Discussion of "AI Exposure and the Adaptive Capacity of Workers in the US Labor Market"

Dimitris Papanikolaou
Kellogg School of Management and NBER

September 27, 2025

The authors introduce a measure that combines occupational exposure to AI with an index of worker characteristics that influence how severely individuals may be affected by potential displacement. Drawing on several national datasets, they construct an occupation-level index of adaptive capacity based on four variables: net liquid wealth, geographic density, age, and skill transferability. These factors are intended to capture an individual's capacity to absorb economic shocks from AI-driven changes in labor demand. A central result is that workers' capacity to adapt is strongly and positively correlated with AI exposure. This suggests that many of the occupations most exposed to AI tend to employ workers who are relatively well positioned to handle job transitions. At the same time, the authors identify a sizable group who face both high exposure and low adaptive capacity. These workers are concentrated in customer service, administrative support, and clerical jobs.

My main comment is that exposure of an occupation's tasks to AI is not the same as worker displacement. This distinction matters for two reasons. First, exposure does not necessarily translate into reduced labor demand, since offsetting forces may operate, as I discuss in Section 1. Second, it is important to distinguish occupation-level from worker-level outcomes. Workers can move out of declining occupations, though switching costs vary with age. Conversely, even when new technologies expand demand for an occupation, incumbent workers may still be displaced if they lack the skills required to adapt. I develop this point further in Section 2. Last, I have several secondary comments regarding the construction of the index, which I briefly discuss in Section 3

1 AI task exposure vs Occupation Labor Demand

My first point is that observing AI substituting for certain tasks within an occupation does not, on its own, imply that the overall demand for that occupation will decline. The challenge is that even when technologies are labor-saving at the task level, their aggregate impact on labor demand is ambiguous. Labor-saving technologies may reduce the need for workers in specific tasks, but they can also increase overall productivity, encourage reallocation toward complementary tasks, and ultimately expand labor demand in the aggregate economy.

Hampole, Papanikolaou, Schmidt, and Seegmiller (2025) formalize these ideas using a simple model that links task-specific technological improvements to occupational labor demand. Their framework captures both the direct substitution effect and the indirect productivity and reallocation effects. When a technology broadly substitutes for labor across all tasks of an occupation—as in the case of automatic telephone switching displacing operators—the outcome is a steep reduction in labor demand. In contrast, when a technology applies narrowly—for example, an automation tool that handles expense reporting for economists—the negative effect is attenuated, or even reversed. Because workers can reallocate time and effort across tasks, automation in one area frees up capacity in others and can generate productivity gains that boost demand both within the occupation and in related occupations.

Their model highlights three statistics that summarize the impact of a given technology on labor demand. First, the average exposure of an occupation's tasks to the new technology is typically negatively related to labor demand: if all tasks are even moderately affected, demand falls (as when the introduction of tractors reduced demand for farm labor). Second, the concentration of exposure matters: if technological advances are confined to a narrow set of tasks, workers can shift effort toward other activities, raising productivity and mitigating demand losses. Third, the magnitude of productivity improvements determines the extent to which labor demand increases in occupations that are indirectly affected.

Hampole et al. (2025) provide empirical evidence consistent with these mechanisms using a novel measure of AI task exposure that combines information from AI developers' resumes with task descriptions from O*NET. Their results point to three main findings. First, adoption is concentrated in large, productive firms and is associated with higher sales, profits, and TFP. Second, at the occupational level, higher-wage jobs tend to be more exposed to AI; while average exposure reduces employment, concentrated exposure facilitates reallocation toward complementary tasks and offsets these declines. Third, at the firm level, AI adoption raises overall employment, consistent with productivity-driven labor demand growth. Taken together, their analysis suggests that the positive reallocation and productivity effects can outweigh the direct substitution effects of AI. For instance, although AI exposure is greatest among high-wage occupations, these groups experience modest increases in relative labor demand following adoption.

In sum, an occupation's task exposure to AI is not a sufficient statistic for how labor demand for the occupation is likely to evolve in the future, even if AI substitutes for labor in individual tasks.

2 Occupation vs Worker Outcomes

My second point is that worker-level outcomes can differ markedly from occupation-level outcomes. There are two main reasons for this divergence. First, workers can transition out of affected occupations. For example, employees in occupations at high risk of replacement by AI—such as customer service agents—may move into other occupations that are less exposed. The feasibility of such transitions depends on the breadth of workers' skill portfolios. Those with versatile, transferable skills are more mobile and therefore less vulnerable, while workers with highly specialized skills may struggle to shift into new roles. Second, even when a technology raises labor demand within an occupation, incumbent workers may not benefit if they lack the necessary skills. For instance, the diffusion of computers increased labor productivity overall but disadvantaged workers unable to acquire the requisite digital skills. Age often serves as a useful proxy for adaptability, with younger workers typically better able to adjust to technological displacement than older cohorts.

Most studies in the United States focus on occupation-level outcomes. A notable exception is Kogan, Papanikolaou, Schmidt, and Seegmiller (2023), who provide granular evidence on how technology exposure influences the earnings of individual workers between 1980 and 2010. They distinguish between two types of technological advances: those related to routine tasks, which they interpret as largely labor-saving, and those related to non-routine tasks, which they view as primarily labor-augmenting. This classification rests on the idea that routine tasks are those that can be codified into explicit instructions and therefore lend themselves to automation via software, robotics, or other programmable capital (Autor, Levy, and Murnane, 2003). Non-routine tasks, by contrast, could not be fully performed by machines during this period. To build these measures, they use textual analysis of patent documents combined with detailed occupation task descriptions.

Their analysis is illustrative in contrasting occupation with worker-level outcomes. At the occupation level, they find that industries heavily exposed to labor-saving technologies experience reductions in the wage bill, driven by both lower employment and lower average wages. Conversely, industries more exposed to labor-augmenting technologies see their occupational wage bill rise, largely due to growth in employment.

However, their findings at the worker level do not directly mirror the occupation-level outcomes. On the one hand, they do find that labor-saving technological advances reduce the earnings of incumbent workers in exposed occupation–industry pairs. However, they also show that technology advances that increased occupation labor demand led to (modest) earnings declines for incumbent workers, whereas new entrants experience higher earnings. Digging deeper, Kogan et al. (2023) show that these earnings losses are concentrated among workers who are separated from their employers, often involuntarily. Increased job destruction—defined as switching employers while simultaneously suffering large earnings losses—accounts for more than half of the average displacement effects that they document.

Turning attention to heterogeneity, the negative effects of labor-saving technologies are widespread: they appear across services and manufacturing, manual and cognitive occupations, and among both college-educated and non-college-educated workers. Notably, the effects are not systematically related to workers' age or their relative wage position after controlling for age. This finding implies that switching costs that increase with age is not the whole story of what is going on in the data (see, e.g., Kambourov and Manovskii, 2009). It is possible that employers selectively lay off workers as a function of their seniority.

The response to labor-augmenting technologies displays considerably more heterogeneity across workers. While these technologies tend to generate smaller average earnings declines than labor-saving ones, they still reduce the earnings of many incumbent workers in exposed occupation–industry cells. By contrast, new entrants into these occupations often benefit: labor-augmenting technologies raise the earnings of entrants, offsetting some of the losses among incumbents. Focusing on heterogeneity within incumbents, Kogan et al. (2023) show that the earnings declines among incumbents are especially pronounced for older workers and for those earning relatively high wages compared to their peers within the same occupation and industry. For these groups, Kogan et al. (2023) find that job destruction—the combination of involuntary employer separations and substantial earnings losses—explains a large share of the adverse effects. The fact that relatively high-paid workers are disproportionately harmed suggests that occupation-specific skills, which may underpin their higher relative wages, are not easily transferable and leave these workers more vulnerable to technological change.

In brief, Kogan et al. (2023) show that worker outcomes can diverge sharply from occupation-level outcomes: even if new technologies increase labor demand for an occupation, individual workers may still be displaced if they lack the skills required to use those technologies. Importantly, this mechanism could partly offset the one described in the previous section. Suppose that AI-exposed occupations experience rising labor demand because AI enables workers to reallocate time toward tasks less affected by automation. Some of these tasks may be entirely new and, crucially, may require skills that incumbent workers do not possess. In such cases, we might observe the pattern documented by Kogan et al. (2023): incumbent workers in AI-exposed occupations are displaced, while younger entrants—who are better able to acquire or already possess the relevant skills—benefit from the increase in labor demand. Thus, even if AI exposure is not a sufficient statistic for changes in labor demand at the occupation level, it can still lead to displacement of incumbent workers even if labor demand for the affected occupations increases.

3 Minor Points

In addition to the points in the previous sections, I have a few smaller comments on the construction of the index. First, on the measurement of wealth: the authors exclude housing from their measure of wealth, arguing that it is illiquid and may limit geographic mobility (Fonseca and Liu, 2024). While this is a reasonable choice, housing is also the primary form of saving for many households (Poterba, Venti, and Wise, 2011). It would therefore be useful to consider an alternative version of the index that includes housing equity.

Second, on worker age: although younger workers may indeed find it easier to switch occupations, their human capital represents a much larger share of total wealth compared to workers approaching retirement. This distinction may matter for how vulnerability is measured across age groups.

Third, on aggregation: the weighting of the individual components should ideally reflect their role in mitigating displacement risk. The authors either average the normalized components or apply Principal Components Analysis (PCA), which effectively assigns weights based on comovement. A more informative approach would be to draw on existing empirical evidence about how these characteristics influence workers' ability to withstand displacement.

4 Summary

Overall, identifying which workers are best positioned to withstand shifts in labor demand is a first-order question, and this paper makes an important contribution in that direction. As AI continues to reshape the labor market, it is essential to understand the factors that drive worker displacement and to consider how these effects can be mitigated.

References

Autor, D. H., F. Levy, and R. J. Murnane (2003, 11). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics 118*(4), 1279–1333.

Fonseca, J. and L. Liu (2024). Mortgage lock-in, mobility, and labor reallocation. *The Journal of Finance 79*(6), 3729–3772.

Hampole, M., D. Papanikolaou, L. D. Schmidt, and B. Seegmiller (2025, February). Artificial intelligence and the labor market. Working Paper 33509, National Bureau of Economic Research.

Kambourov, G. and I. Manovskii (2009). Occupational specificity of human capital. *International Economic Review 50*(1), 63–115.

Kogan, L., D. Papanikolaou, L. D. Schmidt, and B. Seegmiller (2023, November). Technology and labor displacement: Evidence from linking patents with worker-level data. Working Paper 31846, National Bureau of Economic Research.

Poterba, J., S. Venti, and D. Wise (2011, Fall). The composition and drawdown of wealth in retirement. *Journal of Economic Perspectives 25*(4), 95–118.