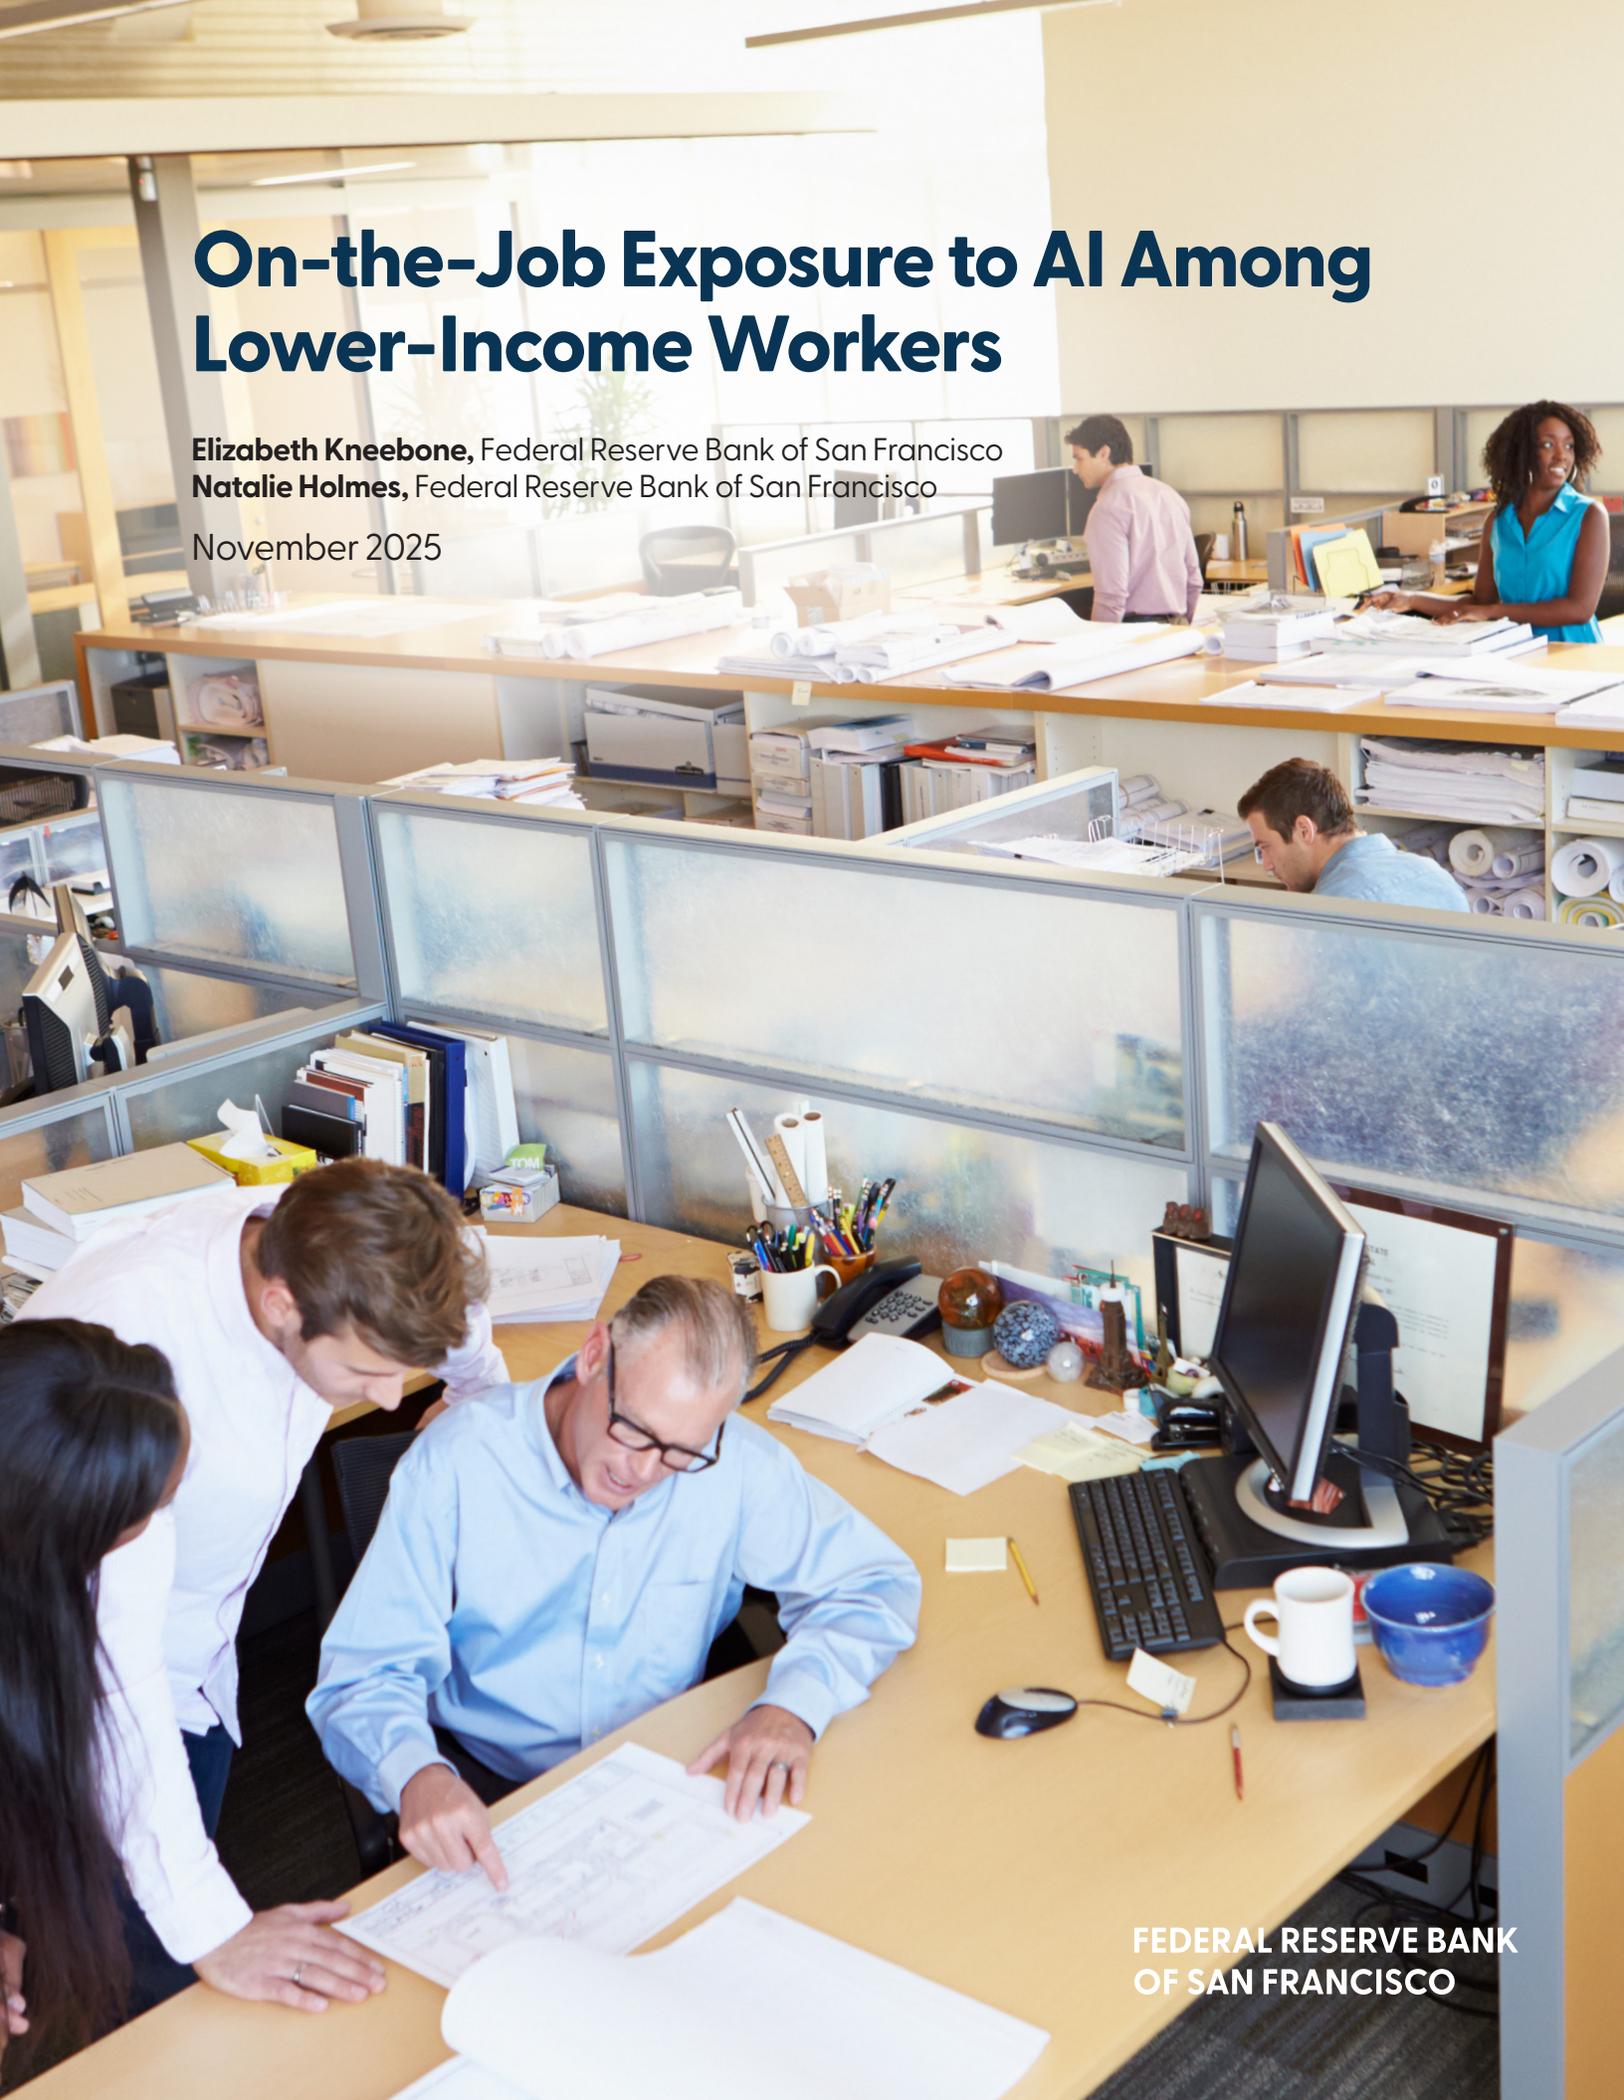


On-the-Job Exposure to AI Among Lower-Income Workers

Elizabeth Kneebone, Federal Reserve Bank of San Francisco
Natalie Holmes, Federal Reserve Bank of San Francisco

November 2025



Acknowledgments

Thank you to the roundtable and listening session participants and members of the Federal Reserve Bank of San Francisco’s Community Advisory Council who shared their valuable insights for this analysis. And thanks to members of the Community Engagement and Analysis team at the Federal Reserve Bank of San Francisco who provided feedback on our methodology and to our colleagues at the San Francisco and Atlanta Federal Reserve Banks who helped organize and facilitate roundtables and listening sessions. Thanks, as well, to Scott Allard, Sergio Galeano, Julia Karon, Sarah Miller, Mark Muro, and Bina Shrimali for their comments on earlier versions of this draft, and to Crystal Ejanda for her editorial guidance and production assistance.

Disclaimer

The views expressed in this report are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System.

Suggested Citation

Kneebone, Elizabeth and Natalie Holmes. 2025. “On-the-Job Exposure to AI Among Lower-Income Workers.” Federal Reserve Bank of San Francisco Community Development Research Brief 2025-03. doi: 10.24148/cdrb2025-03.

Executive Summary

To better understand the potential impacts of generative AI (gen AI) on the economy, this analysis uses quantitative methods to assess the extent to which workers are likely to be exposed to AI on the job, paying particular attention to workers in lower-income households, the occupations and industries in which they work, and how exposure varies across different parts of the country. It also draws on qualitative insights to understand how the impacts of AI integration are showing up in real time and how workforce and training organizations, nonprofits, and employers are adapting.

Key Takeaways

Analysis of 2023 American Community Survey microdata reveals that among workers highly exposed to AI, those in lower-income households:

- Account for more than 6 million (20%) of all AI-exposed workers.
- Are more likely than average to work in Office and Administrative Support occupations and to work in service-oriented industries, such as Health Care and Social Assistance.
- Tend to be older, to have higher levels of educational attainment, and to be higher-earning than lower-income workers as a whole.
- Make up varying shares of the AI-exposed workforce across the country and work in differing mixes of occupations and industries depending on their local labor markets.

Seven roundtable and listening sessions with nearly 60 participants—including employers, workforce system representatives, community-based organizations, and community colleges, among others—yielded the following insights about real-time AI adoption and impacts:

- Adoption and integration of AI into roundtable participants' organizations and operations varied substantially, but for those integrating AI into their operations, early employment impacts were already apparent.

- Respondents saw many ways in which AI could be beneficial to lower-income workers and job seekers but emphasized the need for critical thinking skills to make AI adoption successful.
- Respondents expressed concerns about uneven impacts of AI adoption worsening outcomes for vulnerable workers, unless adequate guardrails and supports are in place.

Contents

- Key Takeaways 2
- Introduction..... 5
- Methods..... 6
- Quantitative Findings..... 9
- Qualitative Findings..... 20
- Conclusion..... 25
- Appendix A: Methods..... 27
- Appendix B: Workers Exposed with High Exposure to AI, by State, 2023 31
- References..... 33
- Notes 35

Introduction

The advent of generative AI is ushering in changes that stand to impact workers, employers, and the economy writ large. There is increasing consensus that generative AI is a “general purpose technology”—that is, one that could transform the economy and society as we know it, akin to the steam engine, personal computers, and the Internet (Calvino, Haerle, and Liu 2025). As with these earlier technological innovations, it is difficult to predict the specific ways in which AI will alter the economic landscape and whether any benefits will be broadly (or more unevenly) shared.

As part of the U.S. central bank, the Federal Reserve Bank of San Francisco works to keep the U.S. economy strong by focusing on stable prices, full employment, safe banking, and secure payments. An important aspect of this work is to understand the economic experiences of employers, workers, and communities across the Twelfth District, including how they may be impacted by emerging technologies like generative AI. Technological advances, for all the innovations and productivity gains they can unlock in the economy, tend to be accompanied by transitions that can shift employment patterns among certain occupations, industries, and populations of workers. While these “transition dynamics” mark periods of adjustments in the economy that typically stabilize over time, understanding the segments of the economy, including the scale and makeup of the workforce, most likely to be impacted by these transitions is important to understanding near-term implications for full employment and household financial stability.

A growing body of research suggests that workers most likely to hold jobs exposed to AI integrations (whether to augment or replace certain tasks) tend to be higher-earning, white-collar workers (Eckhardt and Goldschlag 2025). However, the impacts of AI adoption are likely to be felt across a broader swath of the workforce, especially as the technology continues to evolve. To the extent that impacts are positive, one question is how broadly shared will those benefits be across the income spectrum of workers? The same could be asked of any downside risks: To what extent might earnings or job losses be felt across different populations of workers and earning levels? Our previous research on workers in lower-income households suggests that these workers have the fewest resources to navigate any such AI-related job or earnings losses (Kneebone and Holmes 2025).

The purpose of this analysis is to build on the existing literature on AI exposure to provide insights into how many and what kinds of workers may be relatively highly exposed to AI on the job, paying particular attention to workers at the lower end of the income distribution. After briefly summarizing the methods used in this analysis, we present the findings of a quantitative analysis estimating occupational exposure to AI and assessing employment and demographic characteristics and geographic variation, with a focus on how these patterns differ (or align) for lower-income workers. Given that “exposure” reflects *potential* or prospective change, we also include qualitative insights drawn from roundtables and listening sessions with dozens of workforce stakeholders to understand how they are seeing AI adoption and its employment and workforce impacts evolve in real time. We conclude with considerations of directions for further inquiry.

Methods

This analysis uses a mixed methods approach. For the quantitative analysis, we draw on an O*NET-based ranking of occupational exposure to AI, as well as detailed 2023 American Community Survey (ACS) microdata to identify workers highly exposed to AI on the job.ⁱ We explore employment and demographic characteristics of highly exposed workers, with a particular focus on workers living in low- to moderate-income households (also referred to as lower-income workers). We also draw on qualitative analysis to incorporate more real-time insights into AI adoption and its impacts, described in more detail below.

Estimating AI Exposure

While it is too early to measure the direct effects of AI on employment by occupation, the idea of “exposure” to AI helps to conceptualize and estimate the potential impacts of this evolving technology. Different approaches have been taken to define and measure AI exposure (e.g., Felten, Raj, and Seamans 2023; Kochhar 2023; Schendstok and Schreiner Wertz 2024). By and large, studies measuring exposure have yielded directionally similar results (Eckhardt and Goldschlag 2025).

To determine AI exposure, we use an O*NET-based metric developed by Schendstok and Schreiner Wertz (2024), following classifications developed by Kochhar (2023).ⁱⁱ In

keeping with this previous research, we consider AI exposure to reflect the extent to which work tasks could be either aided or replaced by gen AI tools, including Chat GPT, Claude, Copilot, or Gemini, among others.

O*NET provides ratings of the importance of 41 work activities related to job performance in a given occupation.ⁱⁱⁱ The O*NET-based metric we use in this analysis is based on the following questions:

- What is the likelihood (high, medium, or low) a work activity could be replaced or aided by AI tools at this time?
- What is the relative importance of high-AI-exposure work activities to the occupation in question?

Occupations are then ranked by the relative importance of work activities with a high likelihood of being replaced or aided by AI. For the purposes of this analysis, we consider the top 25% of occupations “highly exposed” to AI (also referred to as high AI exposure occupations). This is not to suggest that other occupations are not exposed to AI; rather, it is meant to help identify occupations that are particularly primed for technological adaptations or substitutions given the nature of primary work tasks, and is consistent with previous research (e.g., Kochhar 2023).

Work tasks with a high likelihood of being supplemented or substituted by AI include such activities as getting information, processing information, scheduling work and activities, controlling machines and processes, and performing administrative activities. (For more details on the methodology used to develop this metric, see [Appendix A.](#))

Note that this analysis does not attempt to measure actual job loss, creation, or transformation, although our qualitative analysis, described below, yields some anecdotal insights into how organizations are seeing AI affect hiring, tasks assignments, and training considerations. We also acknowledge that this is a quickly evolving landscape. As gen AI tools continue to expand and improve, tasks and occupations assessed as having medium or low AI exposure for the purposes of this analysis may become more highly exposed in the future. The development of agentic AI or AI paired with advanced robotics is also beyond this analysis, but both are deserving of their own analysis of potential employment and economic impacts as those technologies continue to evolve.

Identifying Lower-Income Workers

To better understand the extent to which lower-income workers may be affected by AI exposure on the job, we use methods developed in our previous research to identify workers in low- to moderate-income (also referred to as lower-income) households (Kneebone and Holmes, 2025). Considering the earnings of individual occupations or workers can offer a limited view of resource availability or constraints. For example, an occupation may pay low wages, but individuals may hold more than one job to increase their income; or workers may report low earnings but do so because they have other income sources (e.g., another earner in the household). Therefore, we start at the household level. We first identify households with combined resources that qualify them as low- to moderate-income (i.e., below 80% of the area median income). We then identify the workers within the household. For the purposes of this analysis, we consider workers to be those 16 years and older with earned income in the past year from civilian employment.

By this definition, 48.9 million workers lived in a lower-income household in 2023, accounting for 28% of all workers with earned income in the past year. Compared with the workforce as a whole, workers in a lower-income household are more likely to work in occupations including Office and Administrative Support, Sales, Food Preparation and Serving Related, and Health Care Support—all of which report below-average typical annual earnings. Lower-income workers also tend to skew somewhat younger than average and have lower levels of educational attainment. (For more details, see Kneebone and Holmes 2025).

Soliciting Qualitative Insights

To help contextualize our quantitative analysis and capture more real-time insights about how the emergence of AI is impacting workers and employers, we held five roundtables with 46 participants in different regions (New Orleans, Phoenix, Portland, Salt Lake City, and San Francisco) between April and July 2025. We also held two listening sessions with the Federal Reserve Bank of San Francisco's 11-member Community Advisory Council in February and September 2025.

Participants in these roundtables and listening sessions represented a range of perspectives, including workforce boards, training and other community-based organizations, financial institutions, community colleges and higher-education

institutions, philanthropy, and employers representing a range of industries. Participants provided reactions to the findings of our quantitative analysis, and they offered insights into the extent to which—and the ways in which—AI tools have affected their work, organizations, and industries, including implications for employment and efforts to assist lower-income workers in navigating technological transitions related to evolving AI tools. The qualitative data gathered from these participants provide insights into the way this technological transition is playing out for different kinds of workers and organizations that cannot be gleaned from the quantitative data alone.

Quantitative Findings

Among workers highly exposed to AI, more than 6 million live in a lower-income household.

Based on 2023 ACS data, 30.6 million workers—more than 17% of all workers with earned income in the past year—qualified as highly exposed to AI. Among those workers, 20%—or one in five—lived in a lower-income household.

The majority of workers highly exposed to AI—whether lower-income or not—were employed in the private sector (Table 1). However, compared with all workers with high-AI-exposure jobs, those in lower-income households were more likely to be either self-employed or to be employed in the nonprofit sector.

Table 1. Workers Highly Exposed to AI by Class of Employment, 2023

	Private	Government	Self-	Nonprofit
Workers Highly Exposed to AI	Employment		Employed	
All	67.5%	15.7%	8.9%	7.7%
Workers in Lower-Income Households	64.4%	14.4%	11.6%	9.2%

Source: Author analysis of 2023 American Community Survey microdata.

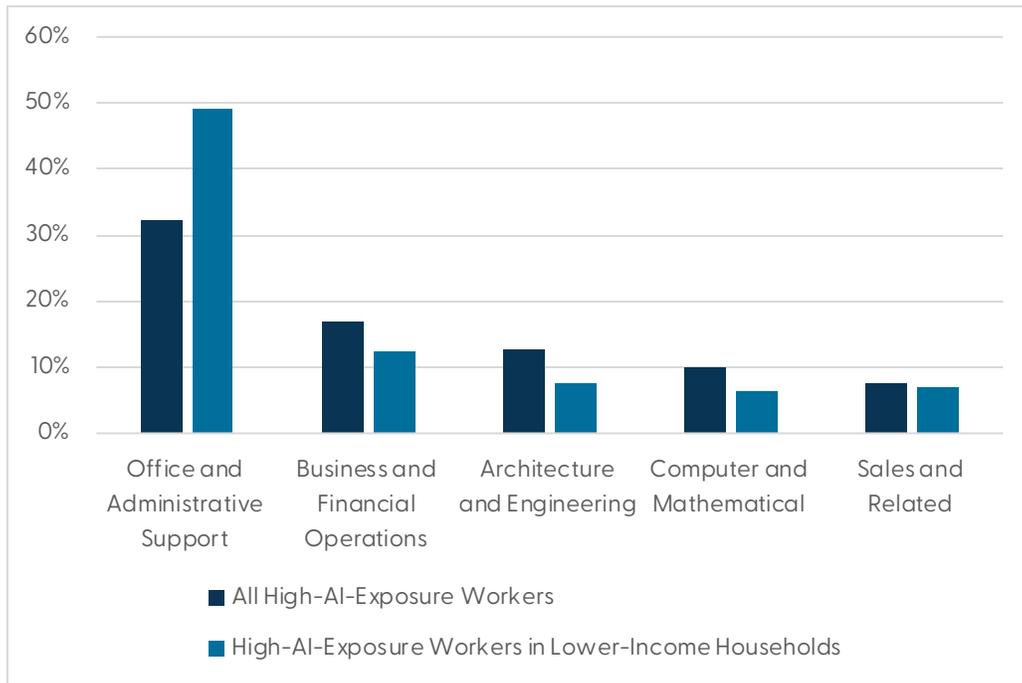
Note: All differences between all workers and workers in lower-income households are significant at the 95% confidence level.

Even greater variation emerges when considering the kinds of tasks and jobs performed by workers in lower-income households, compared with the high-AI-exposure workforce as a whole.

Lower-income workers highly exposed to AI are more likely than average to work in Office and Administrative Support occupations.

Given the nature of the technology, it is to be expected that AI-exposed workers cluster in certain kinds of occupations where required tasks tend to be both computer-based and conducive to technological augmentation or automation; however, AI-exposed workers in lower-income households proved to be especially concentrated in certain kinds of jobs in 2023. The leading occupations subject to high AI exposure were the same regardless of worker income status (Figure 1): Based on 2023 ACS data, Office and Administrative Support occupations led the list for workers with job tasks especially likely to be impacted by AI, followed by Business and Financial Operations and Architecture and Engineering occupations. Yet, while one in three AI-exposed workers overall worked in office and administrative jobs, fully half of lower-income workers highly exposed to AI worked in those occupations.

Figure 2. Most Common Occupations Held by Workers Highly Exposed to AI, 2023



Source: Author analysis of 2023 American Community Survey microdata.

Note: All differences are significant at the 95% confidence level. “Other” not shown.

The most common types of jobs for lower-income workers in Office and Administrative Support occupations included secretaries and administrative assistants, receptionists, and office clerks (Table 2). Other highly exposed jobs outside of Office and Administrative Support include a range of positions, from human resources workers and computer support specialists to cashiers and engineers. The average annual earnings reported by workers in lower-income households who held these high AI exposure jobs in 2023 ranged from less than \$14,000 (cashiers) to almost \$53,000 (software developers).

Table 2. Most Common Jobs Among Selected Occupation Categories for Lower-Income Workers Highly Exposed to AI, 2023

Office and Administrative Support	Business and Financial Operations	Architecture and Engineering	Computer and Mathematical	Sales and Related
<ul style="list-style-type: none"> • Secretaries and administrative assistants • Receptionists and information clerks • Office clerks 	<ul style="list-style-type: none"> • Accountants and auditors • Human resources workers • Management analysts 	<ul style="list-style-type: none"> • Other engineering technologists and technicians • Civil engineers • Engineers, all other 	<ul style="list-style-type: none"> • Software developers • Computer occupations, all other • Computer support specialists 	<ul style="list-style-type: none"> • Cashiers • Retail salespersons • First-Line supervisors of retail sales workers

Source: Author analysis of 2023 American Community Survey microdata.

Note: The “Secretaries and administrative assistants” category excludes legal, medical, and executive assistants; the “Other engineering technologists and technicians” category excludes drafters.

While occupational designations shed light on the job tasks likely to be highly exposed to AI, understanding potential exposure by industry can offer another important perspective. Just as certain types of jobs may distribute differently across different kinds of industries, the adoption (or potential for the adoption) of new technology can transmit differently across sectors.

High-AI-exposure workers with lower incomes are more likely to work in service-oriented industries such as Health Care and Social Assistance, and that is especially true if they hold administrative support jobs.

In 2023, the largest share of high-AI-exposure jobs was found in the Professional, Scientific, and Technical Services industry (which includes organizations that provide legal services, accounting and payroll services, advertising, and computer systems design services, among others), regardless of worker income status (Table 2). But

patterns diverge in marked ways from there: Compared with all workers in high-AI-exposure jobs, lower-income workers with high AI exposure were more likely to be employed in industries such as Health Care and Social Assistance (which includes establishments such as doctors’ offices and hospitals, as well as community food and housing organizations), Transportation and Warehousing, and Retail. They were also more likely to hold jobs in service-oriented industries, including Educational Services and Administrative and Support and Waste Management Services.

Table 3. Most Common Industries for Workers Highly Exposed to AI, 2023

	Lower-Income Workers Highly Exposed to AI	All Workers Highly Exposed to AI
Professional, Scientific, and Technical Services	17.8%	22.1%
Health Care and Social Assistance	12.6%	9.1%
Finance and Insurance	10.0%	11.8%
Manufacturing	8.3%	12.1%
Educational Services	7.4%	5.7%
Transportation and Warehousing	6.9%	4.3%
Public Administration	6.6%	8.7%
Retail Trade	5.4%	4.0%
Other Services, Except Public Administration	4.2%	2.9%
Administrative and Support and Waste Management Services	4.1%	2.9%

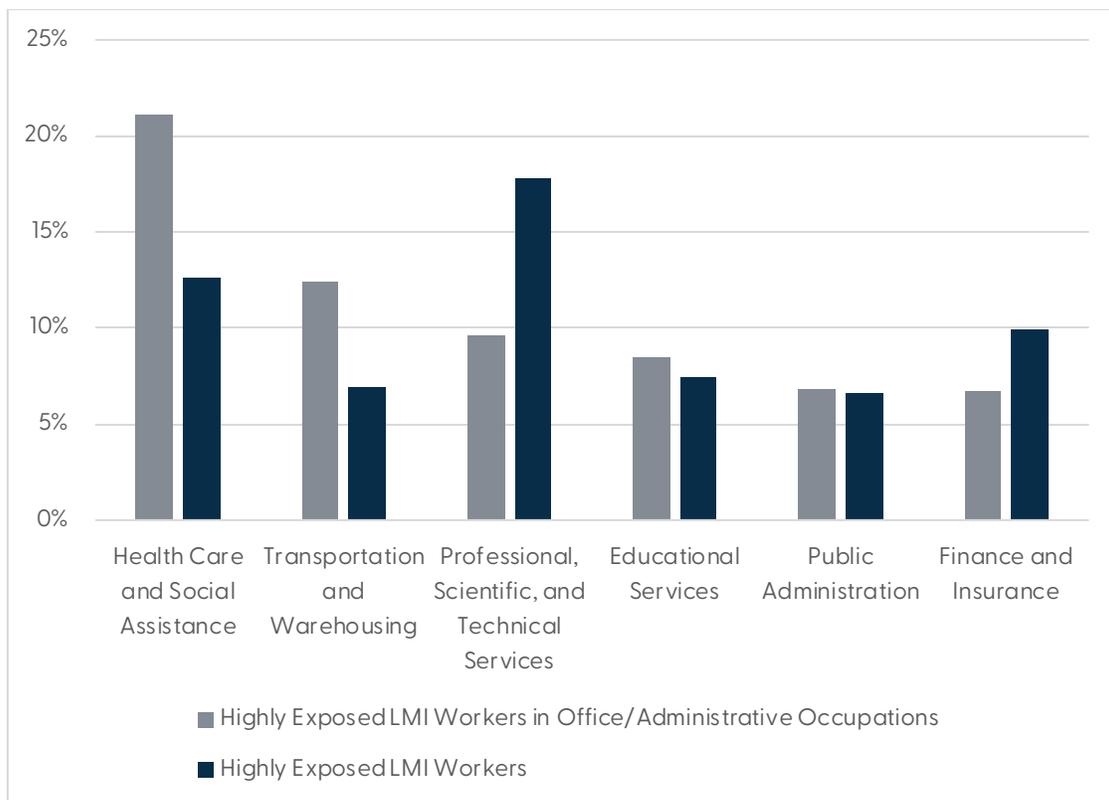
Source: Author analysis of 2023 American Community Survey microdata.

Note: All differences are significant at the 95% confidence level.

However, these patterns can shift when considering the intersection of particular occupations and industries. Take, as an example, the roughly 50% of lower-income, high-AI-exposure workers that work in Office and Administrative Support occupations. The distribution of those workers by industry departs markedly from the overall average for AI-exposed lower-income workers (Figure 2). Rather than Professional, Scientific, and Technical Services being the most common industry for

office and administrative workers, Health Care and Social Assistance leads the list, with more than one in five lower-income workers with high AI exposure holding jobs in that industry. Those workers were also much more likely than average to work for a nonprofit organization: Roughly 21% of lower-income office and administrative workers highly exposed to AI were employed by a nonprofit in 2023, which is more than double the share for AI-exposed lower-income workers as a whole.

Figure 2. Most Common Industries Among Lower-Income Workers Highly Exposed to AI Who Work in Office and Administrative Support Occupations, 2023



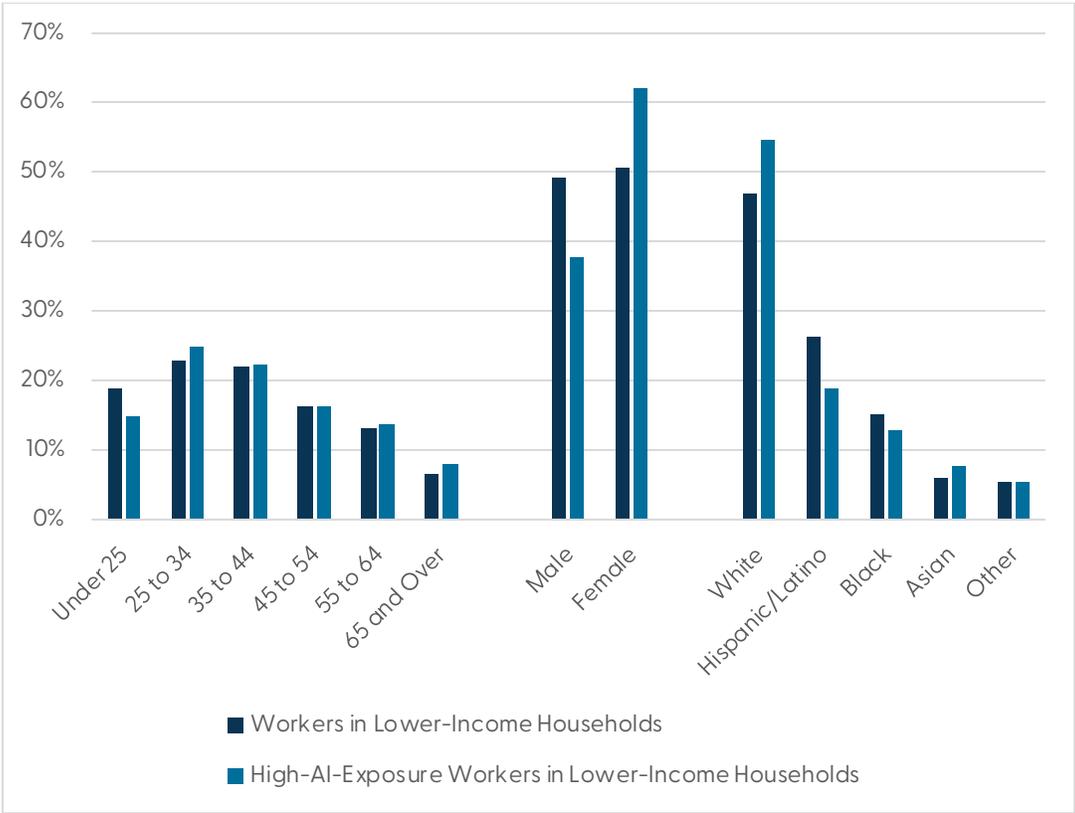
Source: Author analysis of 2023 American Community Survey microdata.

Note: All differences are significant at the 95% confidence level, except in Public Administration.

Lower-income workers highly exposed to AI tend to be older, to have higher levels of educational attainment, and to be higher-earning than lower-income workers as a whole.

As noted above, earlier research on the occupational exposure of all workers to AI found certain groups were more likely to hold high-exposure jobs, including older workers, women, and workers who are White or Asian. The same patterns hold true for workers living in lower-income households (Figure 3). Compared with all lower-income workers, those in jobs highly exposed to AI were more likely to be over age 25. They were also much more likely to be female (62% compared with 51%), and more likely to be White or Asian.

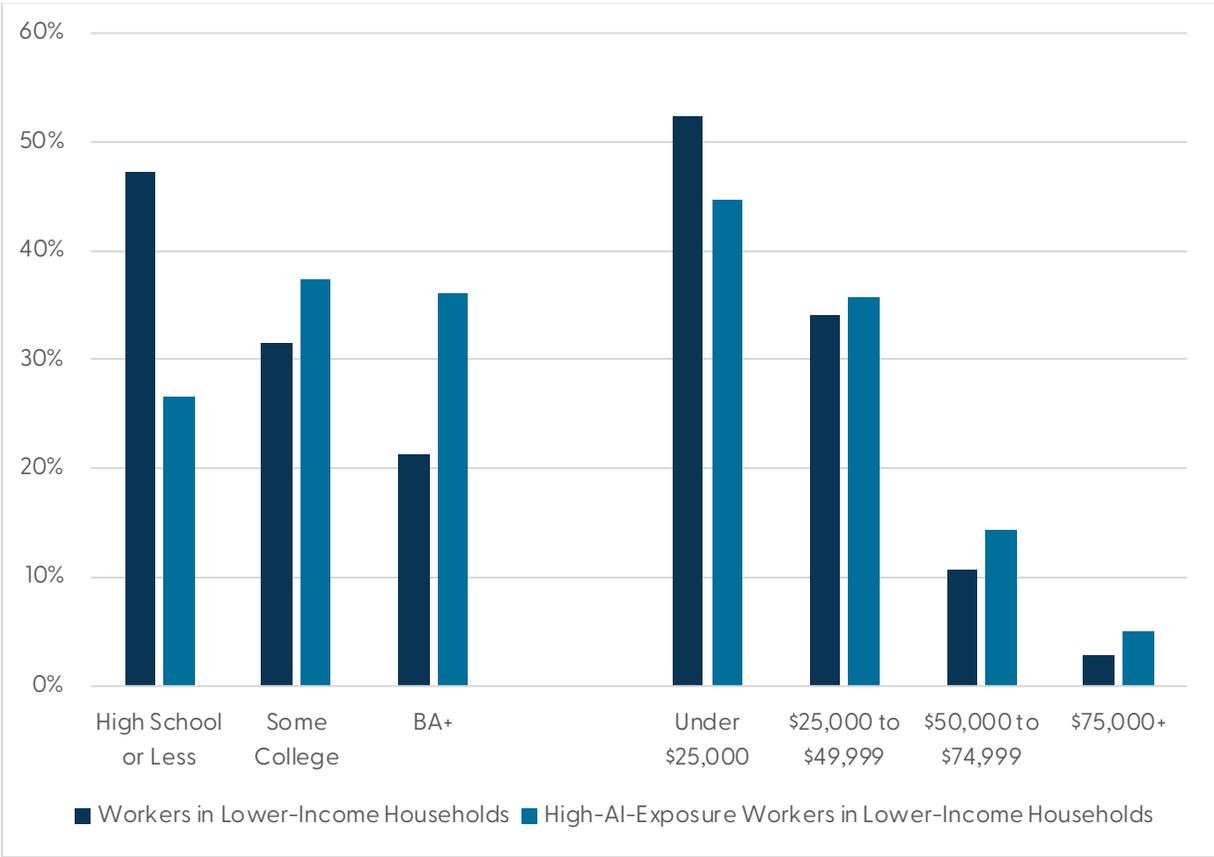
Figure 3. Demographic Characteristics of Lower-Income Workers by AI Exposure, 2023



Source: Author analysis of 2023 American Community Survey microdata.
 Note: All differences are significant at the 95% confidence level, except for the following categories: 35 to 44, 45 to 54, and “Other.”

Educational attainment and earnings differences between all lower-income workers and those highly exposed to AI also mirror previous research on AI exposure among the workforce as a whole. Compared with all lower-income workers, those with jobs highly exposed to AI were more likely to have earned a bachelor’s degree or completed at least some college (which could include associate’s degrees or other certifications). Perhaps commensurate with their somewhat older age profiles and higher levels of education, lower-income workers highly exposed to AI were also more likely to report higher earnings: The median annual income for all lower-income workers was \$23,000 in 2023, but for those with high-AI-exposure jobs it was \$28,000. And 20% of AI-exposed workers in lower-income households earned more than \$50,000 a year in 2023 (Figure 4).

Figure 4. Educational Attainment and Earnings of Lower-Income Workers by AI Exposure, 2023



Source: Author analysis of 2023 American Community Survey microdata.

Note: All differences are significant at the 95% confidence level.

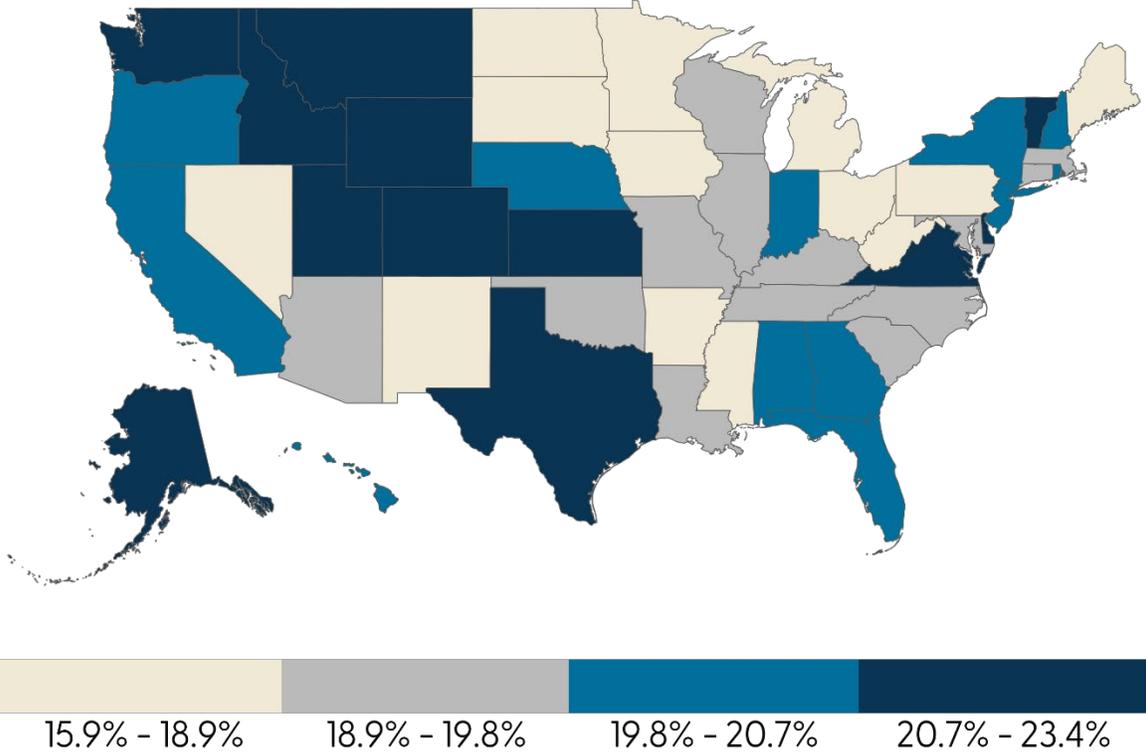
Many of the patterns surfaced in the above findings help explain the differences that emerge in the makeup of lower-income workers highly exposed to AI. For instance, workers employed in the nonprofit sector are more likely to be women and to have relatively higher levels of educational attainment. Similarly, employees in Office and Administrative Support occupations and in the Health Care and Social Services industry skew female and are more likely than lower-income workers, on average, to have some college education and to earn \$50,000 to \$75,000.

The share of workers with high AI exposure who live in lower-income households varies considerably across the country, as do the kinds of occupations and industries in which they are employed.

Just as states and regional labor markets vary in their mix of jobs and industries and in their demographic makeup, the distribution and degree of AI exposure—and the extent to which it affects lower-income workers—ranges across the country. The share of workers in high-AI-exposure jobs ranges across the country (see Appendix B for detailed estimates), as does the share of those workers living in lower-income households (Map 1). In 2023, lower-income workers made up a higher-than-average share of workers with high AI exposures in states such as Delaware (23.4%) and Virginia (22.7%) along the eastern seaboard, Kansas (21.0%) and Nebraska (20.6%) in the Midwest, and in Utah (22.7%), Washington (21.6%), Idaho (21.4%), Alaska (21.3%), Oregon (20.7%), and Hawai'i (20.5%) in the Federal Reserve's Twelfth District.

At the other end of the spectrum, the share of AI-exposed workers living in lower-income households reached as low as 15.9% in North Dakota. Shares ranged to an even greater degree at the sub-state level, exceeding 40% in communities such as Arizona's Gila, Graham, and Greenlee Counties and falling below 5% in Mississippi's South Delta region.

Map 1. Share of Workers Highly Exposed to AI Who Live in a Lower-Income Household, by State, 2023

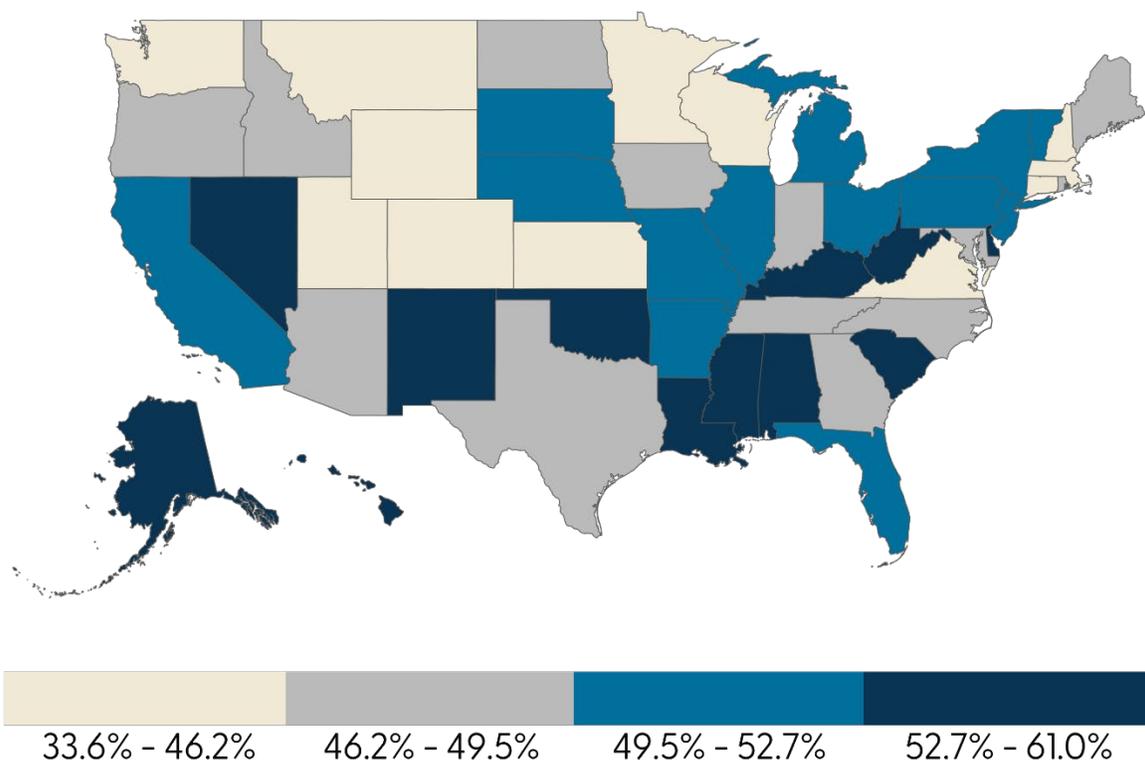


Source: Author analysis of 2023 American Community Survey microdata.

The makeup of AI-exposed jobs among lower-income workers also varies by location. While Office and Administrative Support occupations were the most common in every state in 2023, the prevalence of that category shifted depending on the makeup of the local labor market (Map 2). In states such as Alaska, Louisiana, Mississippi, and Nevada, office and administrative jobs accounted for more than half of all jobs held by AI-exposed workers in lower-income households, and that share reached as high as 61% in Hawai'i. However, in states such as Washington and Utah, office and administrative jobs accounted for below-average shares of jobs held by lower-income workers highly exposed to AI (42.6% and 37.1%, respectively). Those states saw higher-than-average shares of lower-income workers with high AI exposure working in Architecture and Engineering occupations (in Washington's case) or in Business and

Financial Operations (in Utah’s case), as well as in Computer and Mathematical occupations (in both states).

Map 2. Share of Lower-Income Workers Highly Exposed to AI Who Work in Office and Administrative Support Occupations, by State, 2023



Source: Author analysis of 2023 American Community Survey microdata.

The industry makeup of AI-exposed jobs varies as well across states. In all but eight states, Professional, Scientific, and Technical Services was the most common industry for lower-income workers highly exposed to AI, following the national pattern. But the share of jobs in that industry ranged widely, from 31.1% in Washington, D.C. to 12.5% in Indiana. For the eight states that diverged from the national average (Alaska, Arkansas, Hawai'i, Kentucky, Minnesota, South Dakota, and West Virginia), the Health Care and Social Services industry led for percentage of lower-income workers highly exposed to AI on the job, with that share reaching as high as 24.9% in Alaska.

Qualitative Findings

While findings from our quantitative analysis help provide a sense of the potential scale of high AI exposure on the job and offer insights into the makeup and geographic distribution of lower-income workers employed in high-exposure occupations, they do not shed light on the extent to which workers or the organizations that employ them have actually adopted AI tools. The following findings emerged from qualitative engagement with 57 roundtable and listening session participants—drawing on the perspectives of workforce boards, training and other community-based organizations, financial institutions, community colleges and higher education institutions, philanthropy, and employers representing a range of industries—and offer more real-time insights into AI adoption and its early impacts based on their experiences.

Adoption and integration of AI into roundtable participants’ organizations and operations varied substantially.

In our roundtables and listening sessions, nearly all participants characterized their organizations as being in a “learning phase” in terms of exploring and integrating the use of generative AI tools. On one hand, a minority of respondents reported that they were not yet experimenting with new AI tools (or were prohibited from doing so) as their organizations assess risks and try to “understand what is safe,” especially for organizations providing social services who must protect client data. One participant noted they “still don’t know who actually holds the data” that gets entered into these tools, so they were exercising caution.

On the other hand, a few respondents reported developing their own tools to help tailor and deliver their services. For nonprofits working to develop service-oriented or case management tools, they reported having access to flexible philanthropic resources to support that exploration. Those unable to use AI to date were often more reliant on less responsive, more risk-averse sources of funding.

A much more common response from respondents was that they fell somewhere in between these two ends of the spectrum. Many reported seeing individual employees experimenting in ad hoc ways with commonly available free or low-cost tools to find efficiencies in their day-to-day tasks. Others reported taking a more intentional

organization-wide approach to adopting particular tools, encouraging staff to integrate them into routinized tasks, such as notetaking or communications support (e.g., drafting emails or social media text) or in specific functions. In the case of nonprofit respondents, once such function mentioned was fundraising (e.g., using tools to help draft grant applications for nonprofits or assemble necessary documentation for compliance reporting to funders). Businesses we spoke with reported exploring the use of AI to help achieve operational efficiencies such as streamlining customer-service operations (e.g., to more quickly assess and direct clients to the right assistance) and specialized project management. At the same time, employers in the financial services and health care fields noted that they are considering the existing regulatory environment (e.g., as it pertains to protection of patient or client privacy) and how it might evolve as they adopt or tailor these tools.

Many respondents reflected that, whether or not they were actively seeking out these tools, AI was being integrated into common tools and platforms they were already using. One nonprofit leader noted the pace of change and integration, saying, “We were naïve enough to think we had a choice, but it’s all happening so fast anyway.” Other participants wondered whether employees even fully understood the extent to which they may already be using AI in these types of integrated platforms, and if that would only become more the case as adoption grows.

For those integrating AI into their operations, early employment impacts were already apparent.

For nonprofit and business representatives who reported adopting AI tools in their work, one commonly cited driver for doing so was to address resource constraints. Nonprofits reported challenges with an uncertain and increasingly constrained funding environment being a driver of their explorations of ways to use AI to increase efficiencies. Multiple businesses we spoke with noted that their motivations for adopting AI stemmed from efforts to backfill positions they were unable to successfully recruit for and to help retain skilled positions, in many cases by having AI alleviate undesirable administrative burdens from skilled workers.

These organizations also reported a range of ways AI adoption was already reshaping their employment decisions. In some cases, AI tool integration led these employers to shift tasks within a position, freeing up capacity for other tasks by using AI to replace repetitive administrative work. Respondents also reported impacts on

hiring. Some are hiring differently (e.g., recalibrating the level of experience required, knowing the role could be augmented by AI tools), while others said they are holding off on hiring for roles they had planned to add to see if the needs could be met through AI adoption among current staff. In some cases, respondents reported eliminating existing positions because they were able to reduce their headcount through their adoption of AI.

Respondents saw many ways in which AI could be beneficial to lower-income workers and job seekers but emphasized the need for critical thinking skills to make AI adoption successful.

Across regions and types of organizations, participants agreed that facility with AI tools is shaping up to be a new essential skill. As one person put it, "If you work with a computer, you'll need to know how to work with AI." But among the workforce training organizations and other nonprofits offering services to lower-income workers we spoke with, the extent to which they had adapted their services to equip clients with AI-oriented skills varied.

One key question we heard from workforce boards was whether AI-focused training would meet the needs of job seekers or whether it would quickly become obsolete, given the pace of change. Some respondents reported hesitancy to create new AI training opportunities without clarity from employers about their needs. Others were grappling with the extent to which they should focus training in areas unlikely to be affected by AI at this time.

At the same time, some respondents noted the ways in which their case managers were already instructing their job seeking clients how to use AI tools to improve their résumés and tailor their cover letters to address job posting requirements or to prepare for interviews. A community college administrator shared how their institution was working to integrate AI tools into classroom curricula to make sure students were familiar with how to use these tools in the course of their work, with the goal of better preparing them for an AI-integrated workplace.

A common concern among participants was the need to make sure their clients and students were learning how to use these tools as inputs to support their work rather than as the actual work output. One respondent noted, for instance, that AI can support job seekers who are English-language learners (ELLs) by helping to translate

job postings and prepare for interview questions. However, if it is clear that job candidates are simply copying and pasting output without sufficient comprehension, it can actually worsen their chances of getting hired. That concern was not unique to ELLs and was echoed by participants working with a broad range of students and workforce clients.

That concern relates to a point all participants using AI in their work agreed on: Employing AI tools effectively demands critical thinking. Critical thinking skills have long been cited (see, e.g., Fee 2025) as a leading priority for employers in workforce recruitment, and respondents predicted they will be even more in demand as the “human skills” that ensure effective implementation of AI become even more important. In fact, workforce and training respondents shared that they had yet to see notable shifts in job postings or employer requests seeking specific AI-related skills. Rather, respondents emphasized the need to deepen focus on training around those transferable and in-demand “soft skills” that could help students and workers be successful to adapting to a rapidly evolving AI landscape and whatever new tools may emerge.

Respondents expressed concerns about uneven impacts of AI adoption worsening outcomes for vulnerable workers, unless adequate guardrails and supports are in place.

When considering the localized data from this analysis presented in each of these listening sessions, many saw opportunities for AI to help job seekers match to good jobs and be aided in their work. However, participants also worried about how the downside impacts might worsen outcomes for exposed workers. For instance, when it comes to exposure among Office and Administrative Support jobs, some participants expressed concerns that if those jobs are negatively impacted it could mean the erosion of jobs often held by women and workers who tend to be higher-earning than lower-income workers as a whole.

More broadly, participants at every roundtable voiced concerns about older workers—given that AI-exposed workers skew somewhat older—as well as others who are not computer- or AI-“native” and what it might take to help them build AI literacy. At the same time, participants noted their worries about workers at the beginning of

their careers being unable to find entry-level office jobs that would allow them to develop skills necessary to advance up the career ladder.

Participants also pointed to the existing digital divide, sharing concerns that those with digital literacy gaps or without reliable internet access risk being left further behind as these tools become more prevalent. Respondents also highlighted that many lower-income workers access the internet primarily through mobile devices, which may limit functionality for many common work functions.

In each roundtable and listening session, participants reflected that eventually these tools will become integrated across the education and training ecosystem, starting with primary education and extending through post-secondary education and workforce training programs. But during this time of transition, respondents largely agreed that questions remain about how to best help students and workers adapt to these new technologies, especially workers who are not currently engaged with formal education or training but may be experiencing on-the-job impacts—or even employment loss—because of the integration of AI tools.

In addition to exploring new training opportunities and strategies, several participants suggested there may be lessons learned from past technological and industrial transitions, including considering the need for something akin to Trade Adjustment Assistance for the age of AI.^{iv} Participants noted such supports would not only provide a buffer for lower-income workers experiencing AI-related employment impacts, but also for higher-earning workers at risk of downward mobility due to transitional dynamics associated with AI implementation.

Conclusion

The quantitative analysis presented here suggests that, although higher-earning, white-collar workers are among those most highly exposed to AI given the nature of their jobs, that is not where AI's employment impacts are likely to end. More than 6 million lower-income workers number among the highly exposed workforce, accounting for 20% of workers highly exposed to AI at work. These lower-income workers are also made up largely of office workers in administrative support and business and financial operations—just at the lower end of the earnings distribution. They are more likely to be working in the nonprofit sector or self-employed compared with AI exposed workers as a whole. And although AI exposure in these roles may offer opportunities to augment lower-income workers' skills and earnings potential, the stakeholders we spoke with also worry that the nature of the tasks associated with these positions may lend themselves to automation to a degree that will cause the erosions of these positions, which are often among the better-paid jobs lower-income workers tend to hold.

In the roundtables and listening sessions we held, there was a general recognition that individual firms and organizations were making decisions about AI adoption that made sense for optimizing their operations, especially in the case of nonprofits struggling with budget shortfalls amid a shifting funding environment and businesses trying to overcome hiring challenges for certain functions. But questions remained about how those individual decisions might aggregate into systemic impacts that could lead to uneven outcomes of this technological transition. To mitigate potential uneven impacts, stakeholders we spoke with noted the need to think about how to best support impacted workers—across the experience and earnings spectrum—citing the role of training and job-search assistance, as well as wraparound supports to help them weather potential spells of unemployment or retraining.

As this technological landscape evolves, there is an opportunity for research efforts to better understand exposure across a wider range of workers (e.g., expanding the lens to understand what “exposure” might look like in terms of task or timing of adoption in occupations beyond the top 25% explored here, and how the makeup of workers in those occupations might differ). Questions also remain about how exposure translates into impact and the extent to which impact equates to augmentation or replacement. These are outcomes that are likely to vary across different occupations, industries,

and groups of workers and that could shift the population of lower-income workers in potential need of transition assistance—either by diminishing it (in the case of net positive augmentation effects) or expanding it (depending on the number of higher-earning impacted workers who become downwardly mobile). As the workforce support ecosystem adopts different approaches to supporting the education and training of—and providing wraparound supports for—exposed or impacted workers, there are also opportunities to evaluate what proves most effective and share promising efforts that could be scaled to improve outcomes and ensure the benefits of this evolving technology are broadly shared.

Appendix A: Methods

Amid the rapid emergence and adoption of generative AI (gen AI) tools, such as ChatGPT, Claude, Copilot, Gemini, and other gen AI-powered applications, a body of research has evolved in an effort to estimate and assess how this technology could impact—and is already impacting—the economy. Some of this research has focused on assessing employment impacts (e.g., Brynjolfsson, Chandar, and Chen 2025; Wynne and Derr 2025), differences in adoption (e.g., Handa et al. 2025; Appel et al. 2025), geographic differences (e.g., Muro, Methkuppally, and Kinder 2025), and changes in employer demand (e.g., Galeano, Hodge, and Ruder 2025).

Given how nascent this technological transition still is, a growing body of literature has also focused on the concept of “exposure,” considering the tasks required to do a given job and the potential for gen AI to help a person do—or replace the need for a person to do—those tasks. The understanding of AI exposure at the task level can then be aggregated to consider job-level exposure and how exposure might “cluster” within particular occupations, industries, and geographies or for different kinds of workers, populations, and communities.

There are different ways of conceptualizing and measuring exposure (e.g., Felten, Raj, and Seamans 2023; Kochhar 2023; Schendstok and Schreiner Wertz 2024). By and large, studies on this topic have yielded directionally similar results (Eckhardt and Goldschlag 2025).

This analysis uses a mixed methods approach. For the quantitative analysis, we draw on an O*NET-based ranking of occupational exposure to AI, as well as detailed 2023 American Community Survey (ACS) microdata to identify workers highly exposed to AI on the job.^v The ACS microdata also allow us to explore employment and demographic characteristics of highly exposed workers. We pay particular attention to workers living in low- to moderate-income households (also referred to as lower-income workers)—an under explored segment of the workforce in the existing literature on AI exposure. We also draw on qualitative analysis to incorporate more real-time insights into AI adoption and its impacts, described in more detail below.

Estimating AI Exposure

To determine AI exposure, we use an O*NET-based metric developed by Schendstok and Schreiner Wertz (2024), following classifications developed by Kochhar (2023). Specifically, Kochhar (2023) developed categorizations of O*NET skills data based on likely AI exposure for a Current Population Survey-based analysis. Schendstok and Schreiner Wertz (2024) used the same methods but adapted them to ACS microdata. Schendstok and Schreiner Wertz shared their metrics with us for this analysis.

O*NET (i.e., U.S. Department of Labor’s Occupational Information Network, Version 28) provides ratings of the importance of 41 work activities related to job performance in a given occupation, using a scale of 0 (not important) to 7 (extremely important).

The categorization used by Schendstok and Schreiner Wertz and developed by Kochhar first classifies each of the 41 work activities captured in O*NET as having high, medium, or low exposure to AI. Exposure is defined as the likelihood generative AI tools (e.g., ChatGPT, Claude, Copilot, Gemini) could aid or replace a given task. High exposure activities include such tasks as getting information, processing information, scheduling work and activities, controlling machines and processes, and performing administrative activities.

Once activities have been classified as having high, medium, or low exposure, the next step is to determine the relative importance of those tasks to job performance. Schendstok and Schreiner Wertz calculated average importance ratings for each group of work activities tagged as having high, medium, and low exposure. They then normalized those measures to determine the relative importance of each group of activities for each occupation.

Using Schendstok and Schreiner Wertz’s normalized calculations, we rank occupations by the relative importance of work activities with a high likelihood of being replaced or aided by AI. For the purposes of this analysis, we consider the top 25% of occupations “highly exposed” to AI (also referred to as “high AI exposure” occupations). This is not to suggest other occupations are not exposed to AI; rather, it is meant to help identify occupations that are particularly primed for technological adaptations or substitutions given the nature of primary work tasks. It is also consistent with Kochhar, who selected the top quartile of *occupations* as “highly exposed,” although it departs from Schendstok and Schreiner Wertz, who designated the top quartile of *workers* as “highly exposed.”

Note that this analysis does not attempt to measure actual job loss, creation, or transformation (although our qualitative engagements, described below, yield some anecdotal insights into how organizations are seeing AI affect hiring, task assignments, and training considerations). We also acknowledge that this is a quickly evolving landscape, and as AI tools continue to expand and improve, tasks and occupations assessed as having medium or low AI exposure for the purposes of this analysis may become more highly exposed in the future. The development of agentic AI or AI paired with advanced robotics is also beyond this analysis, but both are deserving of their own analysis of potential employment and economic impacts as those technologies continue to evolve.

Identifying Lower-Income Workers

To better understand the extent to which lower-income workers may be affected by AI exposure on the job, we use methods developed in our previous research to identify workers in lower-income households (Kneebone and Holmes 2025). Rather than identify lower-income *people* or *occupations*, which can limit insights into resource availability or constraints a worker faces, we start at the household level. People reporting lower earnings may do so because they work in a low-wage job or because they choose to work less, given other income coming into the household. And an occupation may pay relatively lower wages, but a worker may take on additional employment to increase total earnings. Starting with the household level can provide a fuller picture of combined income—from all earners and income sources—and allow us to identify households that fall in the lower portion of the income distribution, meaning they face greater resource constraints relative to the typical household in their area.

We define lower-income households as those with incomes below 80% of the area median income for a given household size (a common threshold used in a range of government programs and regulations to identify lower-income people, families, and/or communities^{vi}). We then identify the number of workers within the household, which includes all workers aged 16 years and over in the civilian labor force. We include any person aged 16 years and over who reported earned income within the past year in recognition of the greater volatility that exists in employment spells among lower-income workers.^{vii}

By this definition, 48.9 million workers lived in a lower-income household in 2023, accounting for 28% of all workers with earned income in the past year. Compared with the workforce as a whole, workers in a lower-income household (also referred to as “lower-income workers” in this analysis) are more likely to work in occupations including Office and Administrative Support, Sales, Food Preparation and Serving Related, and Health Care Support—all of which report below-average typical annual earnings. Lower-income workers also tend to skew somewhat younger than average and have lower levels of educational attainment. (For more details, see Kneebone and Holmes 2025.)

Soliciting Qualitative Insights

In addition to our quantitative analysis, we also held five roundtables between April and July 2025 with a combined 46 participants. The roundtables took place in New Orleans, Phoenix, Portland, Salt Lake City, and San Francisco. We also held two listening sessions with the Federal Reserve Bank of San Francisco’s 11-member Community Advisory Council in February and September 2025. Participants represented a mix of perspectives and backgrounds and included representatives from workforce boards, training and other community-based organizations and financial institutions, community colleges and higher-education institutions, philanthropy, and employers.

Roundtable participants provided reactions to the findings of our quantitative analysis and helped refine methods and presentation of the data. We also solicited information from our roundtable and listening session participants on the extent to which—and the ways in which—AI tools have affected their work, organizations, and industries. We probed for implications related to employment and asked about efforts to assist lower-income workers in navigating technological transitions related to evolving AI tools. The qualitative data gathered from these participants provide insights into the way this technological transition is playing out for different kinds of workers and organizations that cannot be gleaned from the quantitative data alone.

Appendix B: Workers Exposed with High Exposure to AI, by State, 2023

	All Workers with Earned Income in the Past Year	In Lower-Income Households	All Workers with High AI Exposure	High AI Exposure in Lower-Income Households	Share of AI-Exposed Workers in Lower-Income Households
United States	175,555,600	48,925,050	30,557,273	6,058,299	19.8%
Alabama	2,449,834	658,733	395,081	78,691	19.9%
Alaska	368,685	99,511	69,081	14,735	21.3%
Arizona	3,739,818	1,014,433	673,203	131,777	19.6%
Arkansas	1,473,047	375,539	221,489	40,579	18.3%
California	20,289,092	5,927,560	3,654,066	724,629	19.8%
Colorado	3,384,432	958,072	651,760	140,267	21.5%
Connecticut	1,992,072	561,158	369,238	70,634	19.1%
Delaware	538,374	151,110	89,021	20,849	23.4%
District of Columbia	396,334	118,463	125,830	25,733	20.5%
Florida	11,422,274	3,165,553	1,979,899	400,606	20.2%
Georgia	5,727,743	1,584,411	975,335	194,299	19.9%
Hawai'i	700,104	190,462	116,784	23,978	20.5%
Idaho	1,012,776	284,828	160,543	34,328	21.4%
Illinois	6,779,811	1,909,708	1,171,107	221,405	18.9%
Indiana	3,615,347	963,921	566,138	112,701	19.9%
Iowa	1,760,157	490,112	269,070	48,099	17.9%
Kansas	1,580,114	429,222	259,594	54,517	21.0%
Kentucky	2,239,880	589,547	354,011	68,919	19.5%
Louisiana	2,189,054	604,131	364,954	69,295	19.0%
Maine	765,160	179,273	114,146	20,653	18.1%
Maryland	3,383,105	973,307	697,890	133,950	19.2%
Massachusetts	3,965,018	1,106,536	754,610	149,232	19.8%
Michigan	5,207,894	1,409,022	915,466	167,844	18.3%
Minnesota	3,245,466	882,020	558,422	101,138	18.1%
Mississippi	1,378,914	352,221	187,181	32,211	17.2%

Missouri	3,280,177	872,111	555,485	105,857	19.1%
Montana	608,203	170,014	105,219	22,279	21.2%
Nebraska	1,100,747	296,941	188,061	38,803	20.6%
Nevada	1,672,769	460,968	263,115	49,567	18.8%
New Hampshire	811,377	210,353	132,823	26,884	20.2%
New Jersey	5,099,820	1,484,533	947,514	188,299	19.9%
New Mexico	1,025,282	280,444	183,065	32,535	17.8%
New York	10,246,540	2,837,909	1,912,568	384,072	20.1%
North Carolina	5,591,153	1,548,783	902,715	175,974	19.5%
North Dakota	433,317	110,218	76,926	12,198	15.9%
Ohio	6,215,665	1,653,530	1,019,962	186,702	18.3%
Oklahoma	1,981,451	560,967	317,248	62,308	19.6%
Oregon	2,262,958	654,132	411,952	85,100	20.7%
Pennsylvania	6,803,879	1,798,648	1,174,198	221,476	18.9%
Rhode Island	600,919	157,113	102,200	20,775	20.3%
South Carolina	2,699,483	720,966	436,122	82,846	19.0%
South Dakota	508,333	138,096	74,679	13,011	17.4%
Tennessee	3,683,530	988,163	581,012	110,625	19.0%
Texas	15,954,777	4,715,132	2,719,825	563,672	20.7%
Utah	1,881,359	535,816	339,004	76,804	22.7%
Vermont	358,158	99,385	61,619	13,765	22.3%
Virginia	4,584,871	1,312,868	893,996	203,261	22.7%
Washington	4,170,465	1,196,135	748,889	161,690	21.6%
West Virginia	807,959	196,071	128,342	22,811	17.8%
Wisconsin	3,259,491	861,024	540,529	105,558	19.5%
Wyoming	308,412	85,877	46,286	10,358	22.4%

Source: Author analysis of 2023 American Community Survey microdata.

References

- Appel, Ruth et al. 2025. "Anthropic Economic Index Report: Uneven Geographic and Enterprise AI Adoption." Anthropic.
<https://www.anthropic.com/research/anthropic-economic-index-september-2025-report>.
- Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen. 2025. "Canaries in the Coal Mine? Six Facts About the Recent Employment Effects of Artificial Intelligence." *Stanford Digital Economy Lab*. https://digitaleconomy.stanford.edu/wp-content/uploads/2025/08/Canaries_BrynjolfssonChandarChen.pdf.
- Butcher, Kristin, and Diane Whitmore Schanzenbach. 2018. "Most Workers in Low-Wage Labor Market Work Substantial Hours, in Volatile Jobs." Washington, DC: Center on Budget and Policy Priorities.
- Calvino, Flavio, Daniel Haerle, and Sarah Liu. 2025. "Is Generative AI a General Purpose Technology?: Implications for Productivity and Policy." *OECD Artificial Intelligence Papers*, No. 40, OECD Publishing. <https://doi.org/10.1787/704e2d12-en>.
- Eckhardt, Sarah, and Nathan Goldschlag. 2025. "AI and Jobs: The Final Word (Until the Next One)." *Economic Innovation Group* 2.4.6.
- Fee, Kyle. 2025. "Using the Occupational Mobility Explorer to Understand Which Skills Are Most in Demand." Federal Reserve Bank of Cleveland.
<https://doi.org/10.26509/frbc-cd-20250425>.
- Felten, Edward, Manav Raj, and Robert Seamans. 2023. "Occupational Heterogeneity in Exposure to Generative AI."
<https://ssrn.com/abstract=4414065>. or <http://dx.doi.org/10.2139/ssrn.4414065>.
- Galeano, Sergio, Nyerere Hodge, and Alexander Ruder. 2025. "By Degree(s): Measuring Employer Demand for AI Skills by Educational Requirements." Federal Reserve Bank of Atlanta. <https://doi.org/10.29338/wc2025-01>.
- Handa, Kunal, et al. 2025. "Which Economic Tasks Are Performed with AI? Evidence from Millions of Claude Conversations." *arXiv preprint arXiv:2503.04761*.
<https://arxiv.org/pdf/2503.04761>.

- Kneebone, Elizabeth and Natalie Holmes. 2025. "Understanding the Employment Patterns of Workers in Low- and Moderate-Income Households." Federal Reserve Bank of San Francisco. [doi: 10.24148/cdrb2025-01](https://doi.org/10.24148/cdrb2025-01).
- Kochhar, Rakesh. 2023. "Which U.S. Workers Are More Exposed to AI on Their Jobs?" Pew Research Center report. https://www.pewresearch.org/social-trends/wp-content/uploads/sites/3/2023/07/st_2023.07.26_ai-and-jobs.pdf.
- Muro, Mark, Shriya Methkuppally, and Molly Kinder. 2025. "The Geography of Generative AI's Workforce Impacts Will Likely Differ from Those of Previous Technologies." The Brookings Institution. <https://www.brookings.edu/articles/the-geography-of-generative-ais-workforce-impacts-will-likely-differ-from-those-of-previous-technologies/>.
- Ruggles, Steven, et al. 2025. IPUMS USA: Version 16.0. Minneapolis, MN: IPUMS. <https://doi.org/10.18128/D010.V16.0>.
- Schendstok, Matt and Sydney Schreiner Wertz. 2024. "Occupational Exposure to Artificial Intelligence by Geography and Education." U.S. Treasury Department. <https://home.treasury.gov/system/files/136/AI-Combined-PDF.pdf>.
- Wynne, Mark and Lillian Derr. 2025. "Will AI Replace Your Job? Perhaps Not in the Next Decade." Federal Reserve Bank of Dallas. <https://www.dallasfed.org/research/economics/2025/0603#:~:text=There%20is%20Over%20little%20evidence,About%20the%20authors>.

Notes

ⁱ American Community Survey microdata retrieved from IPUMS USA.

ⁱⁱ Kochhar (2023) developed categorizations of O*NET skills data based on likely AI exposure for a Current Population Survey-based analysis. Schendstok and Schreiner Wertz (2024) used the same categorizations but adapted them to ACS microdata. Schendstok and Schreiner Wertz shared their O*NET-based metrics, mapped to ACS occupation codes, with us for this analysis.

ⁱⁱⁱ O*NET is the U.S. Bureau of Labor Statistic’s Occupational Information Network, and this analysis references Version 28.

^{iv} Trade Adjustment Assistance is a federal program first authorized in 1962. It has four components—for workers, firms, farmers, and communities—tasked with reducing negative impacts of increased imports. For workers, assistance includes job training, relocation allowances, and income supports.

^v ACS microdata retrieved from IPUMS USA.

^{vi} For instance, under the Community Reinvestment Act, communities with median incomes below 80% of the area median are considered low- to moderate-income (see https://www.federalreserve.gov/consumerscommunities/cra_resources.htm#lmi). For U.S. Department of Housing and Urban Development program, such as public housing or vouchers-based housing assistance, 80% of area median income demarcates low-income families and individuals (see <https://www.hud.gov/helping-americans/public-housing>).

^{vii} See, e.g., Butcher, Kristin, and Diane Whitmore Schanzenbach. 2018. “Most Workers in Low-Wage Labor Market Work Substantial Hours, in Volatile Jobs.” Washington, DC: Center on Budget and Policy Priorities.



frbsf.org | [f](#) [📷](#) [in](#) [X](#)