THE
QUARTERLY JOURNAL
OF ECONOMICS

Vol. 139 2024 Issue 3

NEW FRONTIERS: THE ORIGINS AND CONTENT OF NEW WORK, 1940–2018*

David Autor
Caroline Chin
Anna Salomons
Bryan Seegmiller

We answer three core questions about the hypothesized role of newly emerging job categories ("new work") in counterbalancing the erosive effect of task-displacing automation on labor demand: what is the substantive content of new work, where does it come from, and what effect does it have on labor demand? We construct a novel database spanning eight decades of new job titles linked to U.S. Census microdata and to patent-based measures of occupations' exposure

* We thank Daron Acemoglu, David Deming, Brad DeLong, Matt Gentzkow, Maarten Goos, Gordon Hanson, Lawrence Katz, Lynda Laughlin, Magnus Lodefalk, Peter Lambert, Marin Soljačić, Sebastian Steffen, John Van Reenen, three anonymous referees, the expert staff of the U.S. Census Bureau, and innumerable seminar and conference participants for insights and critiques that have vastly improved the article. We thank Benjamin Boehlert, Fiona Chen, Grace Chuan, Raimundo Contreras González, Rebecca Jackson, Zhe Fredric Kong, Jonathan Rojas, Edwin Song, Ishaana Talesara, Liang Sunny Tan, Jose Velarde, Yuting Brenda Wu, Rocky Xie, and Whitney Zhang for expert research assistance. We thank Dimitris Papanikolaou for sharing a database of raw patent texts; Enrico Berkes for providing a data set of historical patent citations for earlier patents; and Sammy Mark, Fredrick Soo, and Mary Wakarindy from Digital Divide Data for hand-keying historical Census Alphabetical Index records. We thank Greg Bronevetsky and Guy Ben-Ishai of Google for facilitating industrial-level usage of the Google Ngram Viewer to analyze the emergence of new job titles. We gratefully acknowledge support from the Carnegie Corporation, Google, Instituut Gak, the MIT Work of the Future Task Force, Schmidt Futures, the Smith Richardson Foundation, and the Washington Center for Equitable Growth.

The Quarterly Journal of Economics (2024), 1399–1465. https://doi.org/10.1093/qje/qjae008.
Advance Access publication on March 15, 2024.

to labor-augmenting and labor-automating innovations. The majority of current employment is in new job specialties introduced since 1940, but the locus of new-work creation has shifted from middle-paid production and clerical occupations over 1940–1980 to high-paid professional occupations and secondarily to low-paid services since 1980. New work emerges in response to technological innovations that complement the outputs of occupations and demand shocks that raise occupational demand. Innovations that automate tasks or reduce occupational demand slow new-work emergence. Although the flow of augmentation and automation innovations is positively correlated across occupations, the former boosts occupational labor demand while the latter depresses it. The demand-eroding effects of automation innovations have intensified in the past four decades while the demand-increasing effects of augmentation innovations have not. *JEL codes:* E24, J11, J23, J24.

I. Introduction

A vast economic literature analyzes how rapidly evolving digital technologies—information and communications technologies, robotics, artificial intelligence—affect employment, skill demands, and earnings levels. Focusing on the substitution of machines for workers in tasks where automation has rising comparative advantage, this research anticipates and interprets the decline of middle-skill employment in high-income countries (a.k.a. job polarization) and attendant effects on wage structure, documents the concentrated impact of industrial robotics on labor demand in heavy manufacturing industries and in manufacturing-intensive communities, and explores how artificial intelligence may change the structure of occupations.

This work is comparatively silent—with key exceptions, discussed below—on the flip side of the ledger: the augmentation of human labor and the generation of new work activities that demand this labor. Economic literature on the impact of technological change on employment has primarily treated the set of human job tasks as finite and static, implying that as automation proceeds, labor is shunted into an ever-narrowing scope of activities, as in Susskind (2020). Casual observation and historical evidence suggest the opposite: as employment in labor-intensive sectors such as agriculture, textiles, and mining has eroded, the scope and variety of labor-demanding activities has expanded, for example, in medicine, software, electronics, health care, finance, entertainment, recreation, personal care, and other domains. Recent work by Acemoglu and Restrepo advances the theoretical frontier on this topic (Acemoglu and Restrepo 2018, 2019a), but a major empirical challenge remains: there is almost no direct, consistent measurement of either the emergence of new work tasks within occupations and industries or of the technological and economic forces that are hypothesized to give rise to them.

This article systematically studies the nature, sources, and consequences of the emergence of new work in the United States between 1940 and 2018. We seek to consistently measure the substantive content of new work over eight decades, explore the forces that explain when and where new work emerges, and assess whether, as hypothesized in recent literature, new work exerts a countervailing force to the employment-eroding effects of task-displacing automation. Our empirical analysis is motivated by a two-sector general equilibrium task model that, building on Acemoglu and Autor (2011) and Acemoglu and Restrepo (2018, 2019b), draws economic linkages between task automation, new-task creation, innovation incentives, and the locus and attendant skill demands of new work. Because our main contributions are empirical, the formal model is developed in Online Appendix Section J, with intuition provided in the body of the article.

Our analysis is grounded in three constructs that we operationalize using purpose-built data. The first is new work itself, by which we mean the introduction of new job tasks or job categories requiring specialized human expertise. Following path-breaking work by Lin (2011), we construct a database of new job tasks introduced over eight decades, sourced from nearly a century of internal reference volumes used by U.S. Census Bureau employees to classify the free-text job descriptions of census respondents into occupation and industry categories in each decade. This *Census Alphabetical Index of Occupations and Industries* (CAI) is updated during the processing of each decade's census to reflect new write-in titles ("micro-titles") detected by census coders.[1] Following Lin (2011), we track the emergence of new micro-titles in each decade by comparing successive editions of the CAI. Consider, for example, the micro-titles of "Technician, fingernail" (added in 2000) and "Solar photovoltaic electrician" (added in 2018), which highlight two salient attributes of new work. First,

1. While census tabulations and public-use data sources report several hundred distinct occupation and industry codes in each census year (which we call macro-titles), these titles reflect concatenations of approximately 30,000 occupational and 20,000 industry-level micro-titles enumerated in the CAI in each decade between 1930 and 2018 (U.S. Census Bureau).

new work typically requires expertise acquired through study, apprenticeship, or practical experience (e.g., in cosmetology, in the electrical trades), though the extent of expertise clearly varies across categories. Second, new work often reflects the development of novel expertise in existing work activities (e.g., electrical trade skills specific to solar installations) or through an increase in the market scale of a niche activity (e.g., nail care) rather than a fundamentally new human endeavor.[2]

The second and third constructs that we operationalize are the flow of augmentation and automation innovations over eight decades. In our terminology, augmentation innovations are technologies that increase the capabilities, quality, variety, or utility of the outputs of occupations, potentially generating new demands for worker expertise and specialization. Conversely, automation innovations are technologies that substitute for the labor inputs of occupations, potentially replacing workers performing these tasks. We construct augmentation and automation measures using natural language processing tools (NLP) to map the full text of all U.S. utility patents issued between 1920 and 2018 to the domain of occupations. Following methods introduced by Kogan et al. (2021), we represent both patent documents and occupational descriptions as weighted averages of word embeddings, which are geometric representations of word meanings. We then characterize the semantic closeness between patent texts and occupational descriptions by measuring their proximity in embedding space. Each patent may be classified as an automation innovation, an augmentation innovation, both, or neither.

To capture augmentation innovations, we harness the full set of micro-titles in the CAI associated with each macro-occupation in each decade to provide a text corpus describing the occupation's outputs. For example, in 1999, the U.S. Patent and Trademark Office (USPTO) granted patent US5924427A for a "Method of strengthening and repairing fingernails." Our procedure links this patent to the census macro-occupation of "Miscellaneous personal appearance workers," which encompasses the micro-title of "Technician, fingernail." Similarly, our algorithm links the 2014 patent US7605498B2 "Systems for highly efficient solar power

2. There are exceptions: in the 1900 census, Orville and Wilbur Wright each reported their occupation as "Merchant, bicycle"; in the subsequent census, they both reported "Inventor, aeroplane" (United States Census 1900, 1910). Four decades later, the title of Airplane Designer was prevalent enough to gain its own entry in the 1950 CAI.

conversion" to the macro-occupation "Electrical and electronics engineers" which includes the micro-title "Solar photovoltaic electrician."

To capture automation innovations, we follow Webb (2020), Dechezleprêtre et al. (2021), Kogan et al. (2021), and Mann and Püttmann (2023) in identifying similarities between the content of patents and the tasks that workers perform in specific occupations—that is, their work inputs. We measure these task inputs using the 1939 *Dictionary of Occupational Titles* (DOT) for the 1940–1980 period, and the 1977 DOT for the 1980–2018 period (U.S. Department of Labor, Employment and Training Administration 1939, 1977), again calculating semantic similarities with the full corpus of utility patents issued between 1920 and 2018. For example, in 1979, the USPTO granted patent US4141082A for a "Wash-and-wear coat." Our algorithm links this patent to the macro-occupation of "Laundry and dry-cleaning workers." Similarly, our algorithm links the 1976 patent US3938435A, "Automatic mail processing apparatus," to the macro-occupation of "Mail and paper handlers." It bears emphasis that we apply fully parallel procedures for classifying augmentation and automation innovations, with the sole difference being the corpora of text used to characterize occupational outputs versus inputs. Accordingly, the distinct classifications we generate for augmentation and automation innovations stem entirely from differences in semantic content between the corpora of task inputs and occupational outputs, rather than procedural artifacts.

The empirical analysis proceeds in three steps. The first step is to characterize the substantive content of new work. We quantify its (approximate) contribution to overall employment growth between 1940 and 2018, and document how its occupational composition and educational requirements have evolved across the two four-decade intervals of the sample. We estimate that the majority of current employment is found in new job specialties (new work) introduced after 1940. But the locus of new-work creation has substantially shifted over the eight decades of our sample. Between 1940 and 1980, a substantial share of new work accrued to middle-paid production and clerical occupations. In the more recent four decades, new work has primarily emerged in high-paid professional occupations and secondarily in low-paid service occupations.

The second step analyzes where new work comes from. We posit that new job tasks (i.e., new work) derive from two primary sources. A first source is augmentation innovations, meaning the introduction of new processes and products (e.g., solar voltaic cells), new services (e.g., fingernail hardening), and entirely new products or industries (e.g., dry-cleaning, commercial air travel), that create novel demands for expertise and specific competencies that correspond to new work tasks and are reified in new job titles in our empirical analysis. Conversely, if automation innovations primarily enable machines to subsume existing tasks, as per Acemoglu and Restrepo (2018, 2019b), these innovations should not spur the introduction of new labor-using tasks. (While it is quite plausible that automation innovations give rise to new machine tasks, we are not able to measure such tasks.)

Consistent with this reasoning, we find that augmentation and automation innovations have distinct, asymmetric relationships to the creation of new work. Augmentation innovations strongly predict the locus of new-task creation, as measured by the emergence of new job titles, across occupations and over time. By contrast, automation innovations do not predict where new work emerges. As we show below, augmentation and automation innovation flows are strongly positively correlated at the occupational level, so this is a stringent test.

The second driver of new work that we consider is fluctuations in market size that increase or depress the value of occupational outputs. We document that the flow of new work responds elastically to shifts in demand for occupational output. Adverse demand shocks, which we identify using the widely studied China trade shock (Autor, Dorn, and Hanson 2016), slow the emergence of new work tasks in exposed occupations, even conditional on exposure to augmentation innovations and contemporaneous changes in employment. Conversely, positive occupational demand shocks, which we identify using shifts in demographic structure following DellaVigna and Pollet (2007), accelerate the emergence of new work in exposed occupations. These demographic shifts help account for the emergence of new job types in lower-paid personal services that have been relatively immune to labor-augmenting innovations but are nevertheless a key locus of new-work creation for noncollege workers over the last four decades.

Our final set of results turns the focus from new work—meaning new occupational titles—to shifts in overall occupational labor demand. While it is easy to verify that new tasks differentially emerge in growing occupations—or equivalently that employment differentially expands in occupations in which new tasks are emerging—this correlation does not directly test the role of innovation in spurring or eroding employment. To provide that test, we regress occupational employment and wage bill changes over two four-decade periods on measures of the flow of augmentation and automation innovations to which these occupations are exposed. We document that employment and wage bills grow in occupations exposed to augmentation innovations and contract in occupations exposed to automation innovations. The pattern of results further suggests that the demand-eroding effects of automation innovations have intensified in the past four decades while the demand-increasing effects of augmentation innovations have not—though we stress that our innovation-exposure measures lack sufficient cardinality to draw definitive conclusions.

Because both innovation and employment are plausibly jointly determined at the occupational level, it is critical to ask whether these findings capture causal relationships. We test for causality by developing an instrumental variables strategy that identifies shifts in the flow of augmentation and automation innovations stemming from breakthrough innovations occurring two decades earlier. Using these breakthrough innovations as instruments for subsequent downstream patent flows, we establish that augmentation innovations cause the emergence of new work and the growth of occupational employment, while conversely, automation innovations do not spur new work but do erode occupational employment.

Our work contributes to three branches of literature on technology, skills, and employment. A first studies the interplay between supply, demand, technologies, and institutions in shaping the long-run evolution of skill demands, occupational structure, and wage inequality.[3] A foundational assumption of this literature is that technological change has non-neutral effects on the skill composition of labor demand. We contribute to this literature by linking changes in the structure of occupational demands to the shifting locus of innovation over eight decades, showing that: new work is a quantitatively large contributor to aggregate employment change, that it emerges where innovative activity is focused, and that the locus of new-work generation

3. This vast literature includes Goldin and Margo (1992), Katz and Murphy (1992), DiNardo, Fortin, and Lemieux (1996), Acemoglu (1998), Autor, Katz, and Krueger (1998), Katz and Autor (1999), Krusell et al. (2000), Card and Lemieux (2001), Goldin and Katz (2008), Autor, Goldin, and Katz (2020), Haanwinckel (2023), and Vogel (2023).

has shifted across recent decades, leading to overall changes in occupational structure.

We secondarily contribute to a contemporary literature that explores how automation technologies substitute for existing work, as measured by occupational structure or job tasks.[4] Our work is closely related to Webb (2020), Dechezleprêtre et al. (2021), Mann and Püttmann (2023), and most directly Kogan et al. (2021), who use NLP tools to identify innovations recorded in patents that potentially substitute for the tasks performed by workers, as well as papers by Brynjolfsson and Mitchell (2017); Brynjolfsson, Mitchell, and Rock (2018); Eloundou et al. (2023); Felten, Raj, and Seamans (2018, 2019, 2023) that predict which occupational tasks can be performed by artificial intelligence. Distinct from this literature, we develop a method to identify innovations that generate new work tasks by augmenting occupational outputs.

Most directly, this article contributes to research on the micro- and macroeconomic origins of new work, including Goldin and Katz (1998), Lin (2011), Acemoglu and Restrepo (2018, 2019a), Atack, Margo, and Rhode (2019), Frey (2019), Atalay et al. (2020), Atalay and Sarada (2020), Deming and Noray (2020), and Kim (2022). We extend the pioneering work in Lin (2011), while expanding its scope to provide direct, representative, and time-consistent measurement of new-task creation across eight decades.[5] Distinct from any prior work of which we are aware, we identify innovations that complement occupational outputs, which we hypothesize (and empirically confirm) spur new-task creation. This technique enables us to document that automation and augmentation have measurable and distinct effects on both new-work creation and occupational employment.

4. Relevant works include Autor, Levy, and Murnane (2003), Autor, Katz, and Kearney (2006), Goos and Manning (2007), Goos, Manning, and Salomons (2009, 2014), Acemoglu and Autor (2011), Autor and Dorn (2013), Michaels, Natraj, and Van Reenen (2014), Akerman, Gaarder, and Mogstad (2015), Bárány and Siegel (2018), Cortes et al. (2020), Acemoglu and Restrepo (2022), and Böhm, von Gaudecker, and Schran (2022).

5. In related work, Acemoglu and Restrepo (2019a) develop a set of ingenious proxies for the appearance of new work based on changes in labor share and in the mix of three-digit occupations (what we call macro-occupations) within industries. Atalay et al. (2020), Atalay and Sarada (2020), and Deming and Noray (2020), measure the appearance of new work by analyzing the text of job advertisements. Kim (2022) finds that greater import competition spurs a relative increase in employment in managerial occupations (which are intensive in new work).

The article proceeds as follows. Section II details our methods for identifying new work, describes how the locus of new work has evolved over 1940–2018, and outlines how we identify augmentation and automation innovations embodied in patents that link to specific occupations. Section III briefly sketches our theoretical model, which is formalized in the Online Appendix. Section IV uses both OLS and instrumental variables methods to test the relationship between innovation, demand shifts, and new-work creation. The final empirical section, Section V, assesses whether augmentation and automation innovations have measurable, countervailing effects on occupational employment and wage bill growth. Section VI concludes.

II. Data and Measurement

Our analysis links data on the emergence of new titles, the flow of augmentation and automation innovations, and the evolution of employment and earnings in industry-occupation cells over eight decades, 1940–2018. To characterize employment and earnings by occupation and industry, we use IPUMS census samples for 1940 through 2000 and census ACS samples for 2014–2018 (Ruggles et al. 2022). Our databases of new titles, augmentation flows, and automation flows are purpose-built for this study and we document them here.

II.A. Measuring New Work

We leverage Census Bureau historical coding volumes for occupations and industries for 1930 through 2018 (industry data starts in 1940) that are released each decade as the CAI. Each decade's index contains approximately 35,000 occupation and 15,000 industry micro-titles, each classified to a more aggregated (macro) census occupation or industry code. These indices serve as reference documents for census coders who classify individual census write-ins for job title and industry of employment, which are always reported as free text fields. This process has been performed consistently for occupations since 1900 and is illustrated for occupations in the American Community Survey (ACS) in Online Appendix Figure A.VIII. Unfortunately, the Census Bureau does not systematically remove potentially extinct titles from the CAI (since they may occasionally reoccur), meaning that the CAI is not suitable for measuring work obsolescence.

When census coders encounter multiple instances of a write-in occupational title that cannot be ascribed to an existing micro-title, they bring it to the attention of Census Bureau managers, who perform an internal review to determine whether the title is sufficiently distinct and prevalent to be added to the index. For example, the micro-occupation-title "Artificial intelligence specialist" was added to the CAI in 2000 and classified to the broader occupational title "Computer scientists and systems analysts," which appears in published census tabulations and public-use data sets. The two-stage process for adding a new title—detection by coders, review by managers—clarifies why, for example, "Mental-health counselor" was added in 1970, "Artificial intelligence specialist" in 2000, and "Sommelier" in 2010. Although workers performed these jobs in earlier decades, these specializations were too rare to warrant inclusion beyond a generic counselor, computer scientist, or restaurant server title. To provide concrete illustration, Online Appendix D2 enumerates examples of new titles introduced between 1940 and 2018 in two office occupations (Office machine operators; Stenographers, typists, and secretaries), one professional occupation (Dentists), and one blue-collar occupation (Automobile mechanics). New title additions typically reflect emerging expertise demands for new technologies, changes in educational requirements, and in the case of dentists, the addition of new clientele (Pediatric orthodontist) and entirely new types of services (e.g., Maxillofacial surgeon).

The Census Bureau does not highlight or separately list titles that are newly added to the CAI, and it frequently renames outdated job titles, adds new phrasings of existing titles, and removes gendered forms (e.g., Key boy). Hence, consistently tracking the flow of new CAI titles requires a mixture of automated and manual steps. We compare title lists across successive CAI decades using fuzzy matching and extensive manual revision to discard false positives stemming from rewording, reformatting of the index, and title additions that do not reflect discernible modifications to preexisting counterparts. We do not, for example, count "Software applications developer" as a new micro-occupation, since "Software developer" was already present when this title appeared. Our overarching aim is to retain new titles that reflect plausibly new work activities, meaning that they add a particular task specialization, work method or tool, or professional or educational requirement (Online Appendix D1 provides details.)

TABLE I

EXAMPLES OF NEW TITLES ADDED TO THE CENSUS ALPHABETICAL INDEX OF OCCUPATIONS BY VOLUME, 1940–2018

Year	Example titles added	
1940	Automatic welding machine operator	Acrobatic dancer
1950	Airplane designer	Tattooer
1960	Textile chemist	Pageants director
1970	Engineer computer application	Mental-health counselor
1980	Controller, remotely piloted vehicle	Hypnotherapist
1990	Circuit layout designer	Conference planner
2000	Artificial intelligence specialist	Amusement park worker
2010	Technician, wind turbine	Sommelier
2018	Cybersecurity analyst	Drama therapist

Notes. The table reports examples of new titles added to the census Alphabetical Index of Occupations volume of year Y corresponding to titles recognized by census coders between the start of the prior decade and the year preceding the volume's release, for example, the 1950 CAI volume includes new titles incorporated between 1940 and 1949.

Table I provides examples of the diversity of micro-titles added to the CAI in each decade from 1940 through 2018 (where, for example, new titles in the 1950 CAI reflect those detected by census staff between 1940 and 1950). The left-hand column reports titles that are plausibly associated with new or evolving technologies: "Airplane designers" in 1950, "Engineers of computer applications" in 1970, "Circuit layout designers" in 1990, and "Technicians of wind turbines" in 2010. Many new titles do not have immediate technological origins, however. As shown in the right-hand column of Table I, these titles appear to reflect changing tastes, income levels, and demographics, including "Tattooers" in 1950, "Hypnotherapists" in 1980, "Conference planners" in 1990, and "Drama therapists" in 2018. These examples motivate looking beyond exclusively technological forces, as we do below, in analyzing the sources of new-work creation.

How representative of new work is the CAI-based measure? Without a comparable source against which to benchmark it, we provide a novel test of external validity by confirming that the CAI captures new occupational titles as they enter common usage in the English language. Using the Google Ngram viewer (Michel et al. 2011), we calculate the frequency in digitized English-language books of both new and preexisting occupational titles in each year between 1940 to 2018.[6] The eight panels of

6. We thank Peter Lambert of LSE for suggesting this analysis.

Online Appendix Figure A.1 report the usage frequency of all titles added to the CAI in each decade relative to those present at the start of the decade. Frequencies are normalized as the ratio of instances of the new title (of length n) in a year (excluding unmatched titles) relative to the total number of n-grams in that year. In each cohort of new titles, the median new title is added to the CAI during approximately the decade before or during which it attains peak usage in the English language, suggesting that the CAI provides a somewhat conservative enumeration of new work.[7] Online Appendix Figure A.V further corroborates that this pattern holds across four broad occupational categories that encompass the totality of employment: blue-collar occupations, professional and information occupations, personal service occupations, and business service occupations. (Online Appendix C details these occupational groups.) These findings substantiate that our database, built from successive CAI editions, provides a representative, time-consistent measure of the emergence of new job titles in the United States over the eight decades of the sample.

II.B. The Evolution of New Work, 1940–2018

What do these data tell us about the evolution of new work? A first finding is that new work is quantitatively important. Our estimates imply that the majority of contemporary work as of 2018—roughly 60%—is found in new job titles added since 1940. Figure I reports the contribution of new work to overall employment between 1940 and 2018 in 12 exhaustive, mutually exclusive broad occupational categories. These are ordered from lowest to highest paying, with farming and mining occupations on the left side of the scale and managerial workers on the right side. We further distinguish between 2018 employment found in occupational titles that existed in 1940 versus 2018 employment in occupational titles that were added thereafter. While the majority of contemporary work is new since 1940, the share of employment in new work differs substantially by occupational category. Among professionals—the occupation that added the most workers during these eight decades—this share is 74%. For the production occupation—which added the second smallest number of workers since 1940 (after farming)—this share is 46%.

7. Titles that become newly prominent in one era do not remain at peak prominence in perpetuity, however, typically declining by 20% to 50% in the three to five decades after entering the CAI.

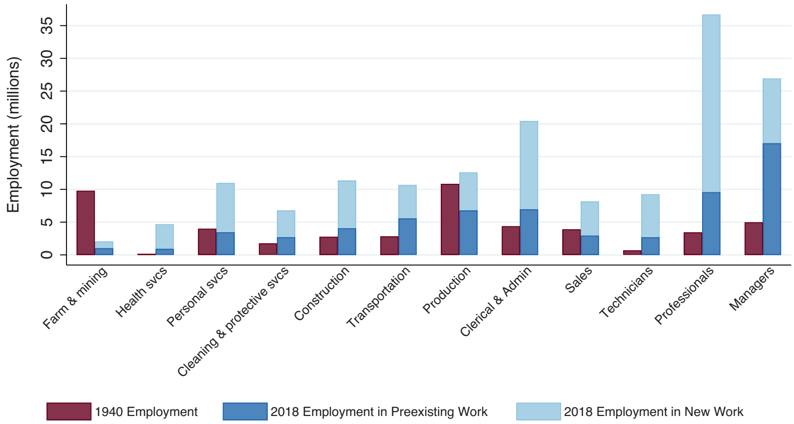

FIGURE I

Employment Counts by Broad Occupation in 1940 and 2018, Distinguishing between Titles Present in 1940 versus Those Added Subsequently

This figure reports the distribution of employment in 1940 and 2018 across broad occupational categories ordered from lowest to highest paying. The first set of bars shows 1940 employment. The second set of bars shows 2018 employment, divided into estimated 2018 employment in occupational titles that existed in 1940 (bottom), and estimated 2018 employment in occupation titles that were added since 1940 (top). Employment in 2018 is estimated by constructing a cumulative new-title share in each broad occupation—summing the number of new titles added over 1940–2018, and dividing this by the total number of titles in the 2018 index, adjusted for (the small number of) titles that were removed. See Online Appendix D3 for additional details.

Because the Census Bureau does not record or report the count of respondents within micro-titles, the estimates in Figure I should be understood as approximate. We follow Lin (2011) in imputing employment counts in new titles by multiplying the new-title share in each census macro-occupation—equal to the ratio of new micro-titles to total micro-titles within a census occupation—by total employment in the occupation.[8]

8. Lin (2011) validates this approach using a special version of the April 1971 Current Population Survey where a subsample of workers are assigned DOT titles. We further explore this imputation in Online Appendix H2 using census Complete Count data for 1940, which contains both macro-titles and the free text write-in micro-titles supplied by census respondents. The count of workers in new titles is strongly increasing in the new-title share—though the slope is below one—and this relationship is more precise when using ordinal share ranks rather than cardinal shares.

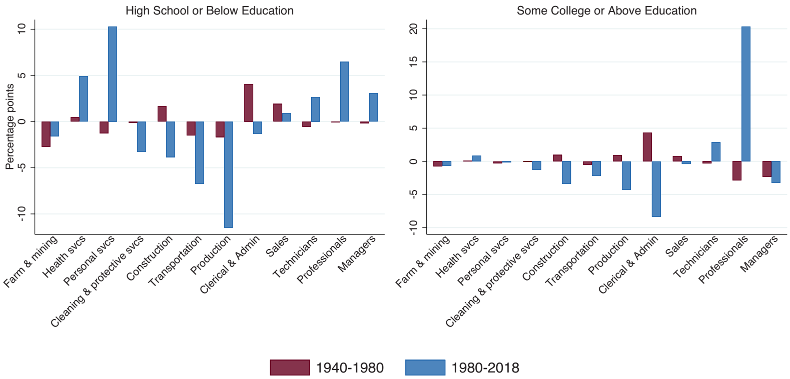

FIGURE II

Difference between the Occupational Distribution of the Flow of New Work versus the Stock of Preexisting Work by Education Group, 1940–1980 and 1980–2018

This figure reports the difference in the share of employment in new work versus the share of employment in existing work separately by education group and period. Each panel reports this difference across 12 broad occupational categories ordered from lowest to highest paying. The first set of bars represent the average difference in employment shares over 1940–1980, while the second set of bars represent the difference over 1980–2018.

A second result is that the occupational distribution of new work has changed markedly over the past eight decades, and in a manner that presages the changing overall shape of employment growth and skill demands. This is shown in Figure II, which plots the contrast between the flow of new work and the stock of preexisting work during 1940–1980 and 1980–2018 for noncollege workers (high school degree or below) and college workers (BA or above).[9] Between 1940 and 1980, the flow of new work largely replicated the stock of existing work. The subsequent four decades display a marked contrast: there is a sharp decline after 1980 in the flow of new work relative to the stock of existing work in traditional middle-skill, middle-pay occupations, including construction, transportation, production, and clerical and

9. Analogous to the procedure for Figure I, we estimate the average educational attainment of workers employed in new titles by taking the average of educational attainment of all workers in their macro-occupation-by-industry cell and weighting across all macro-titles to aggregate to broader occupation categories (where using respondents' macro-occupation and macro-industry to assign characteristics increases the specificity of the imputation).

administrative support. This pattern is consistent with the widely documented phenomenon of occupational polarization unfolding in these decades (Autor, Katz, and Kearney 2006; Goos and Manning 2007; Autor and Dorn 2013; Goos, Manning, and Salomons 2014; Michaels, Natraj, and Van Reenen 2014; Autor 2019).[10]

A third finding is that the pattern of polarization is distinct for college and noncollege workers. Among college-educated workers, the entirety of the decline in new middle-skill occupations is accounted for by a corresponding rise in employment in new professional occupations. Among noncollege workers, however, most of the decline in new middle-skill occupations is offset by a sharp increase in new low-paid personal and health service occupations, accompanied by a modest increase in new professional occupations. Thus, the emergence of new work over the past four decades has led the overall polarization of occupational structure. By implication, employment polarization does not merely reflect an erosion of employment in existing middle-skill work but also a change in the locus of new-work creation.

II.C. Measuring Augmentation and Automation Innovations

Critical to analyzing the relationship between innovation and new work is identifying innovations that potentially complement worker outputs (augmentation) and substitute for worker inputs (automation). We draw on three data sources for this task: the corpus of all U.S. utility patents issued between 1920 and 2018, the CAI, and DOT. The patent data provide a lexicon of potential augmentation and automation innovations for the period of interest. Patent data present the best substantively detailed, time-consistent measure of the flow of innovative activity in the U.S. economy, despite well-known limitations—including not capturing the totality of technological innovation (Moser 2012) and overrepresenting innovations that are difficult to protect as trade secrets.[11]

10. Online Appendix Figure A.VI documents that between the first and second halves of our sample, there is a sharp absolute decline in the flow of new work in traditional middle-skill, middle-pay occupations.

11. Patent texts were obtained by Kelly et al. (2021) from the USPTO patent search website for patents issued from 1976 to 2015, and from Google Patents prior to 1976. We extend the Kelly et al. (2021) sample of patents issued from 2015 to 2018 by scraping the patent text from the Google Patents website. We also scrape Google Patents for the issue date and primary three-digit Cooperative Patent Classification (CPC) code for each patent in our sample. Before 1976,

Identifying patents that may substitute for workers' occupational inputs (automation innovations) requires a text corpus that characterizes the tasks that workers perform in their jobs. For this purpose, we follow Webb (2020) and Kogan et al. (2021) in employing the DOT (U.S. Department of Labor, Employment and Training Administration 1939, 1977), which describes in detail the primary duties of workers in each occupational title. To avoid capturing occupational outputs, we use only task descriptions from the DOT, purging any occupational titles contained in these descriptions. Consistent with Webb (2020) and Kogan et al. (2021), we expect that patents that overlap with workers' occupational tasks reflect innovations that may potentially replace workers in these tasks.

Identifying patents that may complement workers' occupational outputs (augmentation innovations) requires a text corpus that characterizes the outputs of each occupation. We harness the CAI, which supplies tens of thousands of occupation and industry micro-titles in each decade's CAI. We hypothesize that patents that link to the CAI corpus reflect innovations that increase the capabilities, quality, variety, or utility of the outputs of occupations, potentially generating new demands for worker expertise and specialization. This step precisely parallels our method for identifying automation innovations, but we are not aware of any prior work that systematically characterizes innovations that are complementary to labor. The empirical success of this method (demonstrated below) is a key contribution of this article.

To link occupations to relevant patents, we follow Kogan et al. (2021), who measure textual similarities between patent texts and job descriptions using word embeddings, which are geometric representations of word meanings. Distinct from conventional measures of text similarity (e.g., the commonly used bag of words approach found in Gentzkow, Kelly, and Taddy 2019; Atalay et al. 2020; Kim 2022), word embeddings locate closely related words nearby to one another in embedding space. This is valuable because patent texts and occupational descriptions often use dissimilar terms for related concepts. As an example, the verbs *research* and *quantify* are not synonymous, nor do they

each patent document was a single block of text. Subsequently, patent texts were divided into abstract, description, and claims sections. We use the entire text of each patent.

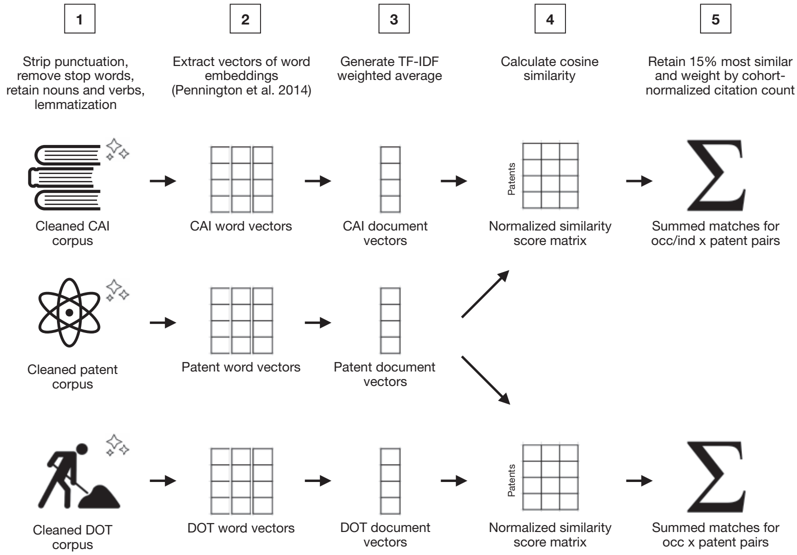

FIGURE III

Linking Patents to Occupations and Industries

This figure summarizes our five-step procedure for linking patents to occupations. The middle row describes the method for cleaning and processing patent data in preparation for creating the linkages. The top row depicts the method for creating augmentation-based patent matches using occupation and industry titles from the CAI. The bottom row depicts the method for creating automation-based patent matches using DOT task descriptions.

have common etymology, but each is among the three closest relatives of the other in English language embedding space.[12]

Figure III provides a schematic overview of this procedure, which we detail in Online Appendix E. Our procedure cleans each patent, CAI, and DOT text document by removing prepositions and other stop words and retaining nouns and verbs, and then represents each document as a term frequency-inverse document frequency (TF-IDF) weighted average of the word embeddings for each word in the document. It next calculates the matrix of cosine similarity scores among the document vectors for each patent issued in a decade and the CAI and DOT document vectors for each occupation (CAI and DOT) and industry (CAI)

12. Calculated using http://vectors.nlpl.eu/explore/embeddings/en/associates/# on the English Wikipedia corpus with queries `research_VERB` and `quantify_VERB` on June 24, 2022.

cell, and computes industry-occupation (CAI) similarity scores by taking an average of the patent-industry CAI similarity score, and the patent-occupation CAI similarity score. Lastly, it retains the top 15% most similar patent-occupation (DOT) and patent-industry-occupation (CAI) matches in a decade, and then sums the weighted count of top 15% matched patents for each occupation-decade (or industry-occupation-decade), with weights equal to each patent's cohort-specific relative citation frequency. We use citation data obtained from the USPTO PatentsView database for patents issued since 1976; for earlier patents, we rely on a database of historical patent citations introduced in Berkes (2018).[13] Because it is infeasible to construct a fully balanced panel of detailed occupations over 1940–2018 without sacrificing substantial resolution, we instead create two balanced panels that cover the first and second halves of our sample. Online Appendix C provides details, and Online Appendix Table A.VI reports descriptive statistics on these augmentation and automation exposure measures.[14] We emphasize that the links we establish between patents and occupations have a many-to-many structure: a given patent may have an augmentation link to one occupation and an automation link to another (or even to the same occupation to which it is augmentation-linked). Most occupations are linked to dozens—and in some cases, hundreds or thousands—of patents in each decade, so no single patent plays an outsized role in determining the augmentation and automation exposure of any given occupation. It is the totality of augmentation and automation links drawn for each occupation that provides identification.

II.D. Occupational Exposure to Innovation

Figure IV previews the substantive content of the automation and augmentation exposure measures by plotting the bivariate relationship between percentiles of each at the level of consistent three-digit (macro) occupations over 1940–1980 (166 occupations) and 1980–2018 (306 occupations), where the occupation-level

13. We thank Enrico Berkes for making these data available to us.

14. When estimating occupation-level models, we establish patent linkages to CAI occupation titles alone; when estimating occupation-by-industry models, we also leverage CAI industry titles to establish patent matches. Unlike the CAI, the DOT contains only occupation-level textual information, so automation exposure measures are always defined at the occupation level.

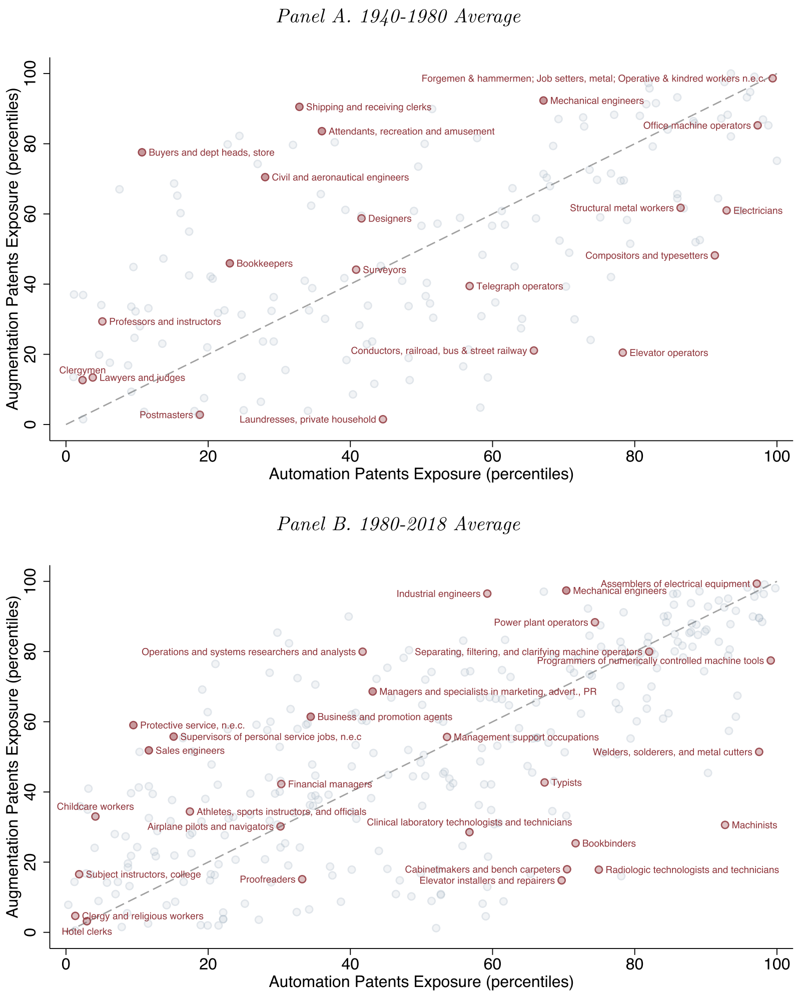

FIGURE IV

The Relationship between Exposure to Automation and Augmentation Patents at the Occupation Level

This figure presents a scatter plot of the relationship between occupational exposure to automation and augmentation patents for 1940–1980 (Panel A) and 1980–2018 (Panel B). Each point corresponds to the average percentile of automation (x-axis) and augmentation (y-axis) exposure of one consistently defined three-digit census occupation, where the average is taken over 1940–1980 ($N = 166$ occupations per year) in Panel A and over 1980–2018 ($N = 306$ occupations per year) in Panel B. The 45-degree line in each panel is plotted with dashes.

augmentation and automation exposure are averaged in the four-decade periods. Immediately evident from the figure is that occupations that are more exposed to augmentation are on average also more exposed to automation. The employment-weighted cross-occupation correlation between augmentation and automation exposure is 0.74 over 1980–2018 and 0.63 over 1940–1980. This positive correlation is expected as many technologies contain both automation and non-automation components. It is also consistent with our conceptual framework, sketched below, which highlights that demand forces that raise the value of occupational output create incentives for the introduction of both augmentation and automation innovations. (This may partly reflect the limitations of our classification procedure for cleanly distinguishing between augmentation and automation innovations.)

While the positive occupational correlation between augmentation and automation exposure is prominent in Figure IV, the off-diagonal components of this figure ultimately provide identifying variation for the analysis that follows. Between 1980–2018, the macro-occupations of Radiologic technologists and technicians, Cabinetmakers and bench carpenters, and Machinists all had high rates of automation relative to augmentation exposure. The same applies to Compositors and typesetters, Elevator operators, and Telegraph operators in 1940–1980. Our conceptual framework predicts that employment in these occupations would tend to erode. Conversely, augmentation outpaced automation between 1980 and 2018 in the occupations of Industrial engineers, Operations and systems researchers and analysts, and (to a lesser degree) Managers and specialists in marketing, advertising, and public relations and Business and promotion agents. In the prior four decades, Shipping and receiving clerks, Buyers and department heads, Civil and aeronautical engineers, and, to a lesser extent, Bookkeepers were all more exposed to augmentation than automation. We would expect these occupations to expand in these subperiods. Conversely, the on-diagonal examples capture occupations with a relatively balanced degree of exposure to both automation and augmentation, either because they are highly subject to both forces (e.g., Assemblers of electrical equipment over 1980–2018 and Office machine operators over 1940–1980) or because they are relatively insulated from both (e.g., Clergy and religious workers in both periods).

Inspection of patents that are augmenting for one occupation but automating for another reveals further logical relationships. For example, patents issued between 2010–2018 that are augmentation-linked to the macro-occupation Computer systems analysts & computer scientists are most often automation-linked to occupations in technicians or clerical and administrative support—occupations that are particularly susceptible to software-based automation during this period of rapid digital growth. Online Appendix Table A.I enumerates examples of individual patents that are augmenting for Computer systems analysts & computer scientists yet automating for other occupations. Example patents include "Methods and apparatus for storing confidential information," which is automation-linked to Billing clerks & related financial records processing; "Direct connectivity system for healthcare administrative transactions," which is automation-linked to Health record technologists & technicians; and "System and method for securing data," which is automation-linked to Office machine operators, computer and peripheral equipment operators, and other telecom operators. We expect that technologies developed in these patents can be harnessed by computer scientists to automate tasks in linked occupations.

Online Appendix Figure A.IV provides insight into augmentation and automation potential at the technology level, distinct from the occupation-level as above. This figure plots the augmentation shares of all linked patents within each three-digit CPC technology class on a zero to one scale, where a share of zero indicates all class patents are automation innovations and a share of one indicates that all class patents are augmentation innovations. These share data are organized into nine broad technology classes and are reported separately for 1940–1980 and 1980–2018. The broad class of Instruments and Information technologies contains on average the most augmenting detailed three-digit CPC classes, followed by Electricity and electronics, and Consumer goods and entertainment. Conversely, Engineering, construction, and mining is the most automating broad class, followed by Chemistry and metallurgy, Health and agriculture, and Manufacturing process. There is also substantial variation in augmentation shares within broad classes and over time: innovations in Transportation and in Lighting, heating, and nuclear have become less augmenting in the most recent four decades relative to the prior four, whereas the opposite is true for Instruments and information, and Electricity and electronics. As detailed in Section IV.B, these secular shifts in the locus of patenting across broad technology classes that have occurred since 1940 affect the intensity of augmentation and automation innovations and shape the set of occupations exposed to these forces.

III. Theoretical Motivation

To formalize the intuitions that guide our empirical analysis, we provide a model in Online Appendix J that considers how three forces shape the endogenous creation of new job tasks, the elimination of old tasks, and the demand for labor at the level of occupations. These forces are augmentation, which generates new labor-using job tasks; automation, which reallocates existing tasks from labor to capital; and shifts in consumer demand, which affect task automation and new-task creation by changing innovation incentives. Our framework generalizes the single-sector setting in Acemoglu and Restrepo (2018) to two sectors with different skill intensities. This serves two purposes. First, it allows us to consider the implications of augmentation and automation for occupational demand (where sectors represent occupations) rather than exclusively for aggregate labor demand. Second, the interaction of these sectors is central for considering the impact of demand shifts on new-work creation.

We posit that new job tasks (i.e., new work) derive from two primary sources. A first is augmentation innovations, meaning new processes and products (e.g., solar voltaic cells), new services (e.g., fingernail hardening), and entirely new products or industries (e.g., dry-cleaning, commercial air travel) that create new demands for expert knowledge and competencies, which are reflected in new job titles in our empirical analysis. Second, even absent specific technological advances, our conceptual framework implies that by raising the value of occupational outputs, positive demand shocks—stemming, for example, from demographic shifts—create incentives for entrepreneurs to introduce both augmentation and automation innovations.[15] Formally, sectoral

15. For example, the Bureau of Labor Statistics reports that solar energy installation managers (O*NET code 47-1011.03) is a recently emerging occupation whose primary task is to "Direct work crews installing residential or commercial solar photovoltaic or thermal systems" in the rapidly expanding renewable energy sector (U.S. Department of Labor: Employment and Training Administration 2022). This title is in turn a specialization of a broader occupational category, first-line supervisors of construction trades and extraction workers (O*NET code 47-1011.00), that is also common in this sector. This title exemplifies the intuition that sectoral demand shifts generate both more work—here greater demand

demand shocks in the model raise the value of sectoral output, spurring entrepreneurs to introduce both new augmentation and automation innovations in that sector. Further, by generating sector-specific capital scarcity, these shocks create an additional incentive for entrepreneurs to prioritize new augmentation innovations, since these are labor-using and capital-saving (as in Acemoglu and Restrepo 2018).

These mechanisms give rise to two empirical implications:

i. Augmentation innovations spur the emergence of new work, whereas automation innovations do not. This asymmetry has both a conceptual and empirical foundation. Conceptually, automation should primarily displace workers from existing tasks rather than directly creating new work tasks. Empirically, our new-work measure enumerates new job titles as they are captured by the Census Bureau, but it does not detect titles that go out of common usage since these are not systematically purged by the Census Bureau. We test for these asymmetric effects in Section IV.

ii. Positive demand shocks spur the introduction of new occupational titles and, further, generate a positive cross-occupation correlation between the flow of augmentation and automation innovations (visible in Figure IV).

A third set of predictions concerns the relationship between innovation and occupational labor demand. Augmentation innovations unambiguously raise occupational labor demand in our conceptual framework since they both create new labor-using tasks that raise (relative) demand for labor (a substitution effect) and increase the value of occupational services, yielding a complementary scale effect. Conversely, substitution and scale effects are partly offsetting in the case of automation innovations: automation of labor-using tasks reduces employment via the substitution effect while the scale effect (stemming from higher productivity) pushes in the opposite direction (as in

for first-line supervisors—and new work, more specialized sectoral expertise, here solar energy installation managers. For negative demand shifts—stemming from import competition, for example—we expect the opposite effect, such that new-work emergence decelerates. We test for the link between demand shifts and new-work emergence, holding constant the direct effect of augmentation innovations, in Section IV.C.

Acemoglu and Restrepo 2018). The structure of consumer preferences in our model guarantees that the substitution effect dominates: automation innovations spur a relative contraction in labor demand in exposed occupations, opposite to the case of augmentation innovations. Section V tests whether, despite their positive cross-occupation correlation, augmentation and automation innovations have measurable, countervailing effects on occupational employment and wage bills (i.e., the product of employment and wages).

This framework implies that the positive effects of augmentation flows and demand shifts on the emergence of new work are causal, as are the countervailing effects of augmentation and automation flows on occupational labor demand. We employ two instrumental variable strategies to test causality. To assess the causal effects of demand shocks on new-title emergence, we exploit exogenous occupational labor demand shifts stemming from trade shocks and demographic changes. To estimate the causal effect of augmentation and automation innovations on occupational labor demand, we isolate unanticipated flows of occupational augmentation and automation innovations stemming from novel upstream innovations occurring in prior decades.

IV. The Effects of Innovation and Demand Shifts on New-Work Creation

This section empirically characterizes the forces that explain where and when new job tasks emerge by relating the emergence of new occupational tasks to the exposure of occupations to (i) augmentation innovations; (ii) automation innovations; and (iii) positive and negative demand shifts. Our focus here is on forces that affect the creation of new work, measured by the emergence of new titles. We take up the effects of new-work creation on occupational labor demand and employment in Section V.

IV.A. Do Augmentation Innovations Spur New Work?

In our conceptual framework, economic forces that complement occupational outputs lead to the demand for new specialties and expertise reflected in new tasks. One of those forces is augmentation. The first hypothesis that we test is that new job tasks emerge differentially in occupations that are more exposed to augmentation innovations. Using the flow of new job titles as a measure of the emergence of new work, we estimate models of the following form:

$$\ln E\left[\text{Newtitles}_{j,t}\right] = \beta_1 \text{AugX}_{j,t} + \beta_2 \text{AutX}_{j,t} + \beta_3 \frac{E_{j,t-10}}{\sum_j E_{j,t-10}} + \delta_t (+\delta_{J,t}), \tag{1}$$

where j indexes consistent census occupations, and t indexes decadal intervals.[16] The dependent variable is a measure of the flow of new work titles emerging in a consistent census occupation in a decade, and the independent variables of interest are $\text{AugX}_{j,t}$ and $\text{AutX}_{j,t}$, measuring occupational exposure to augmentation and automation innovations, respectively, as revealed by textual links between patents and the CAI (augmentation) and DOT (automation). The dependent and key independent variables in year t are measured as cumulative flows over the preceding decade: new titles observed in 1950 are those that emerge between 1940 and 1949; augmentation and automation exposure in 1950 are equal to the log count of linked augmentation and automation patents awarded over 1940–1949. Accordingly, β_1 and β_2 can be read as elasticities.

Decade fixed effects absorb temporal variation in the total number of new titles. We control for the employment share of each occupation j at the start of the decade ($E_{j,t-10}/\sum_j E_{j,t-10}$) to remove any mechanical association between new-title counts and relative occupational employment size.[17] In some specifications, we further add main effects for the 12 consistently defined broad occupational groups, indexed by J, and their interaction with decade fixed effects. All regressions are weighted by start-of-decade occupational employment shares. (All core results for new-title flows and occupational employment changes, found in Tables II, V, and VI, are qualitatively comparable and statistically robust when estimated without weights.) Based on our conceptual framework, we expect $\beta_1 > 0$: more augmentation-exposed occupations will add more new titles.

16. Because the CAI was largely not updated between 2000 and 2010, and was then substantially updated in 2018, the last time interval of our sample is 2000–2018 rather than 2010–2018.

17. Excluding this control leads to larger point estimates for both the augmentation measure and automation measure but does not qualitatively change the results.

TABLE II

OCCUPATIONAL NEW-TITLE EMERGENCE AND AUGMENTATION VERSUS AUTOMATION EXPOSURE, 1940–2018

	(1)	(2)	(3)	(4)	(5)
Panel A: 1940–2018					
Augmentation exposure	17.81***	21.46***		16.85***	21.02***
	(3.52)	(3.74)		(3.96)	(3.54)
Automation exposure			12.75**	1.89	2.35
			(3.93)	(4.52)	(4.07)
Panel B: 1940–1980					
Augmentation exposure	23.46***	27.23***		19.48***	26.11***
	(5.09)	(4.80)		(5.79)	(4.45)
Automation exposure			19.56***	8.15+	9.07+
			(4.21)	(4.85)	(4.87)
Panel C: 1980–2018					
Augmentation exposure	8.14*	12.35***		14.87***	13.86***
	(4.08)	(2.31)		(3.72)	(2.08)
Automation exposure			−1.21	−12.99*	−5.43
			(5.07)	(5.37)	(6.19)
Occ. emp. shares	X	X	X	X	X
Time FE	X		X	X	
Broad occ. × time FE		X			X

Notes. Dependent variable: occupational new-title count. Negative binomial models, coefficients multiplied by 100. N = 1,535 in Panel A; N = 628 in Panel B; N = 907 in Panel C. Twelve broad occupations are defined consistently across all decades. Standard errors clustered by occupation × 40-year period are in parentheses. Observations are weighted by start-of-period occupational employment shares. Augmentation and automation exposure measures correspond to the log of the weighted counts of matched patents. $^{+}p < .10$, $^{*}p < .05$, $^{**}p < .01$, $^{***}p < .001$.

The first panel of Table II reports estimates of equation (1) for the full eight decades of the sample, where all coefficients have been multiplied by 100 for readability. Because the outcome measure in this table, the count of new titles in a macro-occupation in a decade, contains zeros, we fit a negative binomial regression for count data.[18] In subsequent instrumental variables (IV) estimates of these models, we use a combination of log and linear dependent variable models because IV models for count data are less well developed. As will be apparent, our results are robust

18. The working paper version of this article (Autor et al. 2022) estimated count models using the inverse hyperbolic sine (IHS) transformation. Because subsequent work has underscored the sensitivity of the IHS transformation to the treatment of zeros (Chen and Roth 2024; Mullahy and Norton 2022), we use more traditional count models in the final article. This leaves our results substantively unchanged.

to these alternative ways of handling zeros. The first column of Table II reports a specification that excludes major occupation-by-decade main effects. The point estimate of 17.81 (std. err. = 3.52) implies that each 10% increment to augmentation exposure predicts an additional 1.8% faster rate of new-title emergence over the course of a decade. The second column adds 12 occupation dummies and their interactions with decade effects, and hence is identified by variation in new-title emergence rates across detailed occupations within broad occupational categories within decades. The point estimate of 21.46 (std. err. = 3.74) is slightly larger than in the prior column while precision is unaffected.

Our conceptual model implies distinct relationships between augmentation versus automation innovations and the flow of new work: augmentation innovations generate new labor-using tasks, reflected in new titles; automation innovations do not generate such tasks. The next columns of Table II test this prediction. Column 3 removes the augmentation exposure variable (AugX) and replaces it with the automation exposure variable (AutX) using the specification in column (1). Entered independently, the flow of automation innovations also predicts the flow of new titles, with a point estimate on AutX of 12.75 (std. err. = 3.93). This pattern is expected given that automation and augmentation exposures are strongly positively correlated across occupations (see Figure IV). When both innovation measures are simultaneously included in columns (4) and (5), augmentation patents continue to robustly predict the flow of new titles ($\hat{\beta}_1 = 16.85\,(3.96)$ in column (4), and $\hat{\beta}_1 = 21.02\,(3.54)$ in column (5)) whereas automation patents have a small and statistically insignificant relationship with new-title flows ($\hat{\beta}_2 = 1.89\,(4.52)$ in column (4), and $\hat{\beta}_2 = 2.35\,(4.07)$ in column (5)).

Recognizing that count data models may be somewhat functional-form dependent, we present a number of alternative estimators in Online Appendix Table A.II, including Poisson estimates in Panel A; OLS log count estimates (the intensive margin, excluding zeros) in Panel B; and a linear probability model for the occurrence of any new title (the extensive margin) in Panel C. These models support the findings in Table II: augmentation exposure strongly predicts the emergence of new titles on both the intensive and extensive margins; this relationship is generally strengthened by including broad occupation by time effects (drawing identification from within broad occupation by decade comparisons); conditional on augmentation exposure, automation exposure has no robust predictive relationship with new-title emergence and, moreover, is always insignificant and close to zero when controlling for broad occupation by decade effects. In Online Appendix I we further show that our findings are robust to instead using Lin's (2011) measure of new titles for each of the three overlapping sample decades (1980–2000). The comparability of findings is reassuring because for the 1980s and 1990s, Lin (2011) identifies new titles from the DOT rather than the CAI and obtains newly added titles directly from a Conversion Table of Code and Title Changes available for these two decades only.

Leveraging the fact that we have distinct occupational panels for the two halves of our sample, we report estimates for 1940–1980 and 1980–2018 in Table II, Panels B and C. These provide a consistent picture. In all cases, augmentation patents significantly predict faster arrival rates of new occupational job titles. Automation patents are either weakly positive or weakly negative predictors of new-title emergence when augmentation patents are simultaneously included (with one exception in Panel C, column (4)). The precision of the augmentation estimates rises (with almost no change in magnitude) when broad occupation category by decade dummies are included (columns (2) and (5)) to contrast new-title flows across detailed occupations within broad categories in each decade. Thus, despite marked differences in the locus of new-work emergence across these four-decade periods (documented in Figure II), augmentation innovations strongly predict where new titles emerge in both eras. A comparison of the point estimates for 1980–2018 versus 1940–1980 does reveal a substantially shallower, but still robustly significant, elasticity of new-title creation with respect to augmentation innovations in the latter half of the sample ($\hat{\beta}_1^{40-80} = 26.11\,(4.45)$ versus $\hat{\beta}_1^{80-18} = 13.86\,(2.08)$) in the final columns of Panels B and C. This pattern hints that the relationship between innovation and new-task creation may have slowed in recent decades relative to the first four decades after World War II. This deceleration of the new-work creation component of innovation will be a recurrent theme of our analysis, which we synthesize in Section V.B.[19]

19. A potential concern with comparing estimated elasticities across time periods is that the CAI's sensitivity at detecting new titles might fluctuate across decades. This should not bias elasticity estimates, however, if such fluctuations affect only the rate of new-title capture but not its distribution across

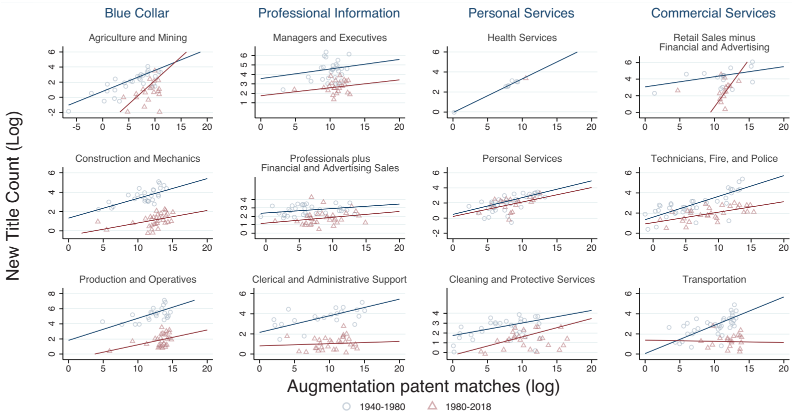

FIGURE V

The Relationship between Exposure to Augmentation Patents and Occupational New-Title Emergence

Each panel is a binscatter showing the relationship between augmentation exposure and new-title emergence within broad occupations. The x-axis plots the logarithm of augmentation patent matches, and the y-axis plots the logarithm of new-micro-title counts. Observations with zero new titles or zero augmentation innovations are dropped in this log-log plot. Circles correspond to data for 1940–1980, and triangles correspond to data for 1980–2018.

Figure V depicts the variation that drives these estimates with a set of binscatters plotting the (log of) the flow of augmentation innovations within the 12 broad occupational categories against the (log of) the flow of new occupational job titles, separately for 1940–1980 and 1980–2018. The predictive relationship between augmentation innovations and the arrival of new titles is a pervasive feature of the data. In 11 of 12 occupations (all but Transportation) and in both halves of the sample, there is a clear, upward-sloping relationship between augmentation innovations and new-title flows. (There are not enough detailed occupations with nonzero new titles to estimate this relationship for Health Service occupations in the 1940–1980 period.)

occupations. Formally, let the number of new occupational titles captured in decade t be proportional to the true (latent) number of new titles: $Y_{j,t} = \phi Y^*_{j,t}$ where $\phi > 0$. If the structural equation relating new-title flows to augmentation flows is $\ln Y^*_{j,t} = \alpha + \gamma \ln \text{Aug}X_{j,t}$ (as estimated above), then $\frac{\partial \ln Y_t}{\partial \ln \text{AugX}_t} = \frac{\partial \ln Y_t}{\partial \ln Y^*_t} \times \frac{\partial \ln Y^*_t}{\partial \ln \text{AugX}_t} = \gamma$.

IV.B. Testing Causal Relationships: Augmentation, Automation, and New Work

The new-title estimates reflect conditional correlations that may not have a causal interpretation. Our conceptual model, for example, implies that demand shifts favoring a given sector generate both new augmentation and new automation patents, where the former catalyzes the introduction of new titles. This raises the concern that the estimates are tainted by simultaneity bias stemming from unmeasured demand shifts. (Though, if variation in innovation rates and new-task creation is primarily driven by demand shifts, we would expect both augmentation and automation innovations to predict new-title emergence—which is not the case.) To more definitively distinguish causality from correlation requires isolating exogenous flows of augmentation and automation innovations. We sketch a strategy for doing so here.

Our IV strategy exploits the occurrence of breakthrough innovations (Kelly et al. 2021), which represent discontinuous advances in the innovation environment that materialize at a specific point in time. Breakthrough patents are both novel and impactful: novel in that they are conceptually distinct from their predecessors, relying less on prior art; impactful in that they influence future scientific advances, catalyzing a host of follow-on innovations. This chain of reasoning implies that breakthroughs, recorded in patents, can potentially be used to identify exogenous flows of follow-on innovations. The identification assumption is that the precise timing of breakthroughs is not anticipated (conditional on preexisting trends), while the arrival of follow-on innovations is causally affected by the timing of antecedent breakthroughs.

Our implementation of this idea harnesses Kelly et al.'s (2021) breakthrough-innovation measure, which operationalizes the novelty and impactfulness of patents using NLP tools. Kelly et al. (2021) compare the similarity of each U.S. utility patent granted relative to those that precede it (backward similarity) to measure its novelty, and compare it to those that follow it (forward similarity) to measure impact. A breakthrough patent is one that has a high ratio of forward to backward similarity, where each is measured over the 5 years before (backward) or 10 years after (forward) the patent date. Our analysis uses the top 10% of breakthrough patents granted in each decade, but

Panel A. The flow of top 10% breakthrough patents, 1900–2002

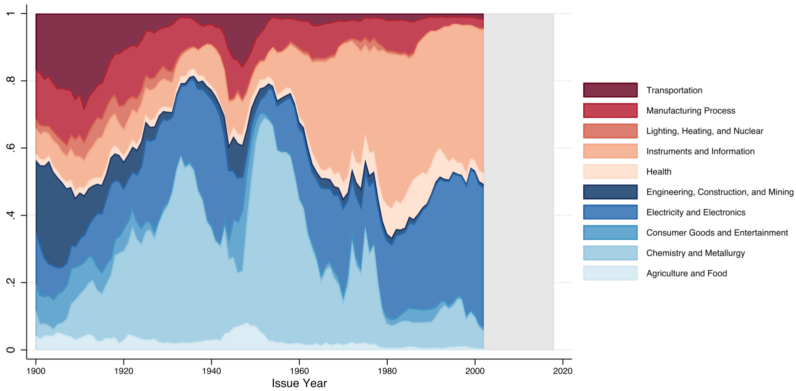

Panel B. The flow of all patents, 1900–2018

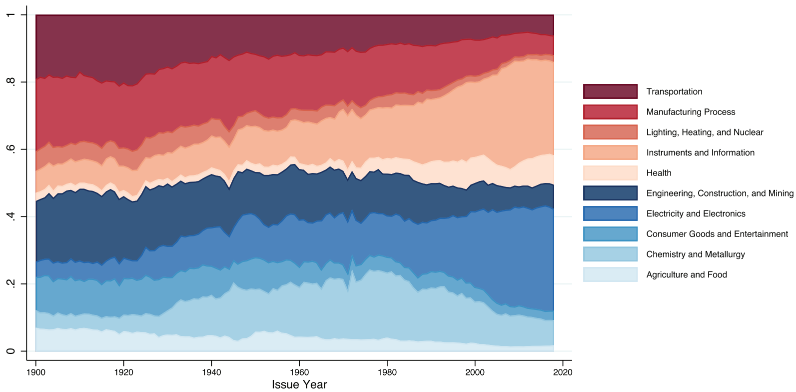

FIGURE VI

Breakthrough-Patent Flows by Class versus All Patent Flows by Class

This figure presents the distribution of breakthrough (Panel A) and total (Panel B) patents by 10 broad technological classes in each year between 1900 and 2002 (breakthroughs) and 1900 and 2018 (all patents). We use the classification of patent three-digit CPC classes into broad technology groups from Kelly et al. (2021), while excluding patents assigned to the broad classes "weapons" and "other," which together make up less than 1% of all patents.

our results are robust to a range of breakthrough thresholds and forward-/backward-similarity window lengths.

The two panels of Figure VI illustrate the substantive difference between breakthrough patents and the full population of patents (inclusive of breakthroughs). The figure reports stacked area plots of the distribution of breakthrough and non-breakthrough patents granted in each year between 1900 and 2002 across 10 broad, exhaustive, and mutually exclusive technology classes.[20] Panel A shows that the locus of breakthroughs has changed dramatically across decades. In the first two decades of the twentieth century, breakthroughs are concentrated in Engineering, Transportation, and Manufacturing processes. Over the next four decades, between 1920 and 1960, the locus of breakthrough innovations shifts increasingly to Chemistry and metallurgy. After 1960, two technology categories grow to dominate breakthrough innovations: Instruments and information, followed by Electricity and electronics. Panel B reports analogous data for overall patenting. Unlike the pattern for breakthrough patents in Panel A, the evolution of overall patenting across domains in Panel B is relatively muted and slow-moving. Nevertheless, consistent with the expectation that breakthrough patents shift the tide of innovation across domains, the class distribution of the full set of patents appears to echo that of breakthroughs with a lag of approximately two decades.

Figure VII, Panel A confirms this pattern by plotting the predictive relationship between breakthroughs within 119 three-digit detailed CPC technology classes between 1920 through 2000 and subsequent patenting within that class two decades later.[21] This figure documents that prior breakthroughs are a strong predictor of subsequent patent flows, even conditional on prior non-breakthrough patent flows. The nine binscatters in the figure—one for each broad technology class, each plotted with a distinct marker—reveal a consistent upward-sloping relationship between detailed class-level patent flows and breakthrough patents occurring in those classes two decades earlier, with an overall slope of $\hat{\beta}_{t-20} = 0.42$ (0.03).[22]

20. These classes are Transportation; Manufacturing process; Lighting, heating, and nuclear; Instruments and information; Health; Engineering, construction, and mining; Electricity and electronics; Consumer goods and entertainment; Chemistry and metallurgy; and Agriculture and food.

21. The predictive relationship is estimated by regressing log patent flows by technology class and decade on log breakthrough patent flows lagged by two decades, controlling for log non-breakthrough patent flow counts (also lagged by two decades) and decade fixed effects. We aggregate several small CPC classes to arrive at 119 classes.

22. The broad classes of Health and Agriculture are combined because the former broad class contains only two detailed three-digit classes.

Panel A. Correlation between patent flows and lagged breakthroughs by technology class

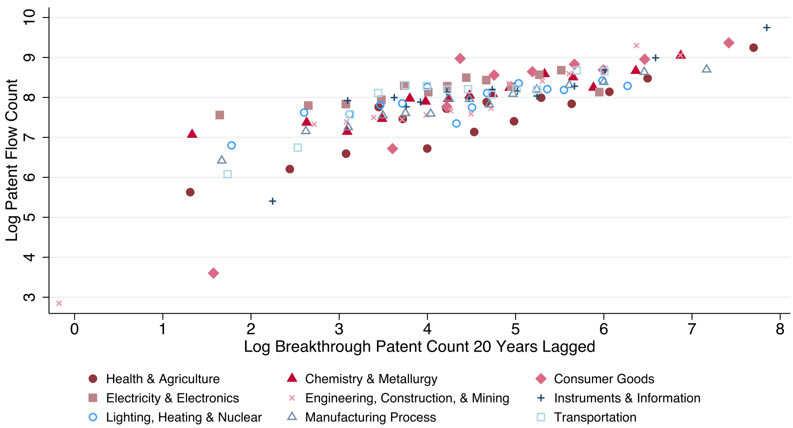

Panel B. Correlation between patent flows and future breakthroughs by technology class

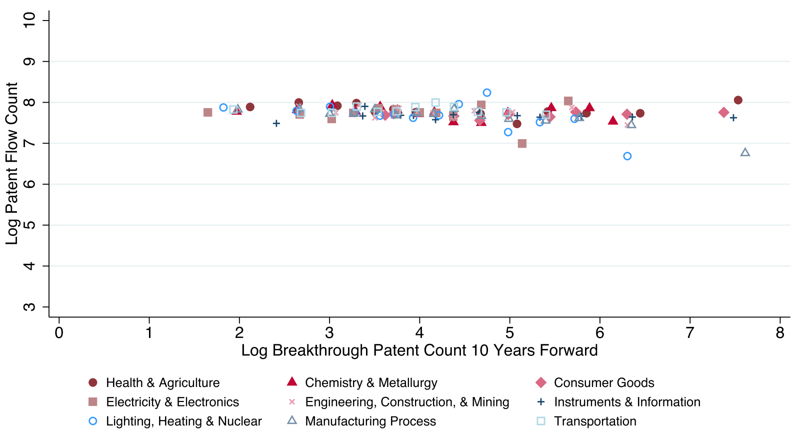

FIGURE VII

Predictive Relationship between Breakthrough Patents and 20-Year Forward Patent Flows and 10-Year Backward Patent Flows by Technology Class

This figure presents conditional correlations between breakthrough technology flows and subsequent patent flows in Panel A, and prior patent flows in Panel B. Panel A regresses log patent flows by technology class and decade on log breakthrough patent flows lagged by two decades, while controlling for log non-breakthrough patent flow counts lagged by two decades and decade fixed effects: $\hat{\beta} = 0.42$, std. err. $= 0.03$, $N = 1{,}051$. Panel B regresses log patent flows by technology class and decade on log breakthrough patent flows one decade forward, while controlling for log non-breakthrough patent flow counts one decade forward and decade fixed effects: $\beta = -0.02$, std. err. $= 0.02$, $N = 1{,}044$. Standard errors are clustered by three-digit technology class.

Might this pattern reflect existing trends whereby breakthroughs occur in already active patenting classes? Figure VII, Panel B replaces the 20-year lagged breakthrough predictor with the 10-year lead of breakthroughs in each class, controlling for the corresponding (log) count of non-breakthrough patents as well as decade fixed effects. The estimate of $\hat{\beta}_{t+10} = -0.02$ (0.02) is small, imprecise, and opposite in sign to the predictive relationship in Panel A. This is consistent with the hypothesis that breakthroughs spur subsequent downstream innovations in the technology classes in which they originate, whereas flows of contemporary non-breakthrough innovations do not spur subsequent breakthroughs in the same classes.

We exploit this mechanism to isolate exogenous variation in the augmentation and automation patent flows to which each occupation is exposed. This variation provides instrumental variables for the endogenous flows of augmentation and automation innovations used above.

i. Following the textual matching procedure detailed in Section II.C, we form automation and augmentation links for each occupation (and later, for industry-occupation cells) using patents lagged by two decades and occupational description texts lagged by zero to three decades, depending on the data source.[23]
ii. For each set of occupation-patent matches from the prior step, we sum the matched patents from two decades earlier that are among the top 10% of breakthrough patents for the decade. This is done twice, once for augmentation matches, $N^{Aug,10}_{j,t-20}$, and once for automation matches, $N^{Aut,10}_{j,t-20}$.

23. Recall that the automation links for 1940–1980 are based on the DOT 1939 volume, while automation links for 1980–2018 are based on the DOT 1977 volume. Implicitly, these automation links are lagged by zero to three decades. Augmentation textual links, by contrast, are based on CAI volumes that are lagged by one decade. Absent a consistent occupation scheme for bridging CAI volumes prior to 1930—which would require considerable coarsening of the data—we cannot apply a longer lag. As documented in Online Appendix Table A.II, Panel D, temporal variation across CAI volumes is not ultimately central to the predictive relationship between patent flows and new-title emergence: we obtain similar results when linking patents to the initial CAI volume for each four-decade subperiod—similar to the DOT. We also obtain similar results when removing new titles from each decade's CAI before obtaining augmentation-linked patents, as shown in Online Appendix Table A.II, Panel E and Online Appendix Table A.XII.

iii. Occupations are then classified as augmentation breakthrough-exposed if they are above the median of the matched augmentation breakthrough count measure, and similarly for automation breakthrough exposure: $B^{Aug,10}_{j,t-20} = \mathbb{1}\left[N^{Aug,10}_{j,t-20} > \tilde{N}^{Aug,10}_{j,t-20}\right]$, $B^{Aut,10}_{j,t-20} = \mathbb{1}\left[N^{Aut,10}_{j,t-20} > \tilde{N}^{Aut,10}_{j,t-20}\right]$, where a tilde denotes the decadal employment-weighted cross-sectional median of a variable in the regression sample. Because manufacturing industries are intrinsically patent intensive, we compute medians separately for groups of occupations that have above- versus below-50% start-of-period employment in manufacturing.

iv. To similarly account for non-breakthrough patent exposure, we perform a parallel procedure that identifies occupations that are highly exposed to all lagged matched patents in each decade: $B^{Aug,100}_{j,t-20} = \mathbb{1}\left[N^{Aug,100}_{j,t-20} > \tilde{N}^{Aug,100}_{j,t-20}\right]$, $B^{Aut,100}_{j,t-20} = \mathbb{1}\left[N^{Aut,100}_{j,t-20} > \tilde{N}^{Aut,100}_{j,t-20}\right]$.

v. The IV models estimate a variant of equation (1) for the effect of automation and augmentation exposure on the flow of occupational new titles, where $B^{Aug,10}_{j,t-20}$ and $B^{Aut,10}_{j,t-20}$ serve as instruments for the endogenous variables $\text{AugX}_{j,t}$ and $\text{AutX}_{j,t}$. In this specification, $B^{Aug,100}_{j,t-20}$ and $B^{Aut,100}_{j,t-20}$ serve as controls for overall prior augmentation and automation exposure. As highlighted above, we use 2SLS models for this exercise, distinct from the negative-binomial count models reported in Table II. (Recall that OLS variants of the new-title models presented in Online Appendix Table A.II strongly corroborate the new-title models using count data models.)

Table III reports first-stage estimates. Columns (1) and (2) tabulate models with only one endogenous variable (AugX in column (1), AutX in column (2)) and the corresponding breakthrough instruments. Since the standard deviation of our augmentation exposure measure is 3.45 (see Online Appendix Table A.VI), the first-stage coefficient in column (1) of 2.26 (std. err. = 0.41) indicates that, accounting for overall occupational patent exposure, occupations that are above the median of augmentation breakthrough exposure experience an increase of 0.64 standard

TABLE III

First-Stage Estimates for New-Titles Regressions, 1940–2018

	(1)	(2)	(3)	
	Aug	Aut	Aug	Aut
Panel A: 1940–2018				
Augmentation IV	2.26***		2.20***	0.23
	(0.41)		(0.43)	(0.32)
Automation IV		2.40***	0.67	2.37***
		(0.32)	(0.65)	(0.36)
F-stat.	30.88	56.59	27.50	28.96
Sanderson-Windmeijer F-stat.	30.88	56.59	49.47	35.39
Panel B: 1940–1980				
Augmentation IV	1.17+		1.70**	−0.01
	(0.62)		(0.54)	(0.33)
Automation IV		3.45***	2.19+	3.58***
		(0.66)	(1.16)	(0.75)
F-stat.	3.63	27.21	8.20	11.46
Sanderson-Windmeijer F-stat.	3.63	27.21	8.94	9.29
Panel C: 1980–2018				
Augmentation IV	2.82***		2.94***	0.78**
	(0.35)		(0.48)	(0.30)
Automation IV		1.66***	−0.82	1.35***
		(0.20)	(0.73)	(0.23)
F-stat.	65.73	68.15	40.91	29.68
Sanderson-Windmeijer F-stat.	65.73	68.15	62.03	30.37
Occ. emp. shares	X	X	X	X
Time FE	X	X	X	X

Notes. Dependent variable: log patent count. N = 1,276 in Panel A; N = 589 in Panel B; N = 687 in Panel C. First-stage estimates for columns (1), (3), and (4) are in Table IV. Specifications additionally control for indicators for being above the cross-sectional median in total augmenting or automating patent matches two decades prior, and whether the occupation has greater than 50% employment in manufacturing industries. Standard errors clustered by occupation × 40-year period are in parentheses. Observations are weighted by start-of-period occupational employment shares. Augmentation and automation exposure measures correspond to the log of the weighted counts of matched patents. $^{+}p < .10$, $^{*}p < .05$, $^{**}p < .01$, $^{***}p < .001$.

deviations in augmentation exposure two decades later. Column (2) reports a similar relationship for automation exposure, with a first-stage coefficient 2.40 (std. err. = 0.32). The third specification in the table, which spans two columns, reports the first stages for augmentation and automation patents estimated jointly. The first-stage coefficients for augmentation and automation in this joint specification are closely comparable to the standalone estimates in column (2) while the cross-instrument predictive relationships are surprisingly modest. In particular,

$B^{Aut,10}_{j,t-20}$ has essentially no predictive power for subsequent augmentation patent flows $\mathrm{AutX}_{j,t}$ while, similarly, $B^{Aug,10}_{j,t-20}$ is not predictive of subsequent augmentation patent flows $\mathrm{AutX}_{j,t}$. The instruments are able to isolate independent variation in augmentation and automation patent flows.

Panels B and C report separate first-stage estimates for the two halves of the sample. While there is strong identification of both endogenous variables in the pooled full-sample specification, the instruments have a weaker first stage in the first half of the sample, 1940–1980, with F-statistics of 8.20 and 11.46 on augmentation and automation, respectively. We treat the second-stage estimates for this period with due caution. In the latter half of the sample, both instruments play a role in identifying each endogenous variable, though the within-category coefficients remain larger and more precise than the cross-category coefficients.

Table IV reports 2SLS estimates of the effects of augmentation and automation patents on the emergence of new work between 1940 and 2018. Instrumenting augmentation patent flows using breakthrough patents from two decades earlier, column (1) obtains a coefficient of $\hat{\beta}^{Aug}_{40-18} = 24.42\,(7.46)$, implying that a 10% increase in augmentation patenting exposure increases new-title emergence by about 2.4% over the course of a decade. When broad occupation by decade effects are added in column (2), this point estimate falls modestly to $\hat{\beta}^{Aug}_{40-18} = 21.30\,(5.81)$ and precision increases. Column (3) removes the augmentation exposure variable and adds the automation exposure measure in its place. Distinct from the OLS estimates, automation patents do not by themselves positively predict the flow of new titles—the point estimate is small and imprecise ($\hat{\beta}^{Aut}_{40-18} = -3.01\ (11.82)$). When instrumented augmentation and automation patent flows are both included in columns (4) and (5), the estimated effect of augmentation patents on new-title emergence is comparable in magnitude and precision to the estimates in earlier columns, while the coefficient on automation patents is imprecise, with an inconsistent sign in both columns. These 2SLS results are overall substantively similar to their non-instrumented counterparts in Table II. (Unreported placebo regressions that instead use IV dummies for exposure to the bottom 10% of augmenting or automating innovations—the least innovative patents—have weaker first-stage predictive power and no clear second-stage relationship.)

TABLE IV

2SLS Estimates for New-Titles Regressions, 1940–2018

	(1)	(2)	(3)	(4)	(5)
Panel A: 1940–2018					
Augmentation exposure	24.42**	21.30***		29.54***	19.52**
	(7.46)	(5.81)		(8.58)	(6.66)
Automation exposure			−3.01	−14.56	4.21
			(11.82)	(11.29)	(9.71)
F-stat. (aug)	30.88	60.16		27.50	37.16
F-stat. (aut)			56.59	28.96	52.05
Panel B: 1940–1980					
Augmentation exposure	64.64*	30.17***		56.79**	24.82*
	(29.15)	(8.79)		(19.88)	(11.78)
Automation exposure			9.80	−17.29	15.75
			(15.16)	(17.23)	(17.74)
F-stat. (aug)	3.63	11.04		8.20	8.72
F-stat. (aut)			27.21	11.46	13.79
Panel C: 1980–2018					
Augmentation exposure	7.80+	10.50+		21.08**	15.36*
	(4.67)	(6.31)		(7.11)	(5.93)
Automation exposure			−22.70	−26.73+	−13.43
			(14.99)	(13.60)	(9.80)
F-stats (aug)	65.73	52.32		40.91	37.44
F-stats (aut)			68.15	29.68	35.90
Occ. emp. shares	X	X	X	X	X
Time FE	X		X	X	
Broad occ. × time FE		X			X

Notes. Dependent variable: log occupational new-title count. N = 1,276 in Panel A; N = 589 in Panel B; N = 687 in Panel C. Coefficients are multiplied by 100. Specifications additionally control for indicators for being above the cross-sectional median in total augmenting or automating patent matches two decades prior, and whether the occupation has greater than 50% employment in manufacturing industries. Standard errors clustered by occupation × 40-year period are in parentheses. Observations are weighted by start-of-period occupational employment shares. Augmentation and automation exposure measures correspond to the log of the weighted counts of matched patents. $^{+}p < .10$, $^{*}p < .05$, $^{**}p < .01$, $^{***}p < .001$.

Following the format of Table II, Panels B and C of Table IV report separate estimates for the two four-decade halves of the sample. Focusing for brevity on the final column of results (column (5)), the estimated causal effect of augmentation innovations on new-title emergence is positive and statistically significant in both 1940–1980 and 1980–2018. Consistent with earlier findings, the point estimates for automation innovations are generally imprecise and not consistently signed. Also echoing earlier results, the point estimates for the effect of augmentation innovations on new-title flows are considerably larger in the earlier 1940–1980 period ($\hat{\beta}^{Aug}_{40-80} = 24.82\,(11.78)$) than in the later 1980–2014 period ($\hat{\beta}^{Aug}_{80-18} = 15.36\,(5.93)$). We offer the caveat that the estimates for the earlier period have a weaker (though still conventionally robust) first stage and hence should be viewed somewhat cautiously.

While the negative binomial analysis in Table II simultaneously accounts for intensive-margin (number of new titles) and extensive-margin (any new titles) responses to augmentation and automation, the 2SLS models compel us to analyze the two separately. Online Appendix Table A.IX explores the extensive-margin effects of augmentation and automation patent flows for the 1940–2018 period, with corresponding first-stage estimates in Online Appendix Table A.VIII. In these linear probability models, the dependent variable is a dummy equal to one if an occupation has a positive count of new titles in the relevant decade (and zero otherwise). The first-stage estimates again reveal strong identification in this slightly larger sample (where occupations with zero new titles are retained). The second-stage estimates for the effect of augmentation patents on new-title flows are positive but relatively imprecise, with the exception of column (2) ($\hat{\beta}^{Aug}_{40-18} = 1.99\,(0.97)$). Similarly, the estimated slopes on automation innovations are imprecise and vary in magnitude across columns. It bears noting that the vast majority of the variation in new-title emergence occurs on the intensive margin: 83% of occupation-year observations have at least one new title, and when weighting by employment shares (as is done in the regressions), this number rises to 94%. It is therefore perhaps unsurprising that the dummy IV linear probability model does not identify this effect precisely. Nevertheless, the OLS and 2SLS estimates for the extensive-margin are broadly consistent: OLS extensive-margin estimates reported earlier in Online Appendix Table A.II, Panel C find quantitatively comparable—but far more precise—patterns than their 2SLS counterparts. (As with the intensive-margin results, 2SLS estimates of the impact of augmentation are slightly larger than corresponding OLS models.)

In summary, 2SLS estimates corroborate earlier non-instrumented results: new augmentation innovations catalyze the emergence of new work across occupations and over time, whereas new automation innovations have no similar effect, suggesting that these measures capture economically distinct components of innovation. One skeptical interpretation is that automation patents are simply a noisy measure of augmentation patents, which causes the automation measure to be significant on its own but nonrobust to the inclusion of augmentation. Militating against this interpretation is that automation patents have significant independent and opposite-signed (relative to augmentation patents) predictive power for employment and wage bill growth in OLS and 2SLS models, as shown in Section V. We view the distinct relationships between augmentation, automation, and new-title flows as (provisional) confirmation of our hypotheses.

The 2SLS estimates also echo another finding from the OLS models: the new-work-generating effects of augmentation innovations appear to have slackened over time—indeed, the point estimates are less than half the size in 1980–2018 than in 1940–1980. This time pattern could potentially reflect a change in the sensitivity of the Census Bureau's procedure for capturing new titles in the CAI. There are two reasons to suspect otherwise. First, the identification of $\hat{\beta}$ stems from cross-occupation comparisons in the flow of augmentation patents and new titles—and there is no expectation that a change in CAI data-collection procedures would weaken these cross-occupation correlations. Second, our evidence below for employment and wage bills—where measurement is highly stable across decades—shows a similar evolution in the latter four decades of the sample.

IV.C. Demand Shocks and New-Work Creation

While technological innovations are one force contributing to new-work creation, our conceptual framework implies that occupational demand shifts also shape the flow of new work: positive demand shifts spur innovations that generate new occupational specialties, and negative demand shifts slow the rate of new-work emergence. We test these predictions, first by exploiting the well-documented China trade shock (Autor, Dorn, and Hanson 2016) to identify the impact of adverse demand shocks on the flow of new titles; and second, by exploiting demographic shifts, following DellaVigna and Pollet (2007), to identify the effect of positive demand shocks on the flow of new titles. In both cases, we identify occupations' differential exposure to positive and negative demand shifts by leveraging their employment distributions across more versus less exposed industries. Although both inward and outward demand shifts are expected to affect the arrival rate of new work, we find it useful to apply countervailing tests because new work never flows in reverse and occupational titles are rarely removed from the CAI. In addition, the two demand shock instruments cover different time intervals.

1. *Trade Shocks and Demand Contraction.* Starting in the early 1990s, import competition from China generated a sizable negative demand shock to many labor-intensive domestic U.S. manufacturing industries (Bernard, Jensen, and Schott 2006; Autor, Dorn, and Hanson 2013, 2021; Autor et al. 2014; Pierce and Schott 2016; Acemoglu et al. 2016). These shocks directly affected product demand across industries and, since occupational employment shares differ across industries, indirectly affected labor demand across occupations. We use this variation to construct a measure of occupational exposure to the China trade shock, leveraging shifts in imports from China among a set of developed countries other than the United States.[24] This trade-exposure measure captures a plausibly exogenous shock to domestic occupational demand stemming from rising Chinese productivity and falling China-facing trade barriers between 1991 and 2014. Occupational exposure to import competition varies substantially both between and within production and non-production occupations over 1990–2014, as highlighted in Online Appendix Figure A.II.

2. *Demographics Shocks and Demand Expansion.* As a source of positive demand shifts, we follow DellaVigna and Pollet (2007) in exploiting changes in the demographic structure of the U.S. population to predict movements in industry-level demands, which in turn affect occupation-level demands.[25] For this exercise, we obtain predicted consumption across product categories for household members of different ages and combine these age-specific coefficients with population data to construct

24. Related work maps trade shocks to labor market outcomes in Brazil, Canada, India, Norway, Germany, Mexico, and elsewhere (Chiquiar 2008; Topalova 2010; Kovak 2013; Dauth, Findeisen, and Suedekum 2014; Balsvik, Jensen, and Salvanes 2015; Branstetter et al. 2019; Devlin, Kovak, and Morrow 2021). Our approach follows Autor, Dorn, and Hanson (2013) and subsequent papers, with the difference that we project the trade shock variation to the occupation level. Online Appendix F1 details the construction of this demand-shift measure.

25. Our work is related to Mazzolari and Ragusa (2013), Leonardi (2015), and Comin, Danieli, and Mestieri (2020), who analyze the causal effects of age, education, and income on employment via consumption.

occupational exposure as a function of changes in predicted consumption by industry over 1980–2018, which are in turn based on changes in population age structure. This demand measure accounts for interindustry input-output linkages, and hence corresponds to final demands. Online Appendix F2 provides details on the measure's construction.

3. *Demand Shifts and New-Work Emergence: Estimates.* We leverage these two sources of variation to estimate the effects of occupational demand shocks on the emergence of new occupational titles using the following count data model:

$$\ln E\left[\text{Newtitles}_{j,t}\right] = \beta_1^k \times \text{DemandX}_{j,t}^k + \beta_2 \text{AugX}_{j,t} + \beta_3 \sum_k \gamma_{jk,t} + X_{j,t}' \beta_4 + \delta_t, \tag{2}$$

where $\text{DemandX}_{j,t}^k = \sum_K \gamma_{jk,t} \hat{D}_{k,t}$ is the shift-share demand measure, equal to the sum of predicted log changes in industry employment $\hat{D}_{k,t}$ in industries k multiplied by the start-of-decade share of employment in occupation j contained in industry k, denoted by $\gamma_{jk,t}$. The sum of occupational exposure weights, $\sum_k \gamma_{jk,t}$, is included to account for the incomplete shares term in shift-share instrumental variables models (Borusyak, Hull, and Jaravel 2022).[26] Additional controls include occupational augmentation exposure, $\text{AugX}_{j,t}$, and a set of decade dummies, δ_t. To purge any potential mechanical association between occupational size and the rate of new-title emergence, $X_{j,t}$ controls for occupations' initial (start of period) employment share and, in one specification, for their contemporaneous employment share change. The coefficient β_1^k corresponds to the estimated effect of demand shift $\text{DemandX}_{j,t}^k$ for $k \in \{C, D\}$ on new-title flows, where C denotes China trade exposure and D denotes exposure to demographic demand. As with earlier models for new-title counts, we estimate this equation with a negative binomial regression.

Our conceptual model predicts that $\beta_1^C < 0$ and $\beta_1^D > 0$: inward occupational demand shifts slow new-work creation and outward occupational demand shifts accelerate it. Estimates of

26. Shares do not sum to one because nonmanufacturing industries are not directly exposed to trade shocks; and public sector services are not enumerated in Consumer Expenditure Survey data and hence do not contribute to the demographic demand-shock measure.

TABLE V

OCCUPATIONAL NEW-TITLE EMERGENCE AND DEMAND CONTRACTIONS AND EXPANSIONS

	(1)	(2)	(3)	(4)	(5)
Panel A: Negative Trade Shock 1990–2000 and 2000–2018					
Import exposure	−9.83	−15.44**	−12.13*	−17.49***	−17.73***
	(6.09)	(5.23)	(5.53)	(5.13)	(5.17)
Augmentation exposure			7.94+	9.38**	8.32**
			(4.60)	(3.00)	(2.91)
Panel B: Negative Trade Shock, Placebo Test 1970–1980 and 1980–1990					
Import exposure	14.50	3.95	11.77	−2.99	−1.76
	(26.06)	(20.40)	(20.47)	(13.24)	(12.53)
Augmentation exposure			19.57***	20.00***	20.60***
			(3.15)	(1.77)	(1.92)
Panel C: Positive Demographic Shock 1980–2000 and 2000–2018					
Demand-shift exposure	18.39***	12.83+	18.53***	14.86*	14.75*
	(5.49)	(6.66)	(5.35)	(6.71)	(6.60)
Augmentation exposure			5.36	11.17***	10.56***
			(4.97)	(2.43)	(2.37)
Time FE	X	X	X	X	X
Occ. emp. shares	X	X	X	X	X
Ind. exposure control	X	X	X	X	X
Broad occ. FE		X		X	X
Δ occ. emp. shares					X

Notes. Dependent variable: occupational new-title count. Negative binomial models, coefficients multiplied by 100. $N = 610$ in Panels A and C; $N = 588$ in Panel B. Panel A estimates the relationship between new titles emerging 1990–2018 and occupational exposure to import competition in the same period. Panel B (a placebo test) estimates the relationship between new titles emerging over 1970–1990 and occupational exposure to import competition over the 1990-2018 period. Panel C estimates the relationship between new titles emerging over 1980–2000 and 2000–2018 and occupational exposure to demographically driven demand shocks. Standard errors clustered by occupation are in parentheses. Observations are weighted by start-of-period occupational employment shares. Augmentation and automation exposure measures correspond to the log of the weighted counts of matched patents. $^{+}p < .10$, $^{*}p < .05$, $^{**}p < .01$, $^{***}p < .001$.

equation (2) in Table V support these predictions. Panel A reports the effect of inward occupational demand shocks (stemming from trade shocks) on the emergence of new occupational titles between 1990–2000 and 2000–2018. The first specification controls for decade main effects, occupational employment shares, and the sum of exposure shares. The column (1) estimate finds an imprecise negative effect of demand contractions on new-title emergence ($\hat{\beta}_1^C = -9.83\,(6.09)$). Column (2) adds dummies for broad occupational groups, so identification stems from contrasts across detailed occupations within the 12 broad occupational categories depicted in Online Appendix Figure A.II—from variation within rather than between rows of the figure. The estimated effect of import exposure increases in absolute magnitude and precision to $\hat{\beta}_1^C = -15.44\,(5.23)$.

Column (3) adds the occupation-level augmentation exposure measure from above to the column (1) model, thus allowing both demand shocks and augmentation innovations to affect new-title emergence. The augmentation variable obtains a positive coefficient of $\hat{\beta}_2 = 7.94\,(4.60)$. Column (4) again includes broad occupation fixed effects: this further increases the magnitude and precision of the import exposure variable ($\hat{\beta}_1^C = -17.49\,(5.13)$) and the augmentation exposure term ($\hat{\beta}_2 = 9.38\,(3.00)$). Column (5) additionally controls for the contemporaneous change in occupational employment shares—which is quite conservative since employment shares are an intermediate outcome that is directly affected by the China trade shock. Nevertheless, the coefficients of interest are hardly affected.

Might these estimates reflect long-standing trends in new-work creation in trade-exposed occupations that predate the China trade shock? Table V, Panel B confronts this concern with a placebo test. The dependent variable in this panel is the flow of new occupational titles from the 20 years prior to the China trade shock, 1970–1980 and 1980–1990, while the independent variable is the post-1990 change in import exposure, as per Panel A. These placebo estimates reveal that recent Chinese import competition has little to no predictive power for the emergence of new titles from 1970 through 1990 in subsequently trade-exposed occupations. Estimates are imprecisely positive in the first three columns and are negative and close to zero in the final two columns, which contain the full complement of controls. These results increase our confidence that the estimates in Panel A reflect the causal impact of demand contractions on the flow of new work.

Table V, Panel C presents a parallel exercise exploiting demographic demand shifts in place of trade shocks, here extending the sample window backward to 1980. Consistent with expectations, positive demand shifts speed the emergence of new occupational titles. In column (1), the point estimate of $\hat{\beta}_1^D = 18.39\,(5.49)$ on the demographic demand index implies that a 10% increase in occupational demand-shock exposure increases the rate of occupational new-title emergence by approximately 1.8%. Subsequent columns of Panel C explore robustness by applying the same specifications used for the trade-based demand shock. Across all

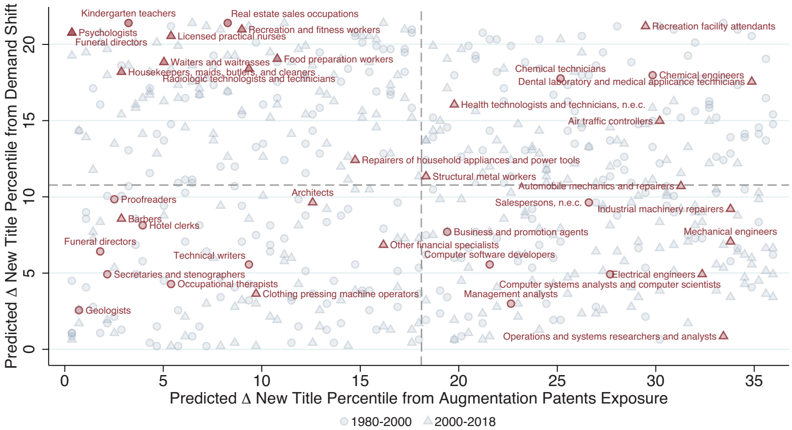

FIGURE VIII

Predicted Δ Occupational New-Title Share Percentile from Exposure to Augmentation versus Exposure to Demand Shifts, 1980–2000 and 2000–2018

This figure plots partial predicted new-title flows from exposure to demand shifts against partial predicted new-title flows due to augmentation exposure (both in percentile terms) for consistent census occupations. Predictions are based on Online Appendix Table A.X, column (4). Dashed lines indicate median predicted changes.

columns, occupations with higher predicted demand growth stemming from demographic change exhibit faster new-title emergence. As in Panels A and B, the coefficient on augmentation exposure is positive and precisely estimated. Moreover, its inclusion has a relatively small effect on the coefficient on the demand shift variable. Figure VIII shows why this is the case: the set of occupations most exposed to demographic demand shifts is substantively different from those most exposed to augmentation innovations. Using the partial predicted effects from a version of Table V, we find that augmentation exposure and demand exposure are essentially uncorrelated over 1980–2018. Figure VIII indicates that demographic demand shifts can help account for the emergence of new titles in personal service and care jobs such as housekeepers, waiters, food preparation workers, and licensed practical nurses (see also Mazzolari and Ragusa 2013; Leonardi 2015; Comin, Danieli, and Mestieri 2020). Conversely, new-title emergence in high-tech jobs, such as electrical engineers and computer systems analysts, is primarily predicted by augmentation exposure.

The evidence in Table V highlights that new work has multiple origins: the flow of augmentation innovations and the operation of demand shifts both shape where new specialized, labor-using tasks emerge. These results also offer a further substantive implication. Much evidence documents that rising import competition has depressed employment in trade-exposed industries and associated occupations during the past three decades (see evidence summarized in Autor, Dorn, and Hanson 2021). The Table V results reveal that this competitive pressure not only numerically reduces employment but also depresses the emergence of new categories of specialized work: it yields fewer jobs and less new work. This distinction is relevant because, as we show in Online Appendix Table A.XV, new work appears to be better remunerated than existing work in the same occupation, possibly because it is more specialized.

V. Augmentation, Automation, and Occupational Labor Demand

What does our evidence on the determinants of new work imply for occupational employment and labor demand more generally? The findings so far do not unambiguously answer this question since employment could potentially contract in occupations where new titles emerge or expand in occupations where tasks are automated. Our conceptual model predicts that because new-task creation generates new demand for human expertise, it will expand employment and wage bills in augmentation-exposed occupations; and conversely because task automation is labor displacing, it will erode employment and wage bills in automation-exposed occupations (see Proposition 2 in Online Appendix J). We explore those implications here.

Figure IX motivates this analysis by documenting the striking association between new-title emergence and occupational employment growth across both halves of our sample, netting out the correlation between start-of-period title count and occupational employment growth (which is generally negative since larger occupations grow more slowly). The locus of new-title emergence and employment growth has shifted substantially across the four-decade intervals of our sample, yet new-title flows strongly predict the set of occupations that gain and lose ground in both periods. Between 1940 and 1980, occupations rapidly adding both employment and new titles include clerical

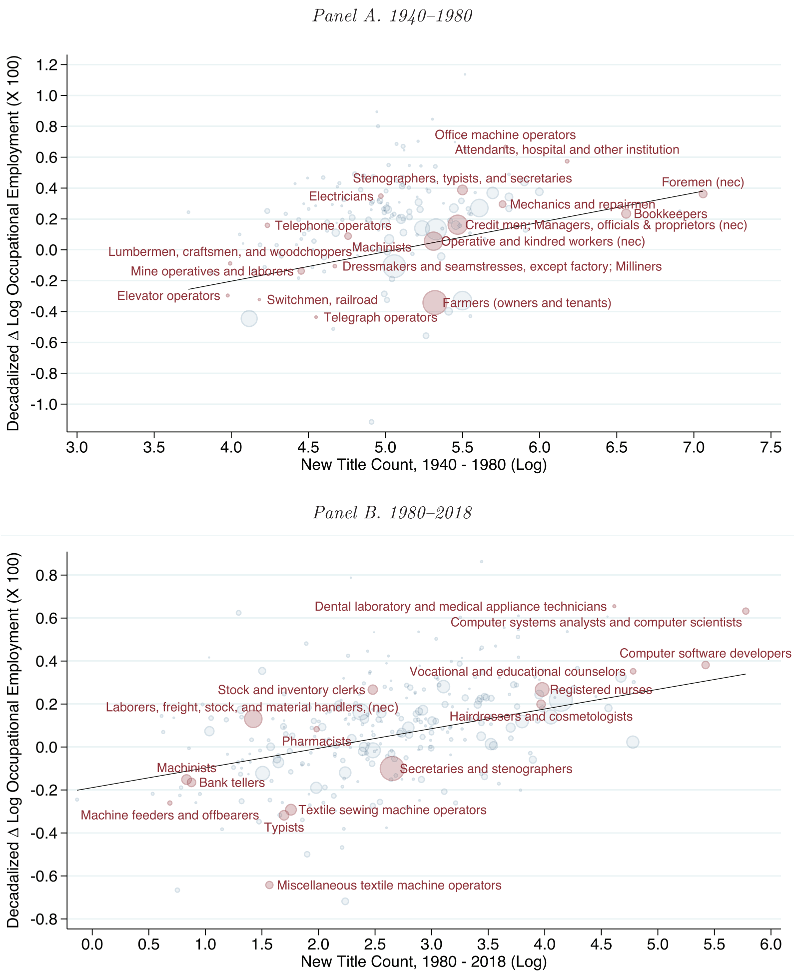

FIGURE IX

Correlations between Employment Growth and New-Title Emergence Rates, 1940–1980 and 1980–2018

This figure illustrates the relationship between log new-title counts and decadalized occupational employment growth, controlling for initial-year log title counts, and using initial-year employment shares as weights. Plotted lines correspond to weighted partial fitted values. The weighted partial correlation between log employment growth and the log of new-title flows is 0.333 for 1940–1980 and 0.482 for 1980–2018. Axis labels are exponentiated.

occupations such as stenographers, typists, secretaries, office machine operators, and bookkeepers, and blue-collar production occupations such as foremen, mechanics and repairmen, and operative workers. Simultaneously, occupations including elevator operators, lumbermen, railroad switchmen, dressmakers, and mining occupations grew slowly and added few titles. In the subsequent four decades, 1980–2018, professional and personal services dominate the occupations rapidly adding employment and new titles: computer software developers, dental laboratory and medical appliance technicians, vocational and educational counselors, registered nurses, and hairdressers and cosmetologists. Conversely, blue-collar production occupations, such as machinists, machine feeders and off-bearers, and white-collar clerical occupations, including bank tellers, typists, secretaries, and stenographers, added relatively few occupational titles and exhibited substantial relative employment declines.

These patterns echo the evidence from Figure II, showing that the locus of new-work creation shifted from the middle-paid center of the occupational distribution before 1980, to the high-paid and low-paid tails of this distribution after 1980. Online Appendix Figure A.III facilitates a more direct comparison of new-title growth rates across periods by applying a consistent set of occupation codes over 1940–2018, albeit at the cost of substantially lower occupational resolution. What stands out from this figure is the shifting fortunes of routine task-intensive occupations—blue-collar occupations such as operative and kindred workers, metal workers, and mechanics, and white-collar occupations such as shipping and receiving clerks, stenographers, typists, and secretaries, and library attendants and assistants. These occupations were well above the median of new-title growth between 1940 and 1980, and then fell substantially in rank (i.e., below the 45-degree line) during the next four decades, coinciding with the decline of routine task-intensive employment (Autor and Dorn 2013).

These correlations between employment growth and new-title flows should not be interpreted as causal. Indeed, the prior results in Table V underscore that the co-occurrence of new occupational tasks and rising occupational employment could reflect the operation of demand shifts. To understand the effects of augmentation and automation on labor demand, we shift focus from new titles to employment and wage bills. Using OLS and 2SLS models, we analyze the correlational and causal relationships between augmentation and automation innovations and changes in occupational employment and wage bills.

V.A. *Employment and Wage Bills: Main Results*

To assess the relationship between augmentation exposure, automation exposure, employment, and wage bills, we estimate models of the form:

$$(3) \quad 100 \times \Delta \ln(Y_{ij,t}) = \beta_1 \text{AugX}_{ij,t} + \beta_2 \text{AutX}_{j,t} + \gamma_{i,t} + \delta_{J,t} + \varepsilon_{ij,t},$$

where the dependent variable is the log change in full-time equivalent employment (annual hours divided by 35 hours × 50 weeks) or wage bill in consistent three-digit industry i by occupation j cells, multiplied by 100 so changes roughly correspond to percentage points. While the new-title models were limited to occupation-by-decade variation in the outcome variable, this estimation is able to exploit richer occupation-by-industry-by-decade variation in employment and wage bills. The key independent variables are $\text{AugX}_{ij,t}$, quantifying exposure to augmentation in industry-by-occupation cells, and $\text{AutX}_{j,t}$, quantifying exposure to automation in occupation cells. The augmentation measure is extended to vary at the industry-occupation level by aggregating patents that exhibit jointly high textual similarity to occupation titles and industry titles (see Online Appendix E). Because the analysis here focuses on employment changes over four-decade intervals (1940–1980 and 1980–2018), we take (the log of) the sum of citation-weighted patent matches for $\text{AugX}_{ij,t}$ and $\text{AutX}_{j,t}$ in each interval. The inclusion of a full set of industry-by-decade dummy variables, $\gamma_{i,t}$ (116 industries for 1940–1980, 206 industries for 1980–2018), means that the coefficients of interest are identified by changes in within-industry occupational employment, holding constant overall industry employment shifts. Additional specifications include broad occupation fixed effects further interacted with time period dummies, $\delta_{J,t}$. Standard errors are clustered on industry-by-occupation cells.

Table VI presents OLS estimates in columns (1) and (2) and 2SLS estimates in columns (3) and (4). The first panel of Table VI fits a stacked long-difference version of equation (3) over eight decades where each observation is the log change in employment in consistent industry-occupation cells in a four-decade interval (1940–1980 or 1980–2018). Column (1) finds that occupations that are more exposed to augmentation exhibit faster

TABLE VI

THE RELATIONSHIP BETWEEN CHANGES IN EMPLOYMENT AND WAGE BILLS AND EXPOSURE TO AUGMENTATION AND AUTOMATION WITHIN INDUSTRY-OCCUPATION CELLS, OLS AND 2SLS STACKED LONG-DIFFERENCE REGRESSIONS, 1940–2018

	OLS		2SLS	
	(1)	(2)	(3)	(4)
Panel A: Decadalized Δln(employment)				
Augmentation exposure	1.51***	1.36***	4.37***	3.63***
	(0.21)	(0.22)	(0.93)	(0.96)
Automation exposure	−2.27***	−1.00*	−4.02***	−4.21***
	(0.27)	(0.40)	(0.62)	(0.93)
R^2	0.53	0.57		
Panel B: Decadalized Δln(wage bill)				
Augmentation exposure	1.50***	1.41***	4.95***	4.04***
	(0.23)	(0.25)	(1.02)	(0.98)
Automation exposure	−2.21***	−0.97*	−3.17***	−3.43***
	(0.28)	(0.42)	(0.67)	(0.97)
R^2	0.53	0.57		
Panel C: Decadalized ΔE[ln(wage bill)]				
Augmentation exposure	1.39***	1.30***	4.15***	3.31***
	(0.20)	(0.21)	(0.91)	(0.93)
Automation exposure	−2.13***	−1.01**	−3.40***	−3.74***
	(0.27)	(0.38)	(0.63)	(0.91)
R^2	0.55	0.58		
Panel D: Decadalized Δln(adjusted wage bill)				
Augmentation exposure	1.62***	1.47***	5.18***	4.36***
	(0.24)	(0.25)	(1.03)	(1.00)
Automation exposure	−2.34***	−0.96*	−3.79***	−3.90***
	(0.28)	(0.44)	(0.66)	(0.99)
R^2	0.52	0.56		
F-stat. (Aug)			127.76	150.56
F-stat. (Aut)			202.85	145.10
Ind. × time FE	X	X	X	X
Broad occ. × time FE		X		X

Notes. $N = 33{,}900$ changes in employment and wage bill in consistently defined census occupations over 1940–1980 and 1980–2018. Dependent variables are decadalized and multiplied by 100 so that growth rates are expressed in per decade percentage points. All employment and wage bill changes are winsorized at the 99th percentile. Wage bill variables sum the product of hours and wages for all workers. Hourly wages are imputed to cells where employment but not wages are observed. 2SLS specifications additionally control for indicators for being above the cross-sectional median in total augmenting or automating patent matches two decades prior. Standard errors are clustered by industry-occupation cell (using Stata commands `reghdfe` and `ivreghdfe`) in parentheses. Observations are weighted by start-of-period employment share for each industry-occupation cell. Long-differences are four-decade changes, 1940–1980 and 1980–2018. Augmentation and automation exposure measures correspond to the log of the weighted counts of matched patents. $^{+}p < .10$, $^{*}p < .05$, $^{**}p < .01$, $^{***}p < .001$.

employment growth and conversely, those more exposed to automation exhibit slower employment growth. Indeed, augmentation and automation have precisely estimated but opposite-signed predictive relationships with occupational employment. In column (1), 10% more augmentation exposure predicts 0.15% more employment growth ($\hat{\beta}_1 = 1.51\,(0.21)$), and 10% more automation exposure predicts 0.23% less employment growth ($\hat{\beta}_2 = -2.27\,(0.27)$). Column (2) adds controls for broad occupation by time-effects, demonstrating that these predictive relationships stem from contrasts within and between broad occupational categories.

To address endogeneity concerns discussed earlier, the second pair of columns in Table VI present 2SLS estimates that use breakthrough patent flows from the prior two decades as instruments for observed augmentation and automation patent flows.[27] First-stage estimates for these models are reported in Online Appendix Table A.III. The 2SLS results confirm the implications of the OLS estimates but with greater magnitude: a 10% increase in augmentation exposure spurs 0.36% additional employment growth ($\hat{\beta} = 3.63$ (0.96)) over the course of a decade in the 2SLS model relative to 0.14% in the corresponding OLS column; a 10% increase in automation exposure spurs a 0.42% reduction in employment growth ($\hat{\beta} = -4.21$ (0.93)) over the course of a decade in the 2SLS model relative to 0.10% in the OLS estimate. The larger point estimates for the automation and augmentation employment elasticities in these 2SLS estimates suggest that attenuation due to measurement error may impart a larger bias to the OLS estimates than does simultaneity. We make the caveat

27. As with earlier 2SLS models, we normalize the breakthrough exposure instruments with sector-specific (manufacturing versus nonmanufacturing) median breakthrough thresholds. Let $N^{A,10}_{ji,t-20}$ equal breakthrough patent exposure count for occupation j in industry i in decade $t-20$, where $A \in \{\text{Aug}, \text{Aut}\}$. Like the endogenous variable, the breakthrough matches $N^{A,10}_{ji,t-20}$ are also aggregated over a four-decade period, except the aggregation window is lagged by 20 years as with our previous 2SLS analysis. Occupation by industry cells are classified as breakthrough-exposed according to $B^{A,10}_{ji,t-20} = \mathbb{1}\left[N^{A,10}_{ji,t-20} > \tilde{N}^{A,10}_{j,i\in\text{manuf},t-20}\right] \times \left[i \in \text{manuf}\right] + \mathbb{1}\left[N^{A,10}_{j,i,t-20} > \tilde{N}^{A,10}_{j,i\notin\text{manuf},t-20}\right] \times \left[i \notin \text{manuf}\right]$. Here, the tilde denotes the employment-weighted cross-sectional median of the variable defined over industry-occupation cells, subsetting according to whether industry i is or is not in the manufacturing sector. We analogously construct $B^{A,100}_{ji,t-20}$, denoting above-median overall (breakthrough + non-breakthrough) patent exposure.

that the 2SLS estimates have broader confidence intervals than their OLS counterpart, so that OLS point estimates generally fall within the 95% confidence bands of their 2SLS counterparts.

The next set of results in Table VI focus on the wage bill, equal to the product of employment and earnings (Panel B). These wage bill estimates prove highly comparable to the corresponding employment results in Panel A: the OLS coefficient for the effect of augmentation innovations on the wage bill in column (2) of Panel B is 1.41 (0.25) versus 1.36 (0.22) on employment in column (2) of Panel A. For the 2SLS estimates in column (4), the point estimates are 4.04 (0.98) for the wage bill (Panel B) versus 3.63 (0.96) for employment (Panel A). The corresponding OLS estimates for automation innovations are −0.97 (0.42) for the wage bill and −1.00 (0.40) for employment, while their 2SLS counterparts are −3.43 (0.97) for the wage bill and −4.21 (0.93) for employment. These patterns are consistent with the interpretation that employment and wage bills move one for one, suggesting a unitary elasticity of substitution across occupations (also assumed in our model, where occupations correspond to sectors). This literal interpretation should be qualified, however. Papers by Böhm, von Gaudecker, and Schran (2022) and Autor and Dorn (2009) demonstrate that contracting occupations (here, those more exposed to automation) tend to retain both more experienced workers and workers with relatively high earnings given their experience, while the opposite occurs in expanding occupations (here, those more exposed to augmentation). These compositional shifts, which are akin to quantity rather than price changes in an earnings equation, cloud inference on the effect of augmentation and automation on wage bills net of these compositional shifts.

To (imperfectly) address this compositional issue, we estimate the effect of augmentation and automation innovations on composition-adjusted wages, as detailed in Online Appendix G. Briefly, we use cross-sectional Mincerian wage regressions in each census year to predict the log hourly wage of each worker based exclusively on her schooling, race, and gender, each interacted with a quadratic in age. We take the cell-level average of these predictions to form the expected wage in each industry-occupation cell ($\widehat{W}_{ij,t}$), purged of industry-occupation premia. This procedure delivers two new wage bill change measures (alongside the observed wage bill change, $\Delta W_{ij,t}$): the expected wage bill change ($\Delta \widehat{W}_{ij,t}$); and the composition-adjusted wage bill change

($\Delta \widetilde{W}_{ij,t}$), equal to the observed log change in employment plus the log difference between the observed and composition-adjusted wage change.

Table VI, Panel C uses as the dependent variable the expected wage, $\Delta \widehat{W}_{ij,t}$. Consistent with Böhm, von Gaudecker, and Schran (2022), expected wage bill changes are less positive for augmentation exposure and more negative for automation exposure than are observed wage bill changes—particularly in 2SLS models. This pattern suggests that adverse compositional shifts in augmentation-exposed occupations partly mask any positive wage relationship, whereas positive compositional shifts in automation-exposed occupations mask negative wage effects. Panel D combines these results to estimate models for composition-adjusted wage bills. Accounting for compositional changes within industry-occupation cells, estimates find that augmentation innovations have a meaningfully larger positive effect on wage bills than on employment. Focusing on 2SLS estimates, Panel A, column (4) finds that a 10% rise in augmentation exposure spurs a 0.36% rise in employment ($\hat{\beta}_1 = 3.63\,(0.96)$), and in Panel D, column (4), a 0.44% rise in the adjusted wage bill ($\hat{\beta}_1 = 4.36\,(1.00)$). By implication, the wage premium in the exposed industry-occupation cell rises by 0.08% ($0.44-0.36 = 0.08$). This difference between employment and wage bill effects, which is pronounced in 2SLS models but nevertheless apparent in OLS models, hints that new work may be more skilled or better remunerated than (simply) more work. We return to this point below.[28]

As a further extension, we explore sectoral impacts. Much of the economic literature on the labor-demand impact of technological innovations focuses on manufacturing, where capital investments are well measured and technologies are often intensively used, for example, in the case of industrial robotics (e.g., Berman, Bound, and Griliches 1994; Doms, Dunne, and Troske 1997; Lewis 2011; Acemoglu and Restrepo 2020; Curtis et al. 2021; Dechezleprêtre et al. 2021; Humlum 2021; Aghion et al. 2022; Hirvonen, Stenhammar, and Tuhkuri 2022). The augmentation and automation exposure measures developed here are not confined to manufacturing since they are based on the semantic content of occupations and innovations. Leveraging this fact, Online Appendix Table A.V reports a corresponding set

28. We do not find that automation has larger negative effects on occupational wage bills than occupational employment in either OLS or 2SLS models.

of OLS and 2SLS employment and adjusted-wage-bill models for the 1940–2018 period, estimated separately for manufacturing and nonmanufacturing. Within both the manufacturing and nonmanufacturing sectors, along employment and wage bill margins, and using both OLS and 2SLS models, we find that augmentation innovations significantly raise occupational labor demand, whereas automation innovations significantly depress it.

V.B. Evidence of Accelerating Automation

We observed that the elasticity of new-work creation, as proxied by new job titles and with respect to augmentation innovation, appears to have declined between 1940–1980 and 1980–2018 (Tables II and IV). We perform a parallel exercise for occupational labor demand in Table VII, which reports OLS and 2SLS models for the effects of augmentation and automation innovations on employment and composition-adjusted wage bills separately for 1940–1980 and 1980–2018. We note that there is no intrinsic statistical relationship between these two margins of response—new-title creation versus occupational labor demand—so we view this test of parallelism as informative. The by-period estimates in Table VII yield three results. First, OLS and 2SLS estimates strongly corroborate the positive impact of augmentation innovations on employment and wage bills in the four-decade intervals under study. Given that the measures from these intervals are built from different patent cohorts, different occupational labor task descriptors (1939 DOT versus 1977 DOT), and different occupational output descriptors (1950–1980 CAI versus 1990–2018 CAI), the stability and robustness of the estimates across epochs buttresses confidence in the findings.

Second, akin to the evidence for new-title flows, both OLS and 2SLS models suggest a fall in the elasticity of employment creation with respect to augmentation innovations (and a more negative elasticity with respect to automation) in the second relative to first half of the sample. Referring to the even columns of Table VII (which include broad occupation by year dummies), the OLS elasticity estimate for the effect of augmentation on employment falls from $\hat{\beta}_1^{40-80} = 2.12\,(0.38)$ to $\hat{\beta}_1^{80-18} = 0.78\,(0.23)$. The corresponding 2SLS elasticity estimate falls from $\hat{\beta}_1^{40-80} = 5.24\,(2.42)$ to $\hat{\beta}_1^{80-18} = 2.79\,(0.93)$. OLS and 2SLS estimates for the augmentation elasticity of the wage bill also find qualitatively similar declines in magnitudes.

TABLE VII

RELATIONSHIP BETWEEN CHANGES IN EMPLOYMENT AND WAGE BILLS AND EXPOSURE TO AUGMENTATION AND AUTOMATION WITHIN INDUSTRY-OCCUPATION CELLS, OLS AND 2SLS STACKED LONG-DIFFERENCE REGRESSIONS, 1940–1980 AND 1980–2018

	1940–1980				1980–2018			
	OLS		2SLS		OLS		2SLS	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Decadalized Δ*ln(employment)*								
Augmentation	1.85***	2.12***	5.78**	5.24*	1.27***	0.78***	3.35***	2.79**
exposure	(0.39)	(0.38)	(2.21)	(2.42)	(0.21)	(0.23)	(0.91)	(0.93)
Automation	−1.47***	−0.27	−3.97*	−2.79	−3.86***	−1.98***	−5.05***	−5.07***
exposure	(0.40)	(0.64)	(1.77)	(2.95)	(0.34)	(0.46)	(0.66)	(0.92)
R^2	0.63	0.66			0.44	0.48		
Panel B: Decadalized Δ*ln(adjusted wage bill)*								
Augmentation	2.04***	2.36***	6.93**	6.40*	1.30***	0.78***	3.86***	3.24***
exposure	(0.45)	(0.45)	(2.33)	(2.54)	(0.23)	(0.23)	(0.96)	(0.96)
Automation	−1.63***	−0.39	-4.04*	−2.64	−3.81***	−1.74***	−4.80***	−4.86***
exposure	(0.41)	(0.69)	(1.76)	(3.05)	(0.39)	(0.50)	(0.67)	(0.95)
R^2	0.61	0.64			0.44	0.49		
F-stats (aug)			34.80	46.34			160.12	150.02
F-stat. (aut)			38.68	15.16			169.77	301.79
Ind. × time FE	X	X	X	X	X	X	X	X
Broad occ. × time FE		X		X		X		X

Notes. $N = 6{,}520$ changes in employment and wage bills in consistently defined census occupations over 1940–1980, and $N = 27{,}380$ changes over 1980–2018. Dependent variable is decadalized and multiplied by 100 so that growth rates are expressed in per-decade percentage points. All employment and wage bills changes are winsorized at the 99th percentile. Hourly wages are imputed to cells where employment but not wages are observed. 2SLS specifications additionally control for indicators for being above the cross-sectional median in total augmenting or automating patent matches two decades prior. Standard errors are clustered by industry-occupation cell (using Stata commands `reghdfe` and `ivreghdfe`) in parentheses. Observations are weighted by start-of-period employment share for each industry-occupation cell. Long-differences are four-decade changes, 1940–1980 and 1980–2018. Augmentation and automation exposure measures correspond to the log of the weighted counts of matched patents. $^{+}p < .10$, $^{*}p < .05$, $^{**}p < .01$, $^{***}p < .001$.

Moreover, coefficient estimates in odd columns, which instead rely on between-broad-occupation variation, are generally larger and much more precise for our automation measure (and for augmentation in the case of IV estimates), but they also deliver the same pattern of decreased augmentation and increased automation intensity when comparing coefficients across time periods.

Finally, we find a rise in the elasticity of employment and wage bill erosion to automation during 1980–2018 relative to 1940–1980. Focusing on employment, the OLS estimate of this elasticity rises in absolute magnitude from $\hat{\beta}_2^{40-80} = -0.27$ (0.64) to $\hat{\beta}_2^{80-18} = -1.98$ (0.46). The corresponding 2SLS estimate grows from $\hat{\beta}_2^{40-80} = -2.79$ (2.95) to $\hat{\beta}_2^{80-18} = -5.07$ (0.92). This pattern is inescapable in Figure X, which plots the conditional relationship between augmentation, automation, and employment growth separately for these four-decade intervals, using the somewhat more precise OLS estimates from columns (1) and (5), which leverage between- as well as within-broad-occupation variation.

In summary, Tables VI and VII, and the supporting material in Online Appendix Tables A.V and A.XV, corroborate a central implication of our conceptual framework: despite their positive rate of co-occurrence across occupations, augmentation and automation have countervailing causal effects on sectoral labor demand. This is true for employment and wage bills, for the 1940–1980 and 1980–2018 periods, for manufacturing and nonmanufacturing, and for OLS and 2SLS estimates, the latter of which leverage breakthroughs occurring two decades earlier as instruments for subsequent downstream innovations.

Our findings are consistent with Kogan et al. (2021) and Webb (2020), who find that occupations that are more exposed to automation (Kogan et al. 2021) or software patents (Webb 2020) have seen falling relative employment in recent decades.[29] Our analysis adds two central findings to this body of work: isolating a countervailing augmentation force, also stemming from innovation, that is strongly predictive of occupational employment growth; and providing causal evidence that such innovations boost labor demand while automation innovations erode it. In the terminology of Acemoglu and Restrepo (2019a), we find that automation is labor displacing and augmentation is labor reinstating. Also consistent with their evidence, the constellation of results for new-task creation and occupational labor demand suggests that the labor-reinstating power of augmentation has attenuated in recent decades while the labor-replacing force of automation has accelerated.

V.C. Is New Work Different from More Work?

Our primary public use census and ACS data files, augmented with new-title counts, document the prevalence of new work in macro-occupations but do not allow us to observe which

29. Mann and Püttmann (2023) present a contrasting finding: service sector employment grows relatively faster in commuting zones exposed to more automation patents. Given the differing unit of analysis—geographies versus occupations—these findings are not directly comparable.

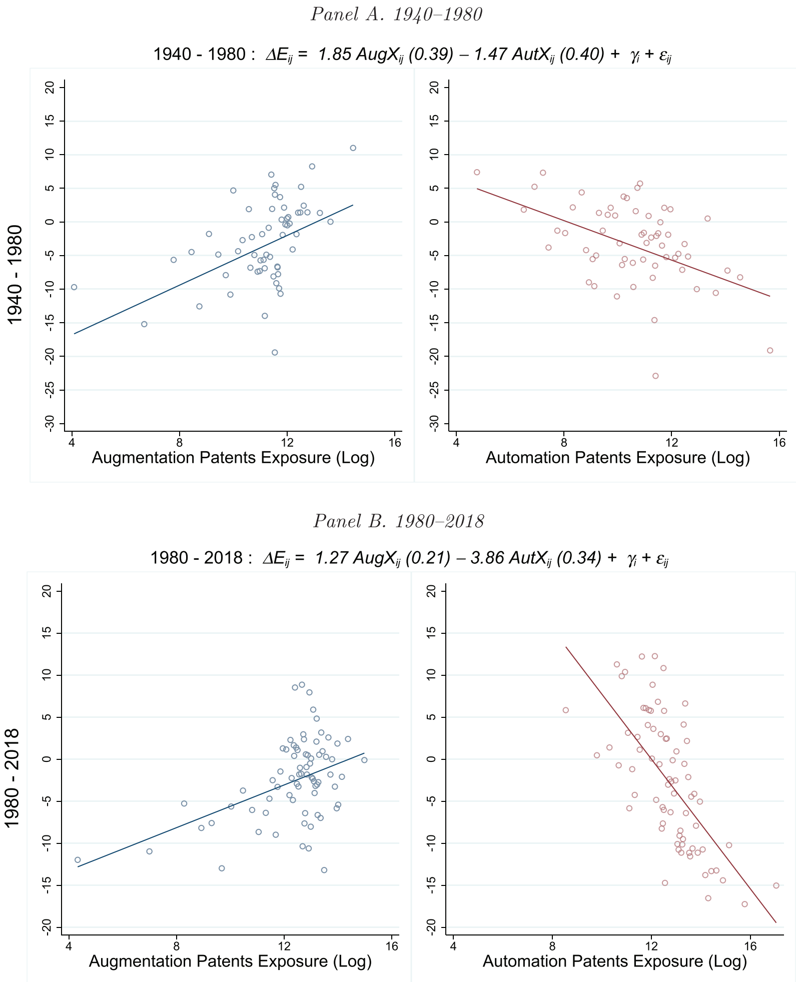

FIGURE X

Conditional Correlations between Automation, Augmentation, and Employment Growth (based on OLS Regressions)

The figure reports binscatters of the employment-weighted conditional correlation between decadalized percent employment growth and exposure to augmentation innovations (left side) and automation innovations (right side). Panel A corresponds to the regression specification in Table VII, Panel A, column (1) using consistently defined census industry×occupation cells over 1940–1980 ($N = 6{,}520$). Panel B corresponds to the regression specification in Table VII, Panel A, column (5) using consistently defined census industry×occupation cells over 1980–2018 ($N = 27{,}380$).

workers are directly employed in that work. These public-use data sources are accordingly ill-suited to estimating wage differentials in new work. Online Appendix H seeks to fill this gap by harnessing data from the 1940 census Complete Count file (Ruggles et al. 2021) to test whether new work is more skilled or better paid than conventional work. The unique virtue of these no-longer-confidential census Complete Count data is that they report all fields for each census respondent, enabling direct observation of which workers are employed in new work, that is, employed in occupational titles that first appeared in the 1940 CAI.

Online Appendix Table A.XV, Panel A explores whether there is positive selection on observable characteristics of workers into new work. Using a classification procedure detailed in Online Appendix H1, we estimate that 2.6% of workers in 1940 were employed in new work titles that were first reported in that decade. Fitting a linear probability model to these data reveals that better-educated workers are significantly more likely to be employed in new work. This is true unconditionally, and also when accounting for a full set of macro-occupation main effects—thus contrasting workers in the same census occupations—and when further accounting for workers' age, sex, race, state, and urban versus rural status. Net of these factors, workers with any secondary education are at least 10% more likely to be employed in new work than those with less than a ninth-grade education (the modal educational category in 1940).

Online Appendix Table A.XV, Panel B explores whether there is a wage differential in new work net of selection on observables. Fitting an OLS cross-sectional wage regression to the 1940 data shows that even after conditioning on a rich set of observables, including macro-occupation main effects, education, race, gender, age, state, and urban/rural status, workers employed in new work earn on average 6.0% more than observably similar workers in preexisting work. Although this evidence does not establish that employment in new work raises wages, it indicates that either workers earn a premium for new work or that new work attracts workers with unobservably higher skill (even within the same macro-occupation and level of education). Either interpretation is consistent with the hypothesis that new work demands specialized skills or expertise relative to conventional work within the same occupations.

VI. Conclusion

Much recent empirical work has documented the displacement of labor from existing job tasks by automation but is mostly silent on the countervailing force of labor reinstatement occurring through the creation of new job tasks or categories requiring specialized human expertise. Using newly constructed measures of the emergence of new work within occupations, as well as flows of innovations that may potentially augment or automate these occupations, we find that the locus of new-work creation has shifted from middle-paid production and clerical occupations in the first four decades after World War II, to high-paid professional and secondarily low-paid services since 1980. We show that new work emerges in response to technological innovations that complement the outputs of occupations and demand shocks that raise occupational demand. Conversely, innovations that automate existing job tasks do not yield new work, while adverse occupational demand shifts slow the rate of new-work emergence. These flows of augmentation and automation innovations are positively correlated across occupations but have countervailing effects on labor demand: augmentation innovations boost occupational labor demand while automation innovations erode it. These effects are present in both four-decade epochs of our sample and are evident in the manufacturing and nonmanufacturing sectors. By harnessing shocks to the flow of augmentation and automation innovations spurred by breakthrough innovations occurring two decades prior, we establish that the effects of augmentation and automation innovations on new-work emergence and occupational labor demand are causal.

The data, methods, and findings in this article provide the foundation for further study of the role of innovation and incentives in automating existing work and catalyzing demands for novel human expertise and specialization. Some core questions that future research may seek to address include:

i. Is automation accelerating relative to augmentation, as many researchers and policy makers fear, and have the loci of automation and augmentation shifted across occupations and skill groups? On the latter question, our evidence is clear: automation has intensified in middle-skill occupations while augmentation has concentrated in professional, technical, and managerial occupations and to a lesser extent in personal service occupations. On the former question the results suggest that in recent decades the demand-eroding effects of automation innovations have intensified, whereas the demand-increasing effects of augmentation innovations have slowed. While the augmentation and automation exposure measures developed here are too novel to warrant definitive conclusions about acceleration or deceleration over the course of many decades, the pattern of results across the margins of both employment and new-work emergence suggest, as in Acemoglu and Restrepo (2019a), that the past four decades have seen relatively more automation and less augmentation than the prior four.

ii. Is new work more labor-augmenting than 'more work'? The finding that augmentation innovations increase occupational wage bills by boosting both employment and wages suggests that new work may be more valuable than more work—plausibly because new work demands novel expertise and specialization that (initially) commands a scarcity premium. Identifying and quantifying such premia will require direct observations of job tasks and earnings of workers engaged in new work, something that is largely infeasible in our census public-use microdata. If, as we suspect, new work provides additional opportunities for skill formation and earnings growth beyond more work, then policies that foster new-work creation may be of particular interest.

iii. While our evidence establishes that augmentation and automation move labor demand in countervailing directions, it does not answer why the locus of innovation has shifted over time nor—more broadly—how elastic the locus and pace of augmentation and automation are to incentives. Theory has much to say on this topic (Acemoglu 2010; Acemoglu and Restrepo 2018; Ray and Mookherjee 2022), but direct measurement is a challenge. The augmentation and automation innovation measures developed here, and tools for identifying them, may prove valuable to this pursuit.

iv. Will rapid advances in artificial intelligence shift the balance of innovation toward more rapid automation across an expanding set of occupational domains? And to the degree that AI is not exclusively automating, what novel specialties will it catalyze and what skill sets will it complement? Early evidence on this question points to a broad range of potential effects, but definitive findings are elusive so far (Brynjolfsson and Mitchell 2017; Felten, Raj, and Seamans 2019; Babina et al. 2020; Webb 2020; Acemoglu et al. 2022; Autor 2022; Brynjolfsson, Li, and Raymond 2023; Eloundou et al. 2023; Noy and Zhang 2023; Peng et al. 2023). AI innovations are however well represented in patents (Webb 2020), suggesting that the classification methodology developed here may offer insight into AI's potential for augmentation and automation.

Massachusetts Institute of Technology and National Bureau of Economic Research, United States
Massachusetts Institute of Technology, United States
Utrecht University, The Netherlands
Northwestern University Kellogg School of Management, United States

Supplementary Material

An Online Appendix for this article can be found at *The Quarterly Journal of Economics* online.

Data Availability

The data underlying this article are available in the Harvard Dataverse, https://doi.org/10.7910/DVN/7RYD2E (Autor et al. 2024).

References

Acemoglu, Daron, "Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality," *The Quarterly Journal of Economics*, 113 (1998), 1055–1089. https://doi.org/10.1162/003355398555838

———, "When Does Labor Scarcity Encourage Innovation?" *Journal of Political Economy*, 118 (2010), 1037–1078. https://doi.org/10.1086/658160

Acemoglu, Daron, and David H. Autor, “Skills, Tasks and Technologies: Implications for Employment and Earnings,” in *Handbook of Labor Economics*, vol. 4 part B, eds. David Card and Orley Ashenfelter, (Amsterdam: North-Holland, 2011), 1043–1171. https://doi.org/10.1016/S0169-7218(11)02410-5

Acemoglu, Daron, David H. Autor, David Dorn, Gordon H. Hanson, and Brendan Price, “Import Competition and the Great US Employment Sag of the 2000s,” *Journal of Labor Economics*, 34 (2016), S141–S198. https://doi.org/10.1086/682384

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo et al.“Artificial Intelligence and Jobs: Evidence from Online Vacancies,” *Journal of Labor Economics*, 40 (2022), S293–S340. https://doi.org/10.1086/718327

Acemoglu, Daron, and Pascual Restrepo, “The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment,” *American Economic Review*, 108 (2018), 1488–1542. https://doi.org/10.1257/aer.20160696

———, “Automation and New Tasks: How Technology Displaces and Reinstates Labor,” *Journal of Economic Perspectives*, 33 (2019a), 3–30. https://doi.org/10.1257/jep.33.2.3

———, “Artificial Intelligence, Automation and Work,” in Ajay K. Agrawal, Joshua Gans, and Avi Goldfarb, editors, *The Economics of Artificial Intelligence*, (Chicago: University of Chicago Press, 2019b), Ch. 8.

———, “Robots and Jobs: Evidence from US Labor Markets,” *Journal of Political Economy*, 128 (2020), 2188–2244. https://doi.org/10.1086/705716

———, “Tasks, Automation, and the Rise in U.S. Wage Inequality,” *Econometrica*, 90 (2022), 1973–2016. https://doi.org/10.3982/ECTA19815

Aghion, Philippe, Celine Antonin, Simon Bunel, and Xavier Jaravel, “Modern Manufacturing Capital, Labor Demand, and Product Market Dynamics: Evidence from France.” LSE Working Paper (2022).

Akerman, Anders, Ingvil Gaarder, and Magne Mogstad, “The Skill Complementarity of Broadband Internet,” *The Quarterly Journal of Economics*, 130 (2015), 1781–1824. https://doi.org/10.1093/qje/qjv028

Atack, Jeremy, Robert A. Margo, and Paul W. Rhode, ““Automation” of Manufacturing in the Late Nineteenth Century: The Hand and Machine Labor Study,” *Journal of Economic Perspectives*, 33 (2019), 51–70. https://doi.org/10.1257/jep.33.2.51

Atalay, Enghin, Phai Phongthiengtham, Sebastian Sotelo, and Daniel Tannenbaum, “The Evolution of Work in the United States,” *American Economic Journal: Applied Economics*, 12 (2020), 1–34. https://doi.org/10.1257/app.20190070

Atalay, Enghin, and Sarada, “Firm Technology Upgrading Through Emerging Work,” Federal Reserve Bank of Philadelphia Working Paper 20–44, 2020. https://doi.org/10.21799/frbp.wp.2020.44

Autor, David H., “Work of the Past, Work of the Future,” *AEA Papers and Proceedings,* 109 (2019), 1–32. https://doi.org/10.1257/pandp.20191110

———“The Labor Market Impacts of Technological Change: From Unbridled Enthusiasm to Qualified Optimism to Vast Uncertainty,” *Brookings Institution: Global Forum on Democracy and Technology*, (2022).

Autor, David H., Caroline Chin, Anna M. Salomons, and Bryan Seegmiller, “New Frontiers: The Origins and Content of New Work, 1940–2018.” NBER Working Paper no. 30389, 2022. https://doi.org/10.3386/w30389

———, “Replication Data for: ‘New Frontiers: The Origins and Content of New Work, 1940–2018’,” (2024), Harvard Dataverse. https://doi.org/10.7910/DVN/7RYD2E

Autor, David H., and David Dorn, “This Job Is ‘Getting Old’: Measuring Changes in Job Opportunities Using Occupational Age Structure,” *American Economic Review*, 99 (2009), 45–51. https://doi.org/10.1257/aer.99.2.45

———, “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market,” *American Economic Review*, 103 (2013), 1553–1597. https://doi.org/10.1257/aer.103.5.1553

Autor, David H., David Dorn, and Gordon H. Hanson, “The China Syndrome: Local Labor Market Effects of Import Competition in the United States,” *American Economic Review*, 103 (2013), 2121–2168. https://doi.org/10.1257/aer.103.6.2121

———, “The China Shock: Learning from Labor-Market Adjustment to Large Changes in Trade,” *Annual Review of Economics*, 8 (2016), 205–240. https://doi.org/10.1146/annurev-economics-080315-015041

———, “On the Persistence of the China Shock,” *Brookings Papers on Economic Activity* (2021), 381–447. https://doi.org/10.1353/eca.2022.0005

Autor, David H., David Dorn, Gordon H. Hanson, and Jae Song, “Trade Adjustment: Worker-level Evidence,” *Quarterly Journal of Economics*, 129 (2014), 1799–1860. https://doi.org/10.1093/qje/qju026

Autor, David H., Claudia Goldin, and Lawrence F. Katz, “Extending the Race between Education and Technology,” *AEA Papers and Proceedings*, 110 (2020), 347–351. https://doi.org/10.1257/pandp.20201061

Autor, David H., Lawrence F. Katz, and Melissa S. Kearney, “The Polarization of the US Labor Market,” *American Economic Review*, 96 (2006), 189–194. https://doi.org/10.1257/000282806777212620

Autor, David H., Lawrence F. Katz, and Alan B. Krueger, “Computing Inequality: Have Computers Changed the Labor Market?,” *Quarterly Journal of Economics*, 113 (1998), 1169–1213. https://doi.org/10.1162/003355398555874

Autor, David H., Frank Levy, and Richard J. Murnane, “The Skill Content of Recent Technological Change: An Empirical Exploration,” *Quarterly Journal of Economics*, 118 (2003), 1279–1333. https://doi.org/10.1162/003355303322552801

Babina, Tania, Anastassia Fedyk, Alex Xi He, and James Hodson, “Artificial Intelligence, Firm Growth, and Product Innovation,” SSRN Working Paper no. 3651052, 2020. https://doi.org/10.2139/ssrn.3651052

Balsvik, Ragnhild, Sissel Jensen, and Kjell G. Salvanes, “Made in China, Sold in Norway: Local Labor Market Effects of an Import Shock,” *Journal of Public Economics*, 127 (2015), 137–144. https://doi.org/10.1016/j.jpubeco.2014.08.006

Bárány, Zsófia L., and Christian Siegel, “Job Polarization and Structural Change,” *American Economic Journal: Macroeconomics*, 10 (2018), 57–89. https://doi.org/10.1257/mac.20150258

Berkes, Enrico, “Comprehensive Universe of U.S. Patents (CUSP): Data and Facts.” University of Maryland Working Paper, 2018.

Berman, Eli, John Bound, and Zvi Griliches, “Changes in the Demand for Skilled Labor within U.S. Manufacturing: Evidence from the Annual Survey of Manufactures,” *Quarterly Journal of Economics*, 109 (1994), 367–397. https://doi.org/10.2307/2118467

Bernard, Andrew B., J. Bradford Jensen, and Peter K. Schott, “Survival of the Best Fit: Exposure to Low-Wage Countries and the (Uneven) Growth of U.S. Manufacturing Plants,” *Journal of International Economics*, 68 (2006), 219–237. https://doi.org/10.1016/j.jinteco.2005.06.002

Böhm, Michael J., Hans-Martin von Gaudecker, and Felix Schran, “Occupation Growth, Skill Prices, and Wage Inequality,” CESifo Working Paper no. 7877, 2022. https://doi.org/10.2139/ssrn.3468012

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, “Quasi-Experimental Shift-Share Research Designs,” *Review of Economic Studies*, 89 (2022), 181–213. https://doi.org/10.1093/restud/rdab030

Branstetter, Lee G., Brian K. Kovak, Jacqueline Mauro, and Ana Venancio, “The China Shock and Employment in Portuguese Firms,” NBER Working Paper no. 26252, 2019. https://doi.org/10.3386/w26252

Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond, “Generative AI at Work,” NBER Working Paper no. 31161, 2023. https://doi.org/10.3386/w31161

Brynjolfsson, Erik, and Tom Mitchell, “What Can Machine Learning Do? Workforce Implications,” *Science*, 358 (2017), 1530–1534. https://doi.org/10.1126/science.aap8062
Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock, “What Can Machines Learn, and What Does It Mean for Occupations and the Economy?,” *AEA Papers and Proceedings,* 108 (2018), 43–47. https://doi.org/10.1257/pandp.20181019
Card, David, and Thomas Lemieux, “Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis,” *Quarterly Journal of Economics*, 116 (2001), 705–746. https://doi.org/10.1162/00335530151144140
Chen, Jiafeng, and Jonathan Roth, “Logs with Zeros? Some Problems and Solutions,” *Quarterly Journal of Economics*, 139 (2024), 891–936. https://doi.org/10.1093/qje/qjad054
Chiquiar, Daniel, “Globalization, Regional Wage Differentials and the Stolper–Samuelson Theorem: Evidence from Mexico,” *Journal of International Economics*, 74 (2008), 70–93. https://doi.org/10.1016/j.jinteco.2007.05.009
Comin, Diego A., Ana Danieli, and Martí Mestieri, “Income-Driven Labor Market Polarization,” Federal Reserve Bank of Chicago Working Paper no. 2020-22, 2020. https://doi.org/10.21033/wp-2020-22
Cortes, Guido Matias, Nir Jaimovich, Christopher J. Nekarda, and Henry E. Siu, “The Dynamics of Disappearing Routine Jobs: A Flows Approach,” *Labour Economics*, 65 (2020), 101823. https://doi.org/10.1016/j.labeco.2020.101823
Curtis, E. Mark, Daniel G. Garrett, Eric C. Ohrn, Kevin A. Roberts, and Juan Carlos Suárez Serrato, “Capital Investment and Labor Demand,” NBER Working Paper no. 29485, 2021. https://doi.org/10.3386/w29485
Dauth, Wolfgang, Sebastian Findeisen, and Jens Suedekum, “The Rise of the East and the Far East: German Labor Markets and Trade Integration,” *Journal of the European Economic Association*, 12 (2014), 1643–1675. https://doi.org/10.1111/jeea.12092
Dechezleprêtre, Antoine, David Hémous, Morten Olsen, and Carlo Zanella, “Induced Automation: Evidence from Firm-Level Patent Data,” University of Zurich, Department of Economics, Working Paper no. 384, 2021. https://doi.org/10.2139/ssrn.3835089
DellaVigna, Stefano, and Joshua M. Pollet, “Demographics and Industry Returns,” *American Economic Review*, 97 (2007), 1667–1702. https://doi.org/10.1257/aer.97.5.1667
Deming, David J., and Kadeem Noray, “Earnings Dynamics, Changing Job Skills, and STEM Careers,” *Quarterly Journal of Economics*, 135 (2020), 1965–2005. https://doi.org/10.1093/qje/qjaa021
Devlin, Allison, Brian K. Kovak, and Peter M. Morrow, “The Long-Run Labor Market Effects of the Canada-US Free Trade Agreement,” Carnegie Mellon University Manuscript, 2021.
DiNardo, John, Nicole M. Fortin, and Thomas Lemieux, “Labor Market Institutions and the Distribution of Wages, 1973–1992: A Semiparametric Approach,” *Econometrica*, (1996), 1001–1044. https://doi.org/10.2307/2171954
Doms, Mark, Timothy Dunne, and Kenneth R. Troske, “Workers, Wages, and Technology,” *Quarterly Journal of Economics*, 112 (1997), 253–290. https://doi.org/10.1162/003355397555181
Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock, “GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models,” arXiv Working Paper, 2023. https://doi.org/10.48550/arXiv.2303.10130
Felten, Edward W., Manav Raj, and Robert Seamans, “A Method to Link Advances in Artificial Intelligence to Occupational Abilities,” *AEA Papers and Proceedings*, 108 (2018), 54–57. https://doi.org/10.1257/pandp.20181021
———, “The Effect of Artificial Intelligence on Human Labor: An Ability-Based Approach,” *Academy of Management Proceedings,* (2019), 15784. https://doi.org/10.5465/AMBPP.2019.140

———“Occupational Heterogeneity in Exposure to Generative AI,” SSRN Working Paper 4414065, 2023. https://dx.doi.org/10.2139/ssrn.4414065

Frey Carl, Benedikt, *The Technology Trap*. (Princeton, NJ: Princeton University Press, 2019).

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy, “Text as Data,” *Journal of Economic Literature*, 57 (2019), 535–574. https://doi.org/10.1257/jel.20181020

Goldin, Claudia, and Lawrence F. Katz, “The Origins of Technology-Skill Complementarity,” *Quarterly Journal of Economics*, 113 (1998), 693–732. https://doi.org/10.1162/003355398555720

———, *The Race between Education and Technology* (Cambridge, MA: Harvard University Press, 2008).

Goldin, Claudia, and Robert A. Margo, “The Great Compression: The Wage Structure in the United States at Mid-Century,” *Quarterly Journal of Economics*, 107 (1992), 1–34. https://doi.org/10.2307/2118322

Goos, Maarten, and Alan Manning, “Lousy and Lovely Jobs: The Rising Polarization of Work in Britain,” *Review of Economics and Statistics*, 89 (2007), 118–133. https://doi.org/10.1162/rest.89.1.118

Goos, Maarten, Alan Manning, and Anna Salomons, “Job Polarization in Europe,” *American Economic Review*, 99 (2009), 58–63. https://doi.org/10.1257/aer.99.2.58

———, “Explaining Job Polarization: Routine-Biased Technological Change and Offshoring,” *American Economic Review*, 104 (2014), 2509–2526. https://doi.org/10.1257/aer.104.8.2509

Haanwinckel, Daniel, “Supply, Demand, Institutions, and Firms: A Theory of Labor Market Sorting and the Wage Distribution,” NBER Working Paper 31318, 2023. https://doi.org/10.3386/w31318

Hirvonen, Johannes, Aapo Stenhammar, and Joonas Tuhkuri, “New Evidence on the Effect of Technology on Employment and Skill Demand,” SSRN Working Paper 4081625, 2022. https://doi.org/10.2139/ssrn.4081625

Humlum, Anders, “Robot Adoption and Labor Market Dynamics,” University of Chicago Working Paper, 2021.

Katz, Lawrence F., and David H. Autor, “Changes in the Wage Structure and Earnings Inequality,” in *Handbook of Labor Economics,* vol. 3, O. Ashenfelter and D. Card, eds. (Amsterdam: Elsevier, 1999), 1463–1555. https://doi.org/10.1016/S1573-4463(99)03007-2

Katz, Lawrence F., and Kevin M. Murphy, “Changes in Relative Wages, 1963–1987: Supply and Demand Factors,” *Quarterly Journal of Economics*, 107 (1992), 35–78. https://doi.org/10.2307/2118323

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy, “Measuring Technological Innovation over the Long Run,” *American Economic Review: Insights*, 3 (2021), 303–320. https://doi.org/10.1257/aeri.20190499

Kim, Gueyon, “Trade-Induced Adoption of New Work,” IZA Discussion Paper no. 15165, 2022. https://doi.org/10.2139/ssrn.4114724

Kogan, Leonid, Dimitris Papanikolaou, Lawrence D.W. Schmidt, and Bryan Seegmiller, “Technology-Skill Complementarity and Labor Displacement: Evidence from Linking Two Centuries of Patents with Occupations,” NBER Working Paper no. 29552, 2021.

Kovak, Brian K., “Regional Effects of Trade Reform: What Is the Correct Measure of Liberalization?,” *American Economic Review*, 103 (2013), 1960–1976. https://doi.org/10.1257/aer.103.5.1960

Krusell, Per, Lee E. Ohanian, José-Víctor Ríos-Rull, and Giovanni L. Violante, “Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis,” *Econometrica*, 68 (2000), 1029–1053. https://doi.org/10.1111/1468-0262.00150

Leonardi, Marco, “The Effect of Product Demand on Inequality: Evidence from the United States and the United Kingdom,” *American Economic Journal: Applied Economics*, 7 (2015), 221–247. https://doi.org/10.1257/app.20130359

Lewis, Ethan, “Immigration, Skill Mix, and Capital Skill Complementarity,” *Quarterly Journal of Economics*, 126 (2011), 1029–1069. https://doi.org/10.1093/qje/qjr011

Lin, Jeffrey, "Technological Adaptation, Cities, and New Work," *Review of Economics and Statistics*, 93 (2011), 554–574. https://doi.org/10.1162/REST_a_00079
Mann, Katja, and Lukas Püttmann, "Benign Effects of Automation: New Evidence from Patent Texts," *Review of Economics and Statistics*, 105 (2023), 562–579. https://doi.org/10.1162/rest_a_01083
Mazzolari, Francesca, and Giuseppe Ragusa, "Spillovers from High-Skill Consumption to Low-Skill Labor Markets," *Review of Economics and Statistics*, 95 (2013), 74–86. https://doi.org/10.1162/REST_a_00234
Michaels, Guy, Ashwini Natraj, and John Van Reenen, "Has ICT Polarized Skill Demand? Evidence from Eleven Countries over Twenty-Five Years," *Review of Economics and Statistics*, 96 (2014), 60–77. https://doi.org/10.1162/REST_a_00366
Michel, Jean-Baptiste, Yuan Kui Shen, Aviva Presser Aiden, Adrian Veres, Matthew K Gray, Google Books Team, Joseph P Pickett, Dale Hoiberg, Dan Clancy, and Peter Norvig et al., "Quantitative Analysis of Culture Using Millions of Digitized Books," *Science*, 331 (2011), 176–182. https://doi.org/10.1126/science.1199644
Moser, Petra, "Innovation without Patents: Evidence from World's Fairs," *Journal of Law and Economics*, 55 (2012), 43–74. https://doi.org/10.1086/663631
Mullahy, John, and Edward C. Norton, "Why Transform Y? A Critical Assessment of Dependent-Variable Transformations in Regression Models for Skewed and Sometimes-Zero Outcomes," NBER working paper no. 30735 (2022). https://doi.org/10.3386/w30735
Noy, Shakked, and Whitney Zhang, "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence," *Science*, 381 (2023), 187–192. https://doi.org/10.1126/science.adh2586
Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer, "The Impact of AI on Developer Productivity: Evidence from GitHub Copilot," arXiv working paper, 2023. https://doi.org/10.48550/arXiv.2302.06590
Pennington, Jeffrey, Richard Socher, and Christopher Manning, "GloVe: Global Vectors for Word Representation," in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP),* Association for Computational Linguistics, Doha, Qatar, 1532–1543. https://doi.org/10.3115/v1/D14-1162
Pierce, Justin R., and Peter K. Schott, "The Surprisingly Swift Decline of US Manufacturing Employment," *American Economic Review*, 106 (2016), 1632–1662. https://doi.org/10.1257/aer.20131578
Ray, Debraj, and Dilip Mookherjee, "Growth, Automation, and the Long-Run Share of Labor," *Review of Economic Dynamics*, 46 (2022), 1–26. https://doi.org/10.1016/j.red.2021.09.003
Ruggles, Steven, Catherine A. Fitch, Ronald Goeken, J. David Hacker, Matt A. Nelson, Evan Roberts, Megan Schouweiler, and Matthew Sobek, "IPUMS Ancestry Full Count Data: Version 3.0 [dataset]," 2021.
Ruggles, Steven, Sarah Flood, Ronald Goeken, Megan Schouweiler, and Matthew Sobek, "IPUMS USA: Version 12.0," 2022. https://doi.org/10.18128/D010.V12.0
Susskind, Daniel, "A Model of Task Encroachment in the Labour Market," Oxford University Discussion Paper 819, 2020.
Topalova, Petia, "Factor Immobility and Regional Impacts of Trade Liberalization: Evidence on Poverty from India," *American Economic Journal: Applied Economics*, 2 (2010), 1–41. https://doi.org/10.1257/app.2.4.1
U.S. Department of Lab or, Employment and Training Administration, *Dictionary of Occupational Titles* (Washington, DC: U.S. Government Printing Office, 1939).
———, *Dictionary of Occupational Titles*, 4th ed. (Washington, DC: U.S. Government Printing Office, 1977).
U.S. Census, "Database with Images: Orville Wright in Entry for Milton Wright, 1900," 1900.

———, "Database with Images: Orville Wright in Household of Milton Wright, Dayton Ward 5, Montgomery, Ohio, United States," 1910.

U.S. Census Bureau, *Census of Population: Alphabetical Index of Industries and Occupations*, US Bureau of the Census, 1915–2018.

U.S. Department of Labor: Employment and Training Administration, "O*NET OnLine: Solar Energy Installation Managers, 47-1011.03," 2022.

Vogel, Jonathan, "The Race between Education, Technology, and the Minimum Wage," NBER Working Paper no. 31028, 2023. https://doi.org/10.3386/w31028

Webb, Michael, "The Impact of Artificial Intelligence on the Labor Market," SSRN Working Paper, 2020. https://dx.doi.org/10.2139/ssrn.3482150