R. and Phipps, A. (2017). The changing occupational distribution by college major. In *Skill Mismatch in Labor Markets*. Emerald Publishing Limited.

Rosen, S. (1983). Specialization and human capital. *Journal of Labor Economics*, 1(1):43–49.

Ross, M. B. (2017). Routine-biased technical change: Panel evidence of task orientation and wage effects. *Labour Economics*, 48:198–214.

Roy, A. D. (1951). Some thoughts on the distribution of earnings. *Oxford economic papers*, 3(2):135–146.

Şahin, A., Song, J., Topa, G., and Violante, G. L. (2014). Mismatch unemployment. *American Economic Review*, 104(11):3529–64.

Schaal, E. (2017). Uncertainty and unemployment. *Econometrica*, 85(6):1675–1721.

Shi, S. (2001). Pricing and signaling with frictions. *Journal of Economic Theory*, 98:232–260.

Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American economic review*, 95(1):25–49.

Stanchi, F. (2019). Creative destruction? the effects of e-commerce on demand and productivity for department stores. *Working Paper*.

Thompson, B. and Daniel, L. G. (1996). Factor analytic evidence for the construct validity of scores: A historical overview and some guidelines.

Tinbergen, J. (1974). Substitution of graduate by other labour. *Kyklos: international review for social sciences*.

Tinbergen, J. (1975). *Income differences: recent research*.

Tippins, N. T. and Hilton, M. L. (2010). *A database for a changing economy: Review of the Occupational Information Network (O* NET)*. National Academies Press.

Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, pages 641–654.

Xia, P., Zhang, L., and Li, F. (2015). Learning similarity with cosine similarity ensemble. *Information Sciences*, 307:39–52.

Yamaguchi, S. (2012). Tasks and heterogeneous human capital. *Journal of Labor Economics*, 30(1):1–53.