How Will Artificial Intelligence Affect Employment?
Implications for Equity

Jennifer L. Steele
American University
Washington, DC
steele@american.edu

Association for Public Policy Analysis and Management
DRAFT IN PROGRESS
November 2024

Abstract

Analyzing U.S. Department of Labor O*NET data on the task composition of 867 jobs, I estimate the fraction of each job that may be vulnerable to automation with artificial intelligence (AI)-enabled tools in the next generation. I characterize 41 occupational tasks by known automation bottlenecks, including theory of mind, flexible dexterity, and vision/strategy, which are analogous to those used by Frey and Osborne, drawing on earlier work by Autor, Levy, and Murnane (2003). I test two assumptions about bottleneck severity, one with greater protection for manual tasks. In aggregating tasks to the occupational level, I weight them by their relative importance to the job and scale them by their level of sophistication, such that more sophisticated tasks are harder to automate. After examining the distribution of job automation risk by education level, interest category, and sector, I estimate three scenarios for how automation might affect productivity. I show that salary inequality rises in even the most protective scenarios, but especially when top-quartile workers capture most of the automation benefits in their sectors.

Introduction

Educational institutions have always had to prepare students for an economically changing landscape. Robotics and digital automation have been reshaping the U.S. economy since at least the 1980s, as assembly line, secretarial, and other routine work has shifted from humans to machines (Autor, 2014). But institutions of higher education may have felt insulated from these changes, which by default were boosting the value of advanced skills and a college degree (Deming, 2023; Goldin & Katz, 2008).

By hollowing out U.S. manufacturing jobs, which previously had offered a path to the middle class for non-college-educated workers, the first four decades of computer-based automation contributed to widening income inequality and to economic and even political polarization between U.S. workers with and without college degrees (Carnevale & Strohl, 2010; Piketty, 2018).

But generative AI is different. Multimodal large language models (LLMs) like OpenAI's GPT and Anthropic's Claude use natural language programming to generate new text, images, and video in ways that closely mimic human-created content. Trained on terabytes of human-generated output, these technologies rapidly generate text or pictures using many layers of predictive analytics of what should be written or drawn next, word by word and pixel by pixel, in light of the context of adjacent words or pixels (Lee & Trott, 2023). Enabled by massive computing power, they can generate creations in a matter of seconds that might demand many hours of work even for human experts. However, the models operate at the level of item prediction rather than via conceptual logic. For this reason, they are prone to many limitations, including reporting of fully made-up information ("hallucinations"), inability to parse logical arguments or structures, limitations in performing math, and limited ability to reason through complex logic problems (Dziri et al., 2022; Ji et al., 2022; Lohr, 2024). However, given attention that researchers are devoting to these problems, the capabilities of generative AI will surely increase.

In comparison to many technological innovations, generative AI has been rapidly embraced by the public. ChatGPT 3.5 was released as recently as November of 2022, catapulting generative AI into global awareness. According to an August 2024 survey by Bick, Blandin, & Deming (2024), about 39.5%

of U.S. adults aged 65 or younger reported already having used AI. About 28% had used it for work, including more than a fifth of blue-collar workers, and 33% had used it outside of work. Though company's adoption has been slower and more heterogeneous (McElheran et al., 2024), an increasing number are finding ways to increase revenue through generative AI (Chui et al., 2021). One experiment with 5000 call center workers given staggered implementation of customer support scripts found that problem resolutions per hour rose 14 percent, with most gains concentrated among new and less-skilled workers (Brynjolfsson et al., 2023).

So the question is not whether generative AI will begin to automate people's work, but how such automation will affect their jobs, enhancing workers' capabilities or supplanting them. On one hand is evidence like that from the call centers, in which AI complements and enhances workers' productivity, with the least-skilled workers yielding the greatest increases in value. This scenario supports David Autor's (2024) view that AI may shrink earnings inequality, given its applications to white-collar and creative work and its potential to boost the skills of lower-performing workers, as in the call center study.

On the other hand, Daron Acemoglu and colleagues (2022) show that firms exposed to AI increase hiring in AI subfields but scale back otherwise. This finding reinforces the longstanding concern that AI will eventually supplant workers by substituting for their work if firms are not incentivized to protect workers and help them use AI as a complement (Acemoglu & Johnson, 2023). In a study of more than 1000 scientists, Toner-Rodgers (2024) showed that AI access in a chemical research lab fostered a 44% gain in material discoveries, a 39% gain in patent filings, and a 17% gain in product development. However, contrary to the call center study by Brynjolfsson et al. (2023), Toner-Rodgers found that the gains in innovation were driven by the most experienced and high-producing scientists who were able to make efficient use of AI-generated suggestions. It therefore suggests that the ability to use AI as a complement rather than substitute for human-generated tasks may depend on both the complexity of the task in question and the skill level of the worker.

In this paper, I undertake analyses of Department of Labor O*NET data under six plausible scenarios for automation risk and salary redistribution. My aim is to gauge possible effects of each scenario on the inequality of automation vulnerability and salaries. I use importance and level data on 41 occupational tasks, which I characterize as automation bottlenecks if they depend on theory of mind, flexible dexterity, or vision/strategy. These are analogous to bottlenecks used by Frey and Osborne (2017), drawing on earlier work by Autor, Levy, and Murnane (2003).

I test two assumptions about these bottlenecks, the second of which is more protective of manual labor. Then I aggregate them to the occupational level, weighting by their relative importance to the job and scaling by their level of sophistication, such that more sophisticated tasks are harder to automate. I describe automation risks by education level, interest category, and sector, and I estimate three scenarios for how this might affect productivity as proxied by salary. I find that salary inequality is sensitive to how lower-skilled workers are able to capitalize their own skill automation. It is much less sensitive to considerations about which particular tasks represent smaller or larger barriers to automation as generative AI technology gradually improves.

Predicting Automation Risk Across Jobs

The longstanding question of which jobs are vulnerable to automation has recently been addressed in separate papers by Frey and Osborne (2017) and Felten, Raj, and Seamans (2021), both of whom use Occupational Information Network (O*NET) data from the U.S. Department of Labor. Frey and Osborne (2017) score job abilities based on automation bottlenecks: *social intelligence; complex perception and manipulation*; *creative intelligence*. They use data on the prevalence of nine automation technologies to calculate the probability that a given job can be wholly automated by AI or not. A limitation is that their model does not incorporate the most recent gains in AI creativity and treats jobs automation as dichotomous.

Felten, Raj, & Seamans (2021) use O*NET ability importance and level scores, combined with *crowd-sourced views on whether particular abilities are exposed to AI*, to predict AI Occupational Exposure (AIOE) and Industry Exposure (AIIE). A strength of their model is that they treat each job as a basket of characteristics—in this case abilities—that vary in whether AI can emulate them. A limitation is that their focus on exposure rather than vulnerability does not distinguish between work that might be complemented versus substituted by AI and does not consider physical types of automation, such as robotics-enabled surgery or self-driving cars. Their crowd-sourced approach to gauging AI exposure is admirably data-driven but atheoretical, which leads to high exposure scores for jobs like teaching and mental health counseling and social services that in my schema, as well as that of Frey and Osborne (2017) and Autor, Levy, and Murnane (2003) seem less vulnerable to automation because of their interpersonal element.

I construct a hybrid approach in which, like Felten, Raj, and Seamans (2021), I treat each job as a weighted basket of components—in this case, work activities or tasks. Like Frey and Osborne (2017). I theorize about jobs' degree of automation risk as a function of known social, physical, and motivational bottlenecks. My approach differs from that of Felten, Raj, & Seamans (2021) in that I use a theoretically guided approach to quantifying task automation vulnerability. I depart from Frey and Osborne (2017) in that I characterize a job's automation risk flexibly, based on automatability of its component activities, and I account for the creative capacity of the new generation of AI, defined by generative large language models. In this way, my estimates anticipate the effects of automation in domains that are not defined by social perspective taking, flexible manual work, and motivated vision and strategy.

Empirical Approach

I employ occupational descriptors from the U.S. Department of Labor's Occupational Information Network (O*NET), including importance and level scores for 41 work activities across 867 jobs. To capture the level and importance of each of 41 work activities (updated in 2023), I use 8-digit job codes from the Bureau of Labor Statistics' 2018 Standard Occupational Classification (SOC) system. To measure the number of workers in each occupation and their median salaries in 2022, I use 749 unique 6-digit BLS SOC codes, which is the precision level at which occupational size and salary data are available.

Activity rating

For each of the 41 job activities in the O*NET content model, I calculate a score based on three dimensions adapted from Frey and Osborne (2017) and Autor, Levy, and Murnane (2003). Specifically, I rate each task dichotomously on three dimensions:

Theory of Mind: Does the activity ask the worker to theorize about or anticipate what other people are thinking and feeling? Examples would include teaching, mentoring, nursing, and selling.

Flexible Dexterity: Does the activity ask the worker to perform manual tasks that are intricate and unpredictable? Examples would include plumbing, HVAC installation and repair, painting a house, replacing pipes or electrical wiring.

Vision and Strategy: Does the activity ask the worker to create intrinsic goals and objectives, or a plan for achieving them? Examples would include setting an organizational vision, creating a marketing plan, deciding on product launches and phase outs.

Interest category distribution

Because I am concerned with what secondary and postsecondary educators communicate to students about automation risk, I also examine the distribution of automation risk and salary changes by six categories of career interest.

The U.S. Department of Labor bases its career recommendations tools, such as *CareerOneStop* and *My Next Move*, on a career interest model known by the acronym RIASEC (Holland, 1959). The acronym stands for six categories of career interest:

Realistic: practical and hands-on, including physical, mechanical, natural, and outdoor work

Investigative: analytic, scientific, logical, and precision-oriented

Artistic: creative and expressive with images, words, music, objects, food, or other media

Social: understanding, helping, supporting, and engaging with people

Enterprising: involving business, strategy, leadership, negotiation, and competition

Conventional: using systems and standards to manage information, data, and materials

O*NET's RIASEC rates jobs based on the RIASEC category with which they have the highest correspondence, which allows me to examine the distribution of projected automation risk according to the RIASEC category most associated with each job. In Appendix A, I provide bar graphs showing the 2022 prevalence and salary distribution of jobs in each RIASEC category, based on 749 six-digit SOC codes.

Calculating job automation risk scenarios

Because technology is changing rapidly, assumptions about which job tasks can be automated are being constantly challenged. In 2017, Frey and Osborne's analysis of job automation risk assumed that jobs requiring "complex perception" and "creative intelligence" would remain difficult to automate within the next twenty years (p. 27). This assumption has been challenged by recent progress in self-driving technology and generative LLMs.

Even so, it remains true that not all job tasks can be easily automated. To help firms, educators, and workers cultivate skills for which demand is likely to persist, we need a framework for predicting generational automation risk in the age of generative AI.

In this paper, I combine insights from two different approaches to automation projections. Frey and Osborne (2017) adapt the framework from Autor, Levy, & Murnane (2003), in which difficulty of automation is determined by how non-routine a job is and by how cognitively demanding it is. Frey and Osborne's framework does recognize the pattern-recognition capabilities of AI, and it rates jobs according to three automation bottlenecks: perception and manipulation tasks, creative intelligence tasks, and social intelligence tasks. The authors use a Gaussian process, analogous to logit, to classify jobs' probability of full automation based on levels of nine U.S. Department of Labor O*NET ability variables pertaining to these three bottlenecks. I extend their approach by focusing on both level and importance scores for all 41 O*NET job activities (in lieu of abilities), and by updating the three bottlenecks to consider theoretical lessons that have emerged from the new generation of generative AI.

The other automation projection approach I build from is that of Felten, Raj, & Seamans (2021), who use O*NET job abilities to gauge likelihood of AI occupational exposure, which they call AIOE, scaling by the importance and prevalence of each ability in a particular job. Rather than adapting the idea of bottlenecks, they use crowd-sourced survey data from the public to gauge the AI exposure level of occupations, industries, and geographic regions. Felten et al. (2021) focus exposure at the level of 52 O*NET abilities (e.g., Fluency of Ideas, Hearing Sensitivity, Stamina), weighted by job importance and prevalence. They anticipate higher AI exposure for skills that are more cognitive, but my approach differs in that I adjust for importance-weighted levels of task sophistication using O*NET activities. I focus on automation risk for activities instead of abilities because activities are the tasks that comprise the job, making them easier to characterize in terms of automation bottlenecks. The challenge of the AIOE approach is that many of the AI exposure ratings seem implausible, with human services professions like teachers and mental health counselors, and complex strategic jobs like chief executives, physicists, and astronomers having among the highest ratings. This makes sense given that all these fields may use AI as complements for their work. On the other hand, a recent, national survey by Bick, Blandin, and Deming (2024) suggests that 39% of U.S. workers, including 22% of blue collar workers, already report using AI applications at work.

My approach rates 41 O*NET work activities by their automation bottlenecks and scales these ratings by their relative importance and level of sophistication in each job. However, my manner of classifying the automation bottlenecks in various jobs is based not on field-specific technological capabilities, as in Frey and Osborne (2017) nor on crowd-sourced views of the exposure of job abilities to AI, as in Felton, Raj, and Seamans (2021). Rather, it is based on a theoretical understanding, illuminated by generative AI and previous automation waves, of the human skills that remain hard to automate across jobs. It is also based on the notion that whether automation serves as a complement or substitute to a given job depends greatly on the level of sophistication at which the tasks vulnerable to automation are being performed. The ability to set and direct tasks toward goals (guided by vision/strategy) and with human rationales (informed by theory of mind) is what makes automation a complement to rather than a substitute for the task. Thus, by scaling automation risk by task level, I obtain a nuanced calibration of a job's likely automation risk. This also allows me to project salary distributional effects under scenarios that facilitate higher and lower degrees of automation complementarity for lower-skilled workers. This also lets me examine how jobs situated in RIASEC interest categories are differentially affected by automation risk, and how this affects their projected earning distributions.

I first calculate two scenarios for task automation risk, which I then weight by tasks' relative importance and level when scaling up to the level of jobs.

Scenario A

For my first measure of task safety from automation, I give equal weight to each of three plausible automation bottlenecks. For each of 41 work activities, as I assign a dichotomous 1/0 value for whether the activity requires some degree of theory of mind (*tomi*), flexible dexterity (*fdex*), and vision/strategy (*visi*). The sum of these dichotomous variables divided by 3 constitutes $task\_safe\_a_t$, which ranges from 0 to 1, with possible values of 0, 0.33, 0.66, and 1.0, as in equation (1). Subscript t indexes work activities, also known as tasks:

$$task\_safe\_a_t=(tomi_t+fdex_t+visi_t)/3 \qquad (1)$$

For example, a task that has only one of these elements is considered 33% safe and 66% potentially automatable, and a task with all three elements is considered unlikely to be automated.

Scenario B

As Felten, Raj, and Seamans (2021) have noted, nonroutine physical tasks may face higher automation barriers than tasks involving theory of mind or vision/strategy, due to the costs of building hardware with sensors and complex moving parts. In addition, the cost of automating any task may be nontrivial if it requires workers to learn new skills and techniques. To capture the effects of overall automation friction *and* of the additional difficulty of automating flexible dexterity tasks, I construct an alternative measure of task safety, as shown in equation (2):

$$task\_safe\_b_t = 0.2+0.2*tomi_t+0.4*fdex_t+0.2*visi_t \qquad (2)$$

where overall automation friction is represented by a constant of 0.2, and where the bottleneck weight for flexible dexterity, at 0.4, is twice that of theory of mind or vision/strategy.

Job automation risk

For each scenario, A and B, I aggregate all 41 values to the job level, j, weighting each task by its relative importance. Relative importance is the task's absolute importance to the job as a percentage of the total importance of all tasks for that job. This scaling is necessary because total importance scores across all tasks are higher in some jobs than in others. Relative importance of task t to job j is therefore defined as the task's importance score, ranging from 0 to 100, divided by the sum of all importance scores for the job, as in (3):

$$rel\_imp_{jt} = (imp_{jt} \,/\, \Sigma_{t=1\text{-}41}\, imp_{jt})\ *100 \qquad (3)$$

The resulting variable, rel_imp_{jt}, for tasks' relative importance within a job ranges from 0 to 9.3 with a mean of 2.4.

To estimate the fraction of work in a job that is vulnerable to automation, I separately scale the complement of $task\_safe_t$ (scenarios A and B separately) by its relative importance to the job, rel_imp_{jt}. I scale this by complement of the level of sophistication at which the task t is conducted in job j. Level is divided by 100 so that it ranges from 0 to 1. Multiplying a task's importance-weighted safety score by the complement of its level creates a negative, linear relationship between task sophistication and automation risk, reflecting the difficulty of automating complex tasks. This important assumption allows automation risk to fall as task sophistication rises, independent of the relative importance of each task to a job.
The level-and-importance-weighted task automation risk would therefore be 0 for a task at the highest sophistication level of 1, and it would be 1-*task_safe*, weighted by importance, for a task that is executed at level 0.

As shown in equation (4), aggregating these calculations to the job level yields the job's fractional automation risk, $jobaut_j$, indicating the fraction of the job that is likely to be vulnerable to automation within a generation:

$$jobaut_j = \Sigma_{t=1\text{-}41} \{(1\text{-}task\_safe_t)*rel\_imp_{jt}*[1\text{-}(task\_level_{jt}/100)]\} \qquad (4)$$

Similarly, I define the aggregate level at which work is performed in a job as

$$joblevel_j = \Sigma_{t=1\text{-}41} (rel\_imp_{jt} * task\_level_{jt}) \qquad (5)$$

The resulting variable, $joblevel_j$, ranges from 27.9 to 74.5, with a mean of 53.2.

To consider the consequences of job automation risk for worker salaries, we must speculate about the extent to which workers will use automation as a complement their work, shifting the demand curve upward, versus the extent to which automation will substitute for their labor, shifting downward the demand curve for their work. Recall that the story of complementarity is favored by Autor (2024) and observed by Chui et al. (2021), whereas Acemoglu and Restrepo (2020) see broad potential for worker substitution in the absence of policy interventions.

I consider three stylized scenarios for AI complementarity and substitution, as described in Table 1. Note that whereas $jobaut_j$ in column 1 refers to the percentage of jobs tasks that can be automated, $sectaut_s$ in column 2 refers to is job automation risk aggregated across all jobs in a SOC major code, or sector, of which there are 22 in the O*NET data. I construct sector risk by adding the $jobaut_j$ values for all the distinct jobs in a sector, and dividing by the number of unique jobs in the sector.

For the median salary of each job, $medsal_j$, I apply the following formula to create three projected salaries, $projsal_{cj}$, where c indexes scenarios 1-3 as described in Table 1, and j indexes jobs. The formula for the salary projections is given by:

$$projsal_{cj} = medsal_j * \{1+[k*(aut/100)]\} \qquad (6)$$

where k is a multiplier described in Table 1, and where aut is the automation risk variable aggregated at the job level, as in $jobaut_j$, or at sector level s, as in $sectaut_s$.

Table 1. Three scenarios for salary distribution effects of AI by job level quartile

Job level quartile	**Scenario 1:** Disproportionate benefits based on job risk	**Scenario 2:** Disproportionate benefits based on sector risk	**Scenario 3:** Small proportional benefits to all
Top	k = 0.2; *aut*=jobaut$_j$	k = 0.2; *aut*=sectaut$_s$	
Second	k = 0.1; *aut*=jobaut$_j$	k = 0.1; *aut*=sectaut$_s$	k = 0.05; *aut*=jobaut$_j$
Third	k = -0.1; *aut*=jobaut$_j$	k = -0.1; *aut*=sectaut$_s$	
Lowest	k = -0.2; *aut*=jobaut$_j$	k = -0.2; *aut*=sectaut$_s$	

Note: Sectors are defined by SOC major codes

This formula converts automation risk from a percentage to a fraction and scales it by variable k, which is set to 10% or 20% in scenarios 1 and 2 and to 5% in scenario 3, based on the assumption that only a modest percentage of productivity changes would be absorbed into workers' market value, with the remainder likely going toward profit, adaptation costs, and so forth. In scenarios 1 and 2, this 10-20% of task automation is added to workers' salaries in the top two quartiles of job level, and it is subtracted from their salaries in the lower two quartiles.

Scenario 2 differs from scenario 1 in that it is not the job's automation risk that is capitalized into salaries, but that of the entire sector. Scenario 2 thus reflects the notion that work is interdependent, and that managers may benefit automation of jobs that are both more and less vulnerable to automation than their own job.

Scenario 3, on the other hand, envisions that people will benefit equally from the automation of their own job tasks, albeit to a small degree. The multiplier of 0.05 across job level quartiles in scenario 3 captures the notion that everyone is able to use automation as a complement their job tasks, increasing their market value by about 5%. This implies that workers have some control over how they use automation, and that they are able to boost their own market value by usefully leveraging automation.

To document the effects of these proposed changes, I examine how the slope of the relationship between a job's median salary ($medsal_j$) and its aggregate level of complexity ($joblev_j$) would change under each of these three scenarios, and under the two different sets of automation bottleneck assumptions described in $task\_safe_t$ scenarios A and B above. I examine these projected changes overall and by RIASEC interest category.

Findings

First, I consider the projected automation risk across 873 jobs under bottleneck scenarios A and B. Under bottleneck scenario A, in which theory of mind, flexible dexterity, and vision/strategy are equally weighted task bottlenecks, the average job has an automation risk of 32.7, with a range from 17.4 to 49.3. This means that 32.7% of tasks in the average job are vulnerable to automation.

In scenario B, the average job has an automation risk of 27.4, with a range from 14.4 to 41.2, meaning that 27.4 of tasks in the average job are presumed to be at risk. The automation risk scores are lower in scenario B because scenario B assumes that 20% of all tasks are safe due to automation friction. Scenario B also assumes that flexible dexterity tasks are twice as hard to automate as theory of mind and vision/strategy tasks. The full distribution of projected automation risks for 873 jobs under each scenario is shown in Figure 1.

Figure 1. Distribution of job automation risks under bottleneck scenarios A and B

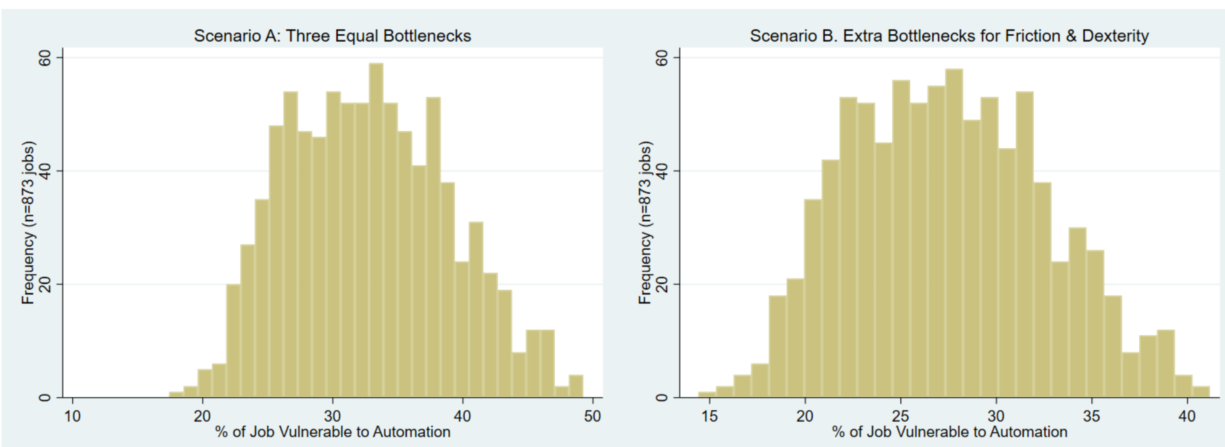

Next, I consider how automation risk is projected to affect workers at different levels of education and in different RIASEC categories of interest. Given that the automation risk scores are constructed to decline as jobs' complexity level increases, we would anticipate that it would covary positively with education level, but the extent of such covariance is not clear a priori. We would also anticipate greater covariance with education level in scenario A than in scenario B, given that manual tasks are considered especially difficult to automate in scenario B.

In Table 2, I show that job automation risk declines almost linearly with the typical educational level of each O*NET job zone. Projecting a tighter bottleneck for flexible dexterity in scenario B does not change the pattern, except that it slightly compresses the mean differences between levels, so that the difference between less than high school and advanced is 13 points under scenario A and only 11 under scenario B. It is also noteworthy that jobs requiring less education show higher standard deviations in automation risks, meaning their risks are more varied than jobs that require higher levels of education.

Table 2. Job automation risk by typical education levels of O*NET job zones

	Scenario A		**Scenario B**		
	mean	sd	mean	sd	n
<HS	39.7	5.4	33.2	4.5	32
HS	37.5	5.0	31.3	4.2	279
Assoc	33.3	4.5	27.8	3.8	207
Bach	29.3	4.2	24.6	3.5	203
Advan	26.5	3.7	22.4	3.1	152
Total	32.7	6.2	27.4	5.1	873

Table 3 presents means and standard deviations for automation risk by RIASEC interest category. Here again, estimates in scenario B are about five points smaller by design, due to the bottleneck for automation friction in scenario B, but they are otherwise similar in their comparative automation risks. The categories with the highest risks of automation, around 35 under scenario A and 29 under scenario B, are realistic (manual and outdoor), artistic, and conventional. It is striking that this finding holds even with the higher protection for flexible dexterity in scenario B.

Table 3. Job automation risk by RIASEC categories

	Scenario A		Scenario B		
	mean	sd	mean	sd	n
Realistic	35.2	5.4	29.3	4.5	379
Investigative	25.9	3.7	21.8	3.1	103
Artistic	35.1	5.2	29.6	4.3	27
Social	30.0	4.8	25.3	4.1	125
Enterprising	30.2	5.3	25.4	4.5	90
Conventional	34.6	6.2	29.1	5.2	149
Total	32.7	6.2	27.4	5.1	873

In Table 4, I present average automation risk estimates under scenarios A and B for 22 work sectors, where the sectors are defined by the two-digit major codes for SOC job classifications. As noted, and by design, automation risks are about five points lower in scenario B than scenario A. But it is notable that the rankings of sectors in terms of automation risks are unchanged between the two scenarios. This occurs even though scenario B doubles the bottleneck for flexible dexterity, thus potentially offering extra protection for workers in manual occupations. In both scenarios in Table 4, automation risks are highest for many manual and routine-oriented sectors including food preparation, production, food service, and office support. Automation risks are mid-range in a wide range of fields including construction, installation, arts, and business. And it is estimated to be relatively low in jobs that require high levels of expertise, including architecture, math and science jobs, and management. Human services like education, social services, and healthcare also fare well in Table 4 because their tasks involve theory of mind as well as the use of vision/strategy to achieve goals.

Automation risk by salaries

To further project the relationship between job automation risk and the distribution of earnings, I graph the risk of automation under scenarios A and B against median salaries for each job in 2022. In Figure 2, I display fitted trend lines of the relationship between automation risk and median salary, showing bottleneck scenario A on the left and scenario B on the right. The scatterplot markers indicate the RIASEC category of each job, making it easy to see that automation risk is highest in realistic and

Table 4. Estimated mean automation risks by 22 SOC major codes

Sector	Scen A	Scen B
Food Preparation	40	33.5
Office and Admin Support	39.3	33.1
Personal Care	38.5	32.3
Production Occupations	37.8	31.5
Farming, Fishing, Forestry	37.7	31.3
Building and Grounds	36.9	30.7
Sales and Related	36.5	30.7
Transportation/Materials Moving	35.6	29.6
Construction and Maintenance	34.9	29.1
Installation, Maintain, Repair	34.7	28.9
Arts, Design, Entertain, Media	34.1	28.7
Healthcare Support	33	27.6
Legal Occupations	30.9	26.3
Business and Finance	30.3	25.7
Protective Services	29.8	24.9
Educational/Library	29.2	24.7
Community/Social Services	29	24.5
Architecture and Design	28.5	23.9
Healthcare Practitioners	28.2	23.6
Computer and Math	27.6	23.4
Life, Physical, Social Scie	27.4	23.1
Management Occupations	27.4	22.9
Total	32.7	27.4

conventional jobs, which also have the lowest pay, and it is lowest in investigative and social jobs, which also show some of the highest salaries. However, because automation risks are higher in scenario A, the relationship is somewhat flatter in scenario B. The patterns in both scenarios suggest that automation patterns are likely to increase income inequality. However, a scenario with general protection through automation friction and particular protection for flexible dexterity suggests that disparities in job loss between higher- and lower-paid workers may be not quite as severe as in the equal bottleneck scenario.

Figure 2. Job automation risk by median salaries and RIASEC categories (n=867 jobs)

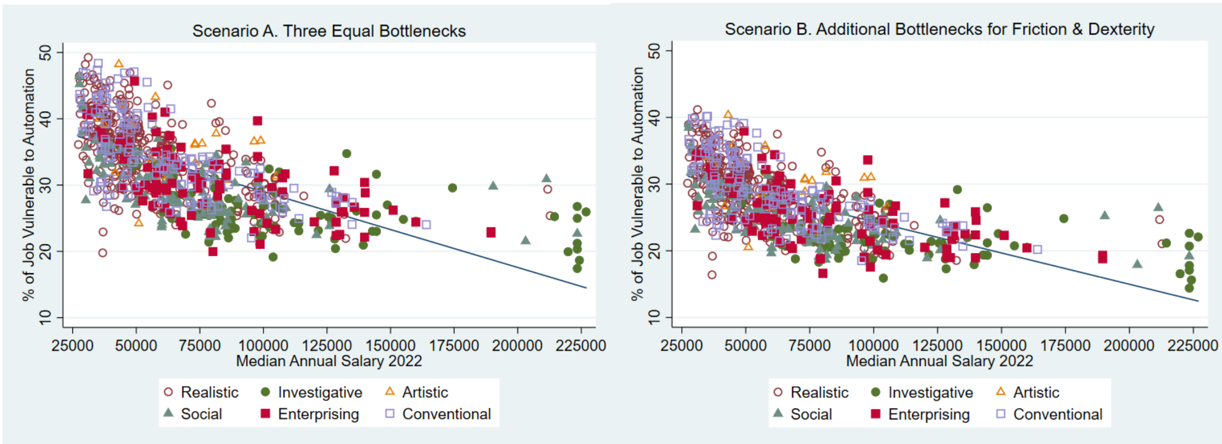

Salary distributional changes by job level

Finally, I consider possible consequences of automation risk for shifts in jobs' median salaries as automation is capitalized into the market value of people's work. Here, I present results for salaries under the distributional assumptions of Table 1. Recall that scenario 1 redistributes 10-20% of job automation risk differentially by job level quartiles within sector. Scenario 2 redistributes 10-20% of *sector* automation risk by job level quartiles within sector. Scenario 3, in contrast, capitalizes 5% of job automation risk into all salaries.

Table 5 captures projected changes in median salaries under each scenario with the equal and additional bottleneck assumptions in scenarios A and B, respectively. Median salaries would rise 3.5% and 3.7%, respectively under the redistribution of sector-level automation risks in scenarios 2A and 2B. This is because those at higher job levels, who have lower automation risks by construction, can capture higher sector-level average risks. They are projected to rise 2% in scenarios 1A and 1B. They would rise least, at 1.3% and 1.4%, respectively, in the equal-capitalization scenarios of 3A and 3B. The different assumptions about bottleneck construction in A and B make very little difference for average projected salary changes.

Table 5. Projected changes in median salaries under automation and distribution assumptions

	Mean	sd	Min	Max	Δ from original
2022 Median	66,801	34,526	27,270	226,880	
A: Equal Bottlenecks					
Scenario 1A	68,114	39,173	23,786	260,745	2.0%
Scenario 2A	69,113	35,889	28,141	235,279	3.5%
Scenario 3A	67,702	38,932	23,362	257,154	1.3%
B: Additional Friction and Dexterity Bottlenecks					
Scenario 1B	68,109	39,477	23,459	262,106	2.0%
Scenario 2B	69,279	35,945	28,223	235,720	3.7%
Scenario 3B	67,768	39,274	23,103	259,075	1.4%

However, my primary interest in the salary projections is how they would affect salary inequity. In Figure 3, I examine how the slope of the bivariate relationship between salaries and job levels would be expected to change in each scenario in Table 5.

Figure 3. Projected changes in slopes of salaries by job level across six scenarios

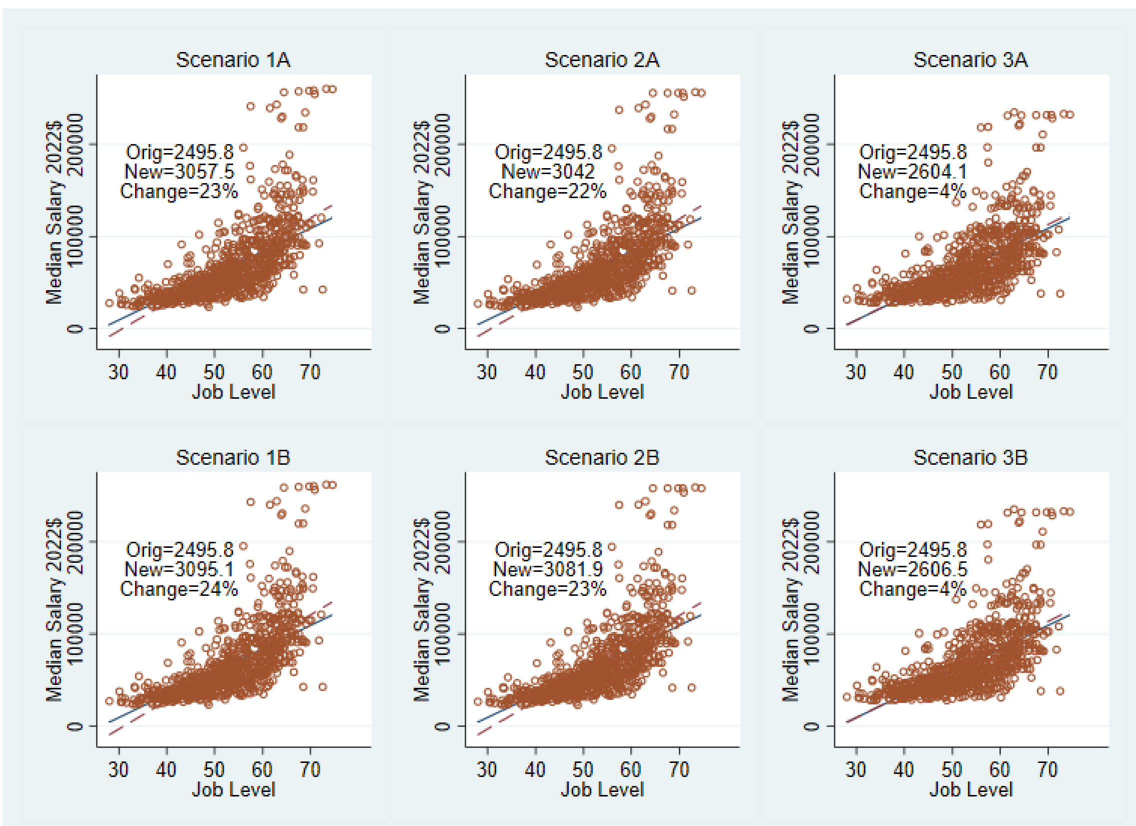

In Figure 3, I find that the disparities in salaries by job level would rise most dramatically, by 23-24%, in the scenario where job-level rather than sector-level automation is capitalized differentially into salaries. This is shown at left in scenarios 1A and 1B. It is unexpected that inequality rises more in the bottom-row B scenarios than the top-row A scenarios, since the B scenarios include greater protections for flexible dexterity and for automation friction overall.

Inequality still rises in scenarios 3A and 3B, where salaries rise by 5% across the board, because a flat percentage change boosts higher salaries more than lower ones in absolute terms. Even so, scenarios 3A and 3B show much less change in salary inequality than the scenarios in which salaries are differentially affected by job level quartile. This suggests that supporting workers in benefitting from their own automation efficiencies may greatly moderate the tendency of AI-enabled automation to exacerbate income inequality. It also shows that inequality is likely to increase even under the most optimistic scenarios.

A related question is how salary inequality would change by RIASEC interest categories. Having this information may allow people to make better-informed decisions about skill acquisition in different career categories, and it may allow educators to better advise and support students. To address this question, I present Figures 4 through 6, in which the relationship between current and projected median salaries are graphed against job levels separately by RIASEC categories. I show anticipated changes in these slopes by RIASEC category for three of the six scenarios above. Figure 4 presents the change in slopes for scenario 1A, with differential capitalization of job automation risk by sector. Figure 5 presents the slope change for scenario 2B, with differential capitalization of sector automation risk by sector. Figure 6 shows the change in slopes for scenario 3B, with uniform 5% capitalization of sector automation risk. Figure 4 assume equal automation bottlenecks, whereas Figures 5 and 6 assume additional bottlenecks for automation friction and flexible dexterity. I present these three scenarios by RIASEC categories because they capture notable variation among the six scenarios in Figure 3.

Figure 4. Scenario 1A: salary-by-job level slope projection by RIASEC category

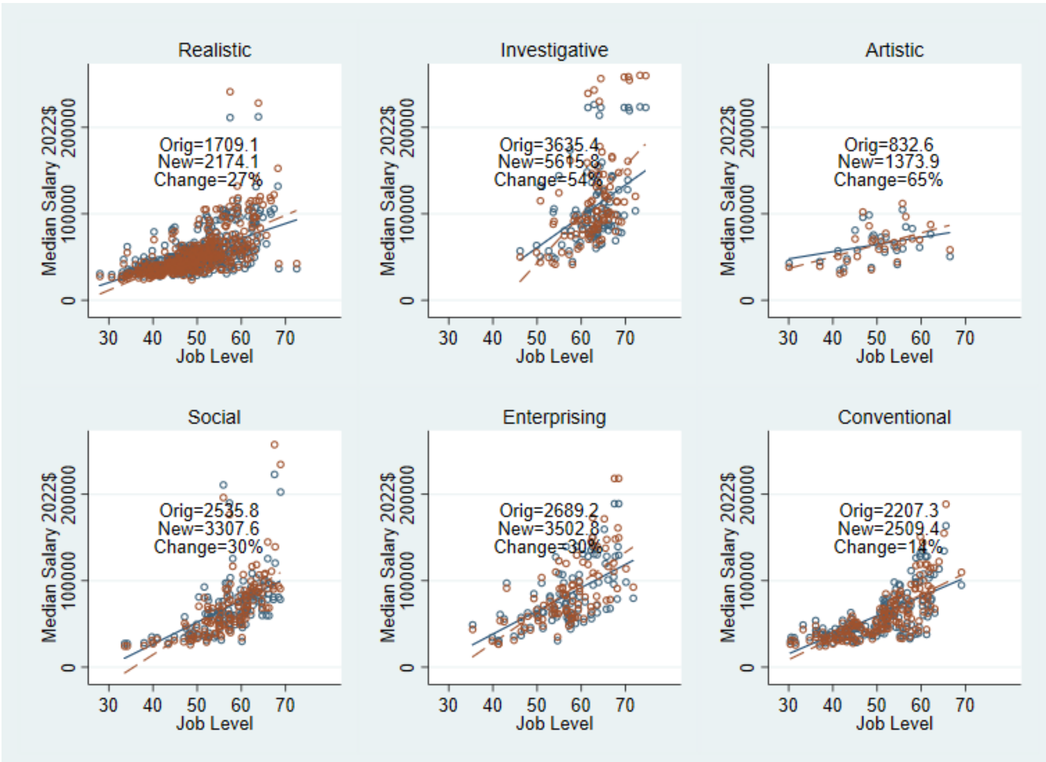

Figure 5. Scenario 2B: salary-by-job level slope projection by RIASEC category

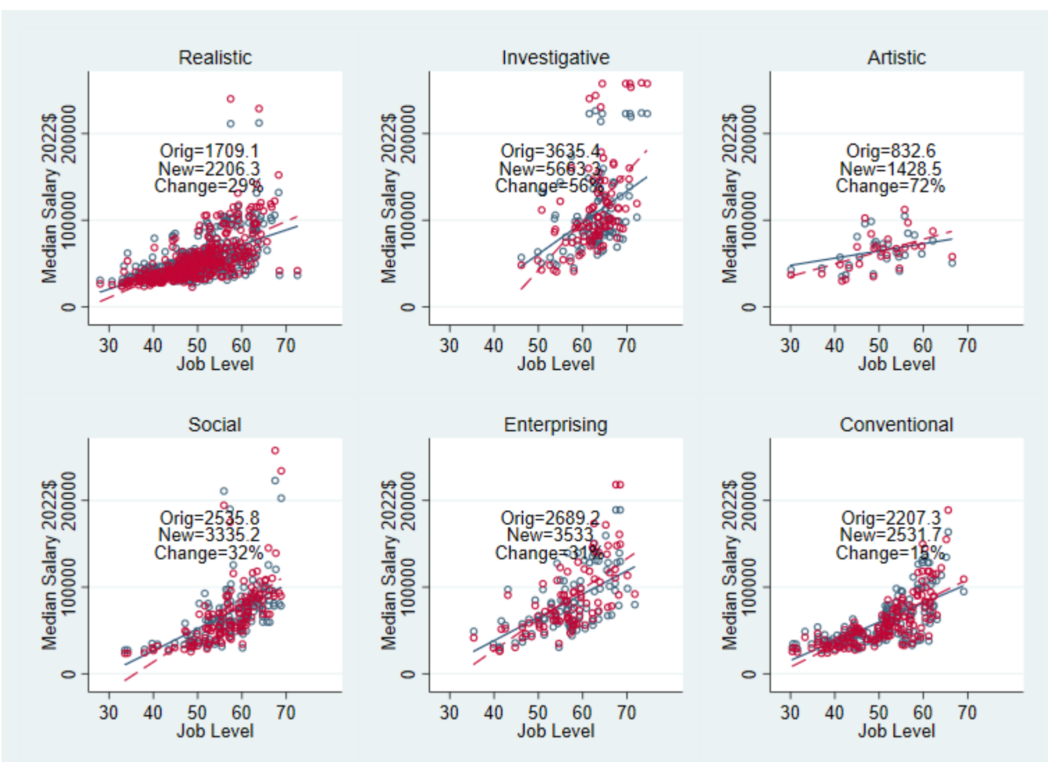

Figure 6. Scenario 3B: salary-by-job level slope projection by RIASEC category

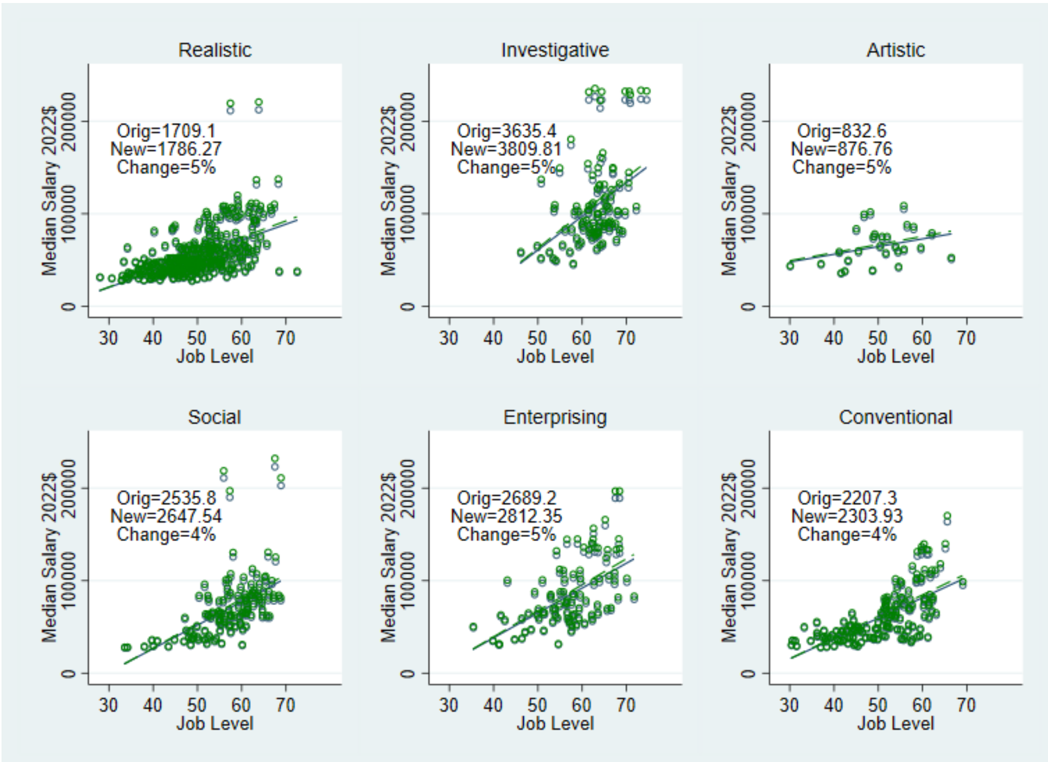

In Figure 4, under scenario 1A, we see that the slopes change at different rates across categories, from a small 11% change in conventional jobs, to a 54% increase in slopes in investigative jobs and 65% increase in artistic jobs. This is notable, as the baseline slope is much flatter in artistic jobs than in investigative jobs, and because I anticipate artistic jobs to be much more vulnerable to automation than investigative jobs, as noted above.

Turning to Figure 5, which assumes additional automation bottlenecks as well as salary redistribution by sector instead of by job, I find patterns very similar to those in Figure 4, but with slightly higher slope increases in all categories except conventional jobs. I now predict a slope change as high as 72% in artistic careers. This suggests that salary inequality within artistic jobs has heightened propensity to rise, given the capacity of generative AI to create original content.

In Figure 6, which assumes additional automation bottlenecks and uniform 5% capitalization of automation risks, the projected relationship of salaries to job level rises by only 4-5% in all RIASEC categories, including the investigative and artistic categories in which it spiked under the redistributive scenarios.

Conclusion

A pressing question with the rise of AI is how its potential to automate white collar analytic and creative tasks will affect the distribution of jobs and salaries in the labor market. The debate is whether AI will largely serve as a complement to existing work, enabling greater productivity and prosperity, or whether it will slowly replace human workers as their baskets of core tasks are automated. My analysis builds on existing papers by Frey and Osborne (2017) and Felton, Raj, and Seamans (2021) but moves beyond them in a few ways. First, I treat jobs as baskets of 41 concrete work activities, whose importance and levels of sophistication have been rated by incumbents in the jobs. I adapt and update Frey and Osborne's (2017) framework of automation bottlenecks as previously developed by Autor, Levy, and Murnane (2003), using a theoretical perspective on tasks that are difficult to automate. I test two different versions of this automation theory—one in which all three bottlenecks are equal, and the other in which flexible dexterity is twice as difficult as theory of mind and vision strategy, and in which all jobs are also protected by frictions in automation, such as tool/task misalignment and worker learning curves. I provide refined estimates of job automation risk by scaling tasks' relative importance by the level of sophistication at which they are carried out. Assuming that the difficulty of automation rises with task complexity, I am able to project the fraction of a job's tasks that are vulnerable to automation and to assess how this vulnerability varies by the jobs' aggregate level of sophistication. I am also able to project possible changes in salary inequality by considering different scenarios for who will use AI as a complement versus a substitute, and to what extent.

I find that differential weighting of automation bottlenecks matters very little in terms of the distribution of automation risks or salary distributions. What matters instead are assumptions about who wins and who loses from automation vulnerability within a sector, and to what extent. I show that when workers at all levels are able to take advantage of their own automation efficiencies, salary inequality may still rise, but by very little.

A key takeaway is that even if people capture the productivity and pay benefits of automation in small-but-equal proportions, returns to expertise will still increase. And if, as seems likely, returns to automation are not equally captured, then returns to expertise will increase more, and the bottom part of the expertise distribution will lose ground. In essence, the problem is not automation *per se.* The problem is heterogeneity in who controls and benefits from the resulting efficiencies.

My projections about the likely share of tasks that can be automated in the next generation appear plausible, with ranges from 17% to 49% of tasks across jobs in the equal bottleneck scenario and 14% to 41% of tasks in the more-conservative scenario. Based on prior research on automation risk (D. H. Autor et al., 2003), and on documented experiences with current AI (Ji et al., 2022; Lohr, 2024; Toner-Rodgers, 2024), it is also likely that complex tasks will remain harder to automate than simple tasks, which is the assumption I use in scaling automation risk by task sophistication.

Implications for educators

It is important to note that my projections about salary distributional changes are modest, assuming that only between 5% and 20% of automation risks are capitalized into salaries. The reality could of course be much higher. In addition, the beneficiaries of automation could be fewer than the top half or top quartile of the job-level distribution. This is why it is so important that educators prepare students to leverage AI so they are equipped to capture and utilize their own automation gains.

Educators may wish to steer students toward interest categories that are less vulnerable to automation, including investigative, social, and entrepreneurial tasks. This recommendation is consistent with recent studies of students' performance on science tests, in which the comparative advantage of children over AI lies in critical thinking and the application of human knowledge to novel scenarios (Zhai et al., 2024).

In addition, educators should work to incorporate hard-to-automate skills like theory of mind, vision/strategy, and even flexible dexterity into courses in a wide array of disciplines. This will allow students to cultivate adaptive, hard-to-automate skills that they can apply across many different careers.

Bibliography

Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2022). AI and Jobs: Evidence from Online Vacancies. *Journal of Labor Economics, 40*(S1). https://doi.org/https://doi.org/10.1086/718327

Acemoglu, D., & Johnson, S. (2023). *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*. Public Affairs. https://www.google.com/books/edition/Power_and_Progress/BV2HEAAAQBAJ?hl=en&gbpv=1

Autor, D. (2024). Applying AI to rebuild middle class jobs. *NBER Working Paper Series*, *w32140*.

Autor, D. H. (2014). Skills, education, and the rise of earnings inequality among the "other 99 percent." *Science*, *344*(6186), 843–851. https://doi.org/10.1126/science.1251868

Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, *118*(4), 1279–1333. https://doi.org/https://doi.org/10.1162/003355303322552801

Bick, A., Blandin, A., & Deming, D. J. (2024). *The rapid adoption of generative AI* (No. 32966; NBER Working Paper). https://doi.org/10.3386/w32966

Brynjolfsson, E., Li, D., & Raymond, L. (2023). Generative Ai at Work. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4426942

Carnevale, A. P., & Strohl, J. (2010). How increasing college access is increasing inequality, and what to do about it. In R. D. Kahlenberg (Ed.), *Rewarding strivers: Helping low-income students succeed in college* (pp. 71–183). The Century Foundation. https://vtechworks.lib.vt.edu/bitstream/handle/10919/83054/IncreasingCollegeAccess.pdf?sequence=1&isAllowed=y

Chui, M., Hall, B., Singla, A., & Sukharevsky, A. (2021). The State of AI in 2021. In *McKinsey & Company* (Issue December 8). https://www.mckinsey.com/capabilities/quantumblack/our-insights/global-survey-the-state-of-ai-in-2021

Deming, D. (2023). Multidimensional Human Capital and the Wage Structure. *SSRN Electronic Journal*, *HKS Workin*, 1–58. https://doi.org/10.2139/ssrn.4383454

Dziri, N., Milton, S., Yu, M., Zaiane, O., & Reddy, S. (2022). On the Origin of Hallucinations in Conversational Models: Is it the Datasets or the Models? *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, *July 10-15*, 5271–5285. https://github.com/McGill-NLP/FaithDial

Felten, E., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, *42*(12), 2195–2217. https://doi.org/10.1002/smj.3286

Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, *114*, 254–280. https://doi.org/https://doi.org/10.1016/j.techfore.2016.08.019

Goldin, C., & Katz, L. F. (2008). *The race between education and technology*. Belknap Press.

Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y., Dai, W., Madotto, A., & Fung, P. (2022). Survey of Hallucination in Natural Language Generation; Survey of Hallucination in Natural Language Generation. *ArXiv*, *Feb.* https://arxiv.org/pdf/2202.03629.pdf

Lee, T. B., & Trott, S. (2023, July 27). Large language models, explained with a minimum of math and jargon. *Understanding AI*. https://www.understandingai.org/p/large-language-models-explained-with

Lohr, S. (2024, July 23). A.I. can write poetry, but it struggles with math. *The New York Times*. https://www.nytimes.com/2024/07/23/technology/ai-chatbots-chatgpt-math.html

McElheran, K., Li, J. F., Brynjolfsson, E., Kroff, Z., Dinlersoz, E., Foster, L., & Zolas, N. (2024). AI adoption in America: Who, what, and where. *Journal of Economics and Management Strategy*, *33*(2). https://doi.org/10.1111/jems.12576

Piketty, T. (2018). Brahmin left vs. merchant right: Rising inequality and the changing structure of political conflict: (Evidence from France, Britian, and the U.S., 1948-2017). In *World Inequality Database Working Paper Series: Vol. 2018/7*. World Inequality Database Working Paper Series No. 2018/7. http://piketty.pse.ens.fr/files/Piketty2018.pdf

Toner-Rodgers, A. (2024). *Artificial intelligence, scientific discovery, and product innovation* (Nov. 6). https://aidantr.github.io/files/AI_innovation.pdf

Zhai, X., Nyaaba, M., & Ma, W. (2024). Can Generative AI and ChatGPT Outperform Humans on Cognitive-Demanding Problem-Solving Tasks in Science? *Science & Education*. https://doi.org/10.1007/s11191-024-00496-1

Appendix A

Figure A1. 2022 prevalence of jobs in each RIASEC category (n=749 six-digit SOC occupations)

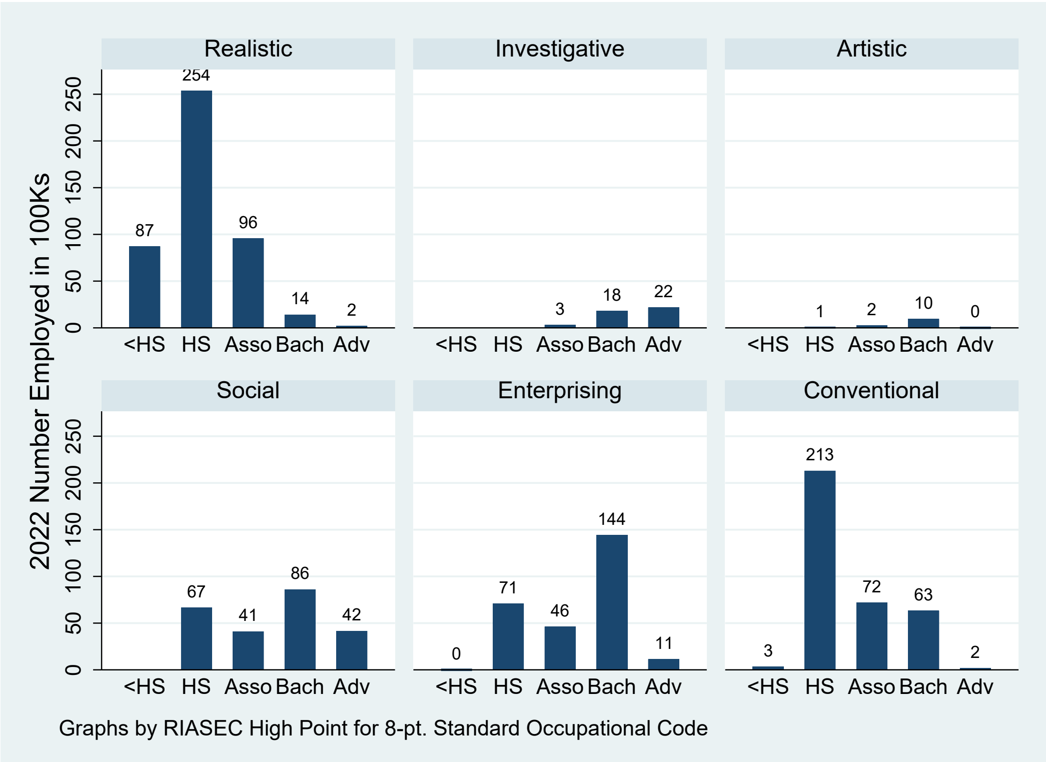

Figure A2. 2022 median salaries of jobs in each RIASEC category (n=749 six-digit SOC occupations)