

# Could Savannah be the next San Jose? The Downstream Effects of Large Language Models\*

**Scott Abrahams**

Louisiana State University  
sabrahams@lsu.edu

**Frank Levy**

Duke University (faculty affiliate),  
MIT (emeritus)  
flevey@mit.edu

January 22, 2025

## Abstract

The anticipated production shock created by large language models (LLMs) may significantly alter the geographic structure of U.S. labor demand. We use established estimates of LLM occupational exposure to develop a map of LLM impacts and find that urban, highly educated coastal metro areas are the most affected. We then consider the downstream effects as places and people adjust to changing employment opportunities. We identify three avenues of adjustment that emerged during the post-1980 decline in manufacturing employment, which we apply to the present situation. Combining our mapping with these avenues, we predict that displaced college graduates will migrate towards smaller, lower exposure urban centers including Rochester, New York and Savannah, Georgia, that demand for a four-year college education will fall, and that the migration patterns and politics of affected persons will dampen rather than exacerbate political polarization—provided that government can successfully moderate the pace of change.

**Keywords:** production shocks, manufacturing employment, geographic polarization, large language models, automation, downstream effects

**JEL Classification Numbers:** J61, N12, O14, O33, O51

---

\*We thank Daniel Rock for sharing data on occupation-level exposure, Emily Pollard for use of her historical occupational employment data and John Boyd, Jack Citrin, Wesley Cohen, Daniel Dougherty, Steve Gooze, Sharique Hassan, Brad Hershbein, Jeffrey Lin, and Theda Skocpol for discussion of various points. All errors and conclusions are our own.

# 1 Introduction

The widespread adoption of large language models (LLMs) into the firm production process may significantly alter the structure of labor demand in the United States. Economists predict that about 15% of all tasks currently performed by human workers could be completed by a generative software tool including OpenAI’s ChatGPT, and 19% of the total workforce is in an occupation where at least half of the tasks involved might be performed by LLMs (Eloundou et al. 2023). Although exposure does not indicate whether an LLM might replace a human worker or raise the worker’s productivity, either outcome implies that a given output could be produced with less labor input.

In the event that the adoption of this technology is both rapid and widespread—far from certain yet certainly possible—preparing the economy for a major national production shock is a first order policy concern. The 2024 Economic Report of the President dedicates an entire chapter to artificial intelligence (AI), stating: “The Federal Government’s role goes beyond ensuring that the gains brought about by AI are widely shared. It must also ensure that the costs to harmed individuals are addressed” (p. 273). While previous work has assessed which occupations are most likely to be exposed to LLMs (Brynjolfsson et al. 2018; Felten et al. 2018; Webb 2019; Eloundou et al. 2023), either as complements or substitutes in production (Acemoglu et al. 2022), historical experience with major shocks to the structure of the economy suggests that any initial employment impact is only the beginning of a process with many disparate downstream effects on geographic inequality (Connor et al. 2023), education, population (Berry & Glaeser 2005), and political behavior (Rodríguez-Pose et al. 2023).

In this paper, we analyze the downstream effects of an earlier major shock to national production—the decline of U.S. manufacturing employment starting in the early 1980s—to predict what we might expect from a large and rapid LLM shock in the absence of effective policy interventions to address individualized or localized harms that may persist (Autor

et al. 2021) even as aggregate national welfare increases. To do so we first trace out the trajectories of U.S. commuting zones (CZs) in terms of the organization of local production, the education composition of the local population, income growth, and voting behavior in presidential elections from 1980 (the peak of manufacturing employment) to 2019 (before the onset of the Covid-19 pandemic).<sup>1</sup> Over that time, manufacturing employment fell by roughly 6.5 million jobs, from 21% to 10% of U.S. workers (Pollard 2019). Accounting for aggregate growth in the labor force over time, the occupational displacement of production workers in the years following 1980 up through the 2008 financial crisis was of a magnitude similar to what LLMs might set in motion.

Beyond the similar potential magnitude of first order effects, the post-1980 manufacturing shock and its downstream effects serve as an apt comparison to LLMs' projected effects because: 1) they marked a permanent shift as opposed to a cyclical movement; 2) they affected different geographic areas unevenly, and 3) they disproportionately affected a specific labor force segment—manufacturing production workers—that had been enjoying relative prosperity in the immediately preceding period.

We can usefully divide the downstream effects of the manufacturing shock into adaptations made by people and adaptations made by places. People adapted to the shift in the labor market along three central avenues: by acquiring more education, by migrating to places with more opportunity, and by changing their political orientation. These three behaviors were loosely linked by the tendency of college graduates to migrate to urban areas with existing concentrations of college graduates (Berry & Glaeser 2005; Goworowska & Gardner 2012). With respect to places, the local economies most hurt by the manufacturing shock could have, in theory, changed what they produced. In practice, migration patterns made it difficult for local economies to attract the college-educated labor that expanding industries required. Local economies were left to maintain their existing industrial structures at lower levels. Those same migration patterns favored urban areas with existing concen-

---

<sup>1</sup>A Commuting Zone is set of counties grouped together to delineate a local economy. The concept was developed by the Economic Research Service of the U.S. Department of Agriculture.

trations of college graduates by both helping existing industries grow and by supporting the creation of “New Work” (Lin 2011).

As the national economy shifted while individual CZs maintained their industrial structures, the areas with the highest 1980 shares of manufacturing workers (and typically lower 1980 shares of college graduates) saw lower income growth rates, slower growth in their college populations, and partially in response to their declining economic standing, shifted from historical support of the Democratic Party to the Republican Party (Choi et al. 2024).

LLMs and the sharp pivot away from manufacturing employment are distinct shocks but the principal possibilities for adaptation—acquiring (or not acquiring) more education, migrating, changing one’s political outlook—are unchanged. As a result, we can draw meaningful comparisons by substituting the details of each situation and examining which populations and areas were (or are likely to be) affected in each case and what incentives and opportunities they might face in adjusting to the significant shift in demand. By combining the occupational measure of exposure to LLMs of Eloundou et al. (2023) with data on the local occupational composition within each metropolitan statistical area (MSA) from the Current Population Survey (CPS), we can characterize the areas most at risk of LLM disruption. Whereas in 1980 the areas most exposed tended to be in the center and north of the country<sup>2</sup> and generally had lower education levels, the areas most exposed to LLMs tend to be urban coastal areas with relatively high education. We consider the implications of these differences for labor mobility, geographic polarization, voting behavior, and the demand for education.

The main contribution of the paper is therefore twofold: we provide a literal map for predicting which geographic areas are most heavily exposed to the technological shock based on their occupational composition, and a figurative map for predicting the avenues through which people are likely to adjust following the initial employment shock that tends to receive the lion’s share of attention in the literature.

---

<sup>2</sup>And manufacturing regions in the Carolinas.

The remainder of the paper proceeds as follows. Section 2 presents a detailed analysis of the 1980 manufacturing shock and the geographic disparities in downstream effects. Section 3 characterizes the areas and education groups most likely to be impacted by the LLM shock. Section 4 extrapolates from the downstream effects of the 1980 production shock to predict what we might expect in turn from the LLM shock. Section 5 discusses the implications for public policy.

## 2 Training data: the national production shock of the 1980s

Like the coming LLM shock, the manufacturing shock's first round was largely man made. The man was Paul Volcker, appointed in 1979 as Chairman of the Federal Reserve with a mandate to break an accelerating inflation that arose from supply shocks in oil and agriculture and an accommodating monetary policy (Blinder 2022, Chapter 7). In the five years prior to Volcker's appointment, the Consumer Price Index had risen by a total 33.7%.<sup>3</sup>

Inflation was a national phenomenon but the strength of the real economy varied by region. The worldwide crop failures that created the agriculture supply shock led to a boom in U.S. agriculture (Ganzel 2009). The 1973 and 1979 OPEC oil boycotts created the oil supply shocks and led to a U.S. energy boom (Graefe 2013). Inflation itself led to a declining international value of the dollar that stimulated demand for U.S. manufactured exports despite rising Japanese competition. The aggregate results were a strong Midwest economy and strong demand for blue-collar workers such that economists began to question the economic value of a college education (Freeman 1976).

In October of 1979, Volcker introduced an extremely tight monetary policy that led to back-to-back recessions in 1980 and 1982. In June of 1982, the one-year *real* interest rate stood at 9.4% and the national unemployment rate at 9.6%.<sup>4</sup> The agriculture and oil booms collapsed and the international value of the dollar rose sharply, undermining manufactured

---

<sup>3</sup>See <https://fred.stlouisfed.org/series/CPIAUCSL>.

<sup>4</sup>For the complete data series, see <https://fred.stlouisfed.org/series/REAINTRATREARATIYE> and <https://fred.stlouisfed.org/series/UNRATE>.

exports.<sup>5</sup> Over time, the initial shock was reinforced by a series of trade shocks (Eriksson et al. 2021) and, eventually, the 2008 financial crisis. The events accelerated a long run shift in the U.S. occupational structure. In the consistent occupational coding developed by Pollard (2019), between 1979 and 1985 alone, “Production Occupations” declined by 1.7 million (-13%) while “Management and Business Occupations” increased by 4.0 million (+29%) and “Professional and Related Business Occupations” increased by 3.4 million (+24%). The latter two categories increasingly required a college degree. By the mid-1980s, parts of the Midwest had come to be known as the “rust belt” and the continental U.S. was described as a “bi-coastal economy” (U.S. Congress Joint Economic Committee 1986).

The other downstream outcomes of these occupational shifts reflected both the manufacturing shock and the way that places and people adapted (or failed to adapt) to changed conditions. With respect to places adapting, the largest concentration of manufacturing job losses was in a set of CZs located primarily in parts of the South, Midwest, and Northeast. Figure 2.1 highlights these areas on a map of the continental United States. The shaded CZs are those in the top 20% in terms of relative manufacturing employment losses from 1980 to 1990, calculated as the change in manufacturing employment divided by total local employment in 1980.<sup>6</sup>

---

<sup>5</sup>For a discussion of this period, see Levy (1999).

<sup>6</sup>The CZs in the top 20% have values of -0.020 or lower. The CZ with the largest loss is Gary, Indiana, with a value of -0.119.

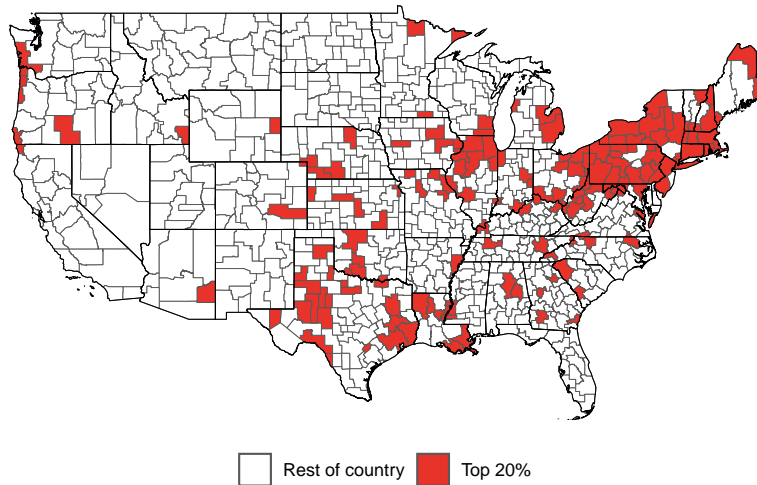


Figure 2.1: The CZs with the largest relative losses in manufacturing employment from 1980 to 1990

Despite the negative shock to manufacturing employment, most of these CZs found it difficult to “reinvent” themselves and, by default, maintained something like their previous industrial structure. Additional evidence on the surprising persistence of local negative shocks in general comes from Bartik (1993); Bound & Holzer (2000); Glaeser & Gyourko (2005); Moretti (2012); Yagan (2019); Notowidigdo (2020); Autor et al. (2021); Hershbein & Stuart (2024).

Figures 2.2 and 2.3 speak to this point with respect to the manufacturing shock. Figure 2.2 plots each CZ’s 1980 share of employment in manufacturing against its share of employment in manufacturing in 2019. There is a clear linear relationship: CZs that were initially concentrated in manufacturing continued to remain so (at lower levels) four decades later despite the dramatic nationwide shift away from manufacturing employment ( $\rho = 0.71$ ).




Figure 2.2: CZ shares of employment in manufacturing in 1980 and 2019

Reinventing the local economy might take significant time and investment—resources not readily available with a shrinking tax base and high unemployment. A CZ might optimally pivot towards the first opportunity available, including low wage service jobs. As one call center consultant said in 2019, “if I had to find 300 people willing to work for \$12 per hour, I would look at places where factories closed.”<sup>7</sup> The economic development director of Davidson County, North Carolina, where the local wooden furniture industry had been decimated, tried pivoting to other industries at first but eventually began to market the area as having a “manufacturing mentality.”<sup>8</sup>

If some CZs found it too costly to adapt, others had no need to adapt because they had concentrations of college graduates that were favored by the occupational shift. Figure 2.3 plots a CZ’s 1980 share of adults who were college graduates against its 2019 share. The scatterplot is nearly a straight line ( $\rho = 0.89$ ) indicating that the CZs with the highest concentrations of college graduates in 1980 were largely the CZs with the highest concentration of college graduates almost forty years later.

<sup>7</sup>Related in a personal communication with the authors, September 2, 2019. See also Autor et al. (2013).

<sup>8</sup>Steve Googe, personal communication with the authors, April 4, 2022. Eventually, his bet paid off.






Figure 2.3: CZ adult population shares of college graduates in 1980 and 2019

Places—local economies—were not adapting to the manufacturing shock as much as the people who lived there were adapting, both through acquiring more education and through migration.

The production shift that accelerated in the initial manufacturing shock led to a sharp reversal in the earnings premium for a four-year college degree. In 1979, the hourly earnings difference between workers with a high school diploma and workers with a bachelor’s degree was 40.1%, the smallest gap since the blue-collar catch-up wage gains of the late 1940s. Starting in 1980, the gap rose steadily, reaching 67.4% by 2017 (Autor et al. 2020b).<sup>9</sup>

At first glance it appears that people were not responding to the incentive. The number of bachelor’s degrees awarded rose only slightly from an average of 950,000 in the early 1980s to only 1,050,000 in the early 1990s. In reality, people were responding but the number of college-age young people was shrinking. The baby-bust cohorts that began in 1965 were now turning 17-18. By 1990, the number of 17–18-year-olds in the population— young people making college-going decisions—had fallen by about 17% even as the rate of young people going to college was rising (Card & Lemieux 2001). This growing rate is

---

<sup>9</sup>Contributing to this reversal was the 1980 legislation deregulating interstate trucking that sharply undercut truck drivers’ pay. See Viscelli (2016).

reflected in Figure 2.4. The figure plots the post-1980 evolution of the BA earnings premium separately for young men juxtaposed with the male rate of BA attainment (left) and for young women juxtaposed with the female rate of BA attainment (right). In both cases the earnings premiums are lagged by four years to reflect what individuals would have observed when making their college-going decisions.<sup>10</sup>

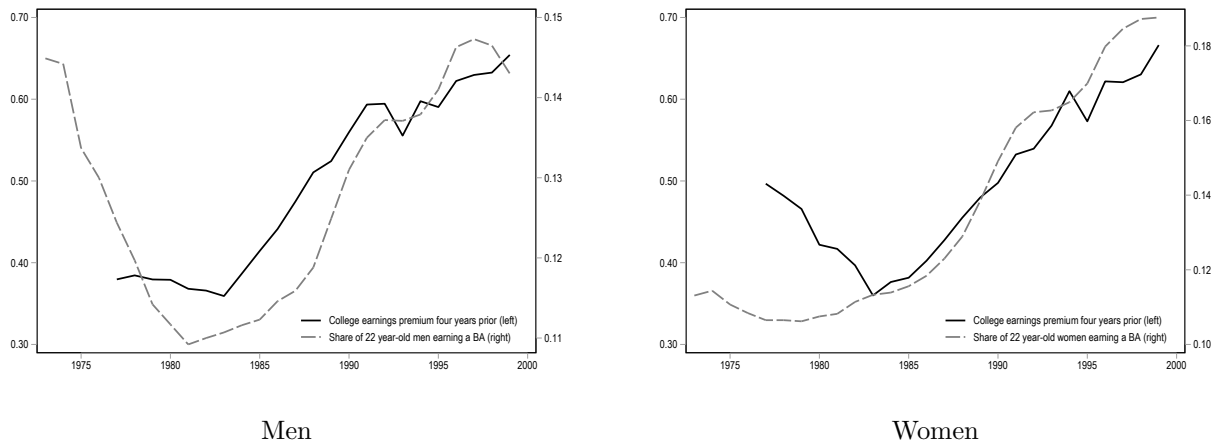


Figure 2.4: BA attainment rates and the college wage premium, 1973 to 1999

The growing college population was, in turn, relatively likely to migrate to economic opportunity. In the period from 1985-1990, among all persons aged 25-39 (in 1985), 13.1% moved to a different state. Within this group, roughly 20% of college graduates moved to another state compared to 11% for non-graduates (Goworowska & Gardner 2012).

Berry & Glaeser (2005) showed that across U.S. cities the share of adults with a college diploma grew faster between 1970 and 1980 in cities with higher 1970 average education levels, a result that continued to apply in subsequent decades. Since that time, multiple studies have explored explanations for this result, ranging from better job matching to greater amenities to the increased possibility that members of two-career couples would both find work (Dauth et al. 2022; Glaeser et al. 2001; Compton & Pollak 2007).

As the nation's occupational structure tilted toward college graduates, the migration of

<sup>10</sup>The higher trend in college attainment for women likely also captures the concurrent underlying trend in changing gender norms in society (including legal norms) and rapidly growing women's participation in the labor force.

college graduates to cities helps explain why local economies did not or could not change their industrial structures (See Figures 2.2 and 2.3 for evidence of this persistence). Again, local economies that lost manufacturing jobs had difficulty keeping college graduates from leaving, which made it harder to recruit new industries (Bound & Holzer 2000). The college graduates gained by large urban areas made them increasingly attractive as places for new industries to expand and for other graduates to relocate to (Glaeser et al. 2001; Florida 2002).

One result of these adaptations was a marked slowdown in the convergence of per capita income across geographic areas. Convergence of average per capita income across states had been a longtime feature of the U.S. economy, but the convergence slowed dramatically after 1980 (Ganong & Shoag 2017). Inequality of per capita income across CZs, measured by the coefficient of variation, had fallen sharply during the 1970s, driven in part by the strong Midwestern economy.<sup>11</sup> Starting in 1980, inequality of per capita income began to move erratically. It then increased steadily after 1993. This trend is illustrated in Figure 2.5, which plots the coefficient of variation for CZ per capita Personal Income by year from 1969 to 2019.




Figure 2.5: The coefficient of variation for per capita Personal Income across CZs over time

<sup>11</sup>The coefficient of variation is the standard deviation divided by the mean.

The role of the manufacturing shock in the end of convergence is illustrated in Figure 2.6, which maps the minimum number of CZs whose Personal Incomes sum to 50% of U.S. Personal Income. In 1980, it required the 30 largest CZs (in terms of Personal Income) to sum to 50% of U.S. Personal Income. By 2019, increased geographic concentration of economic activity meant that the largest 27 CZs now equaled 50% of U.S. Personal Income. The reduction from 30 to 27 CZs was a net result of losing five manufacturing CZs—Kansas City, Milwaukee, Cincinnati, Cleveland, and Buffalo—and gaining two sunbelt cities—Austin and Port St. Lucie.<sup>12</sup> In 1980, per capita Personal Income was 35% larger in the 30 largest CZs than in the rest of the country. In 2019, this gap had increased to 45% (in the new top 27). The difference was due in part to the increasing concentration of college graduates in the largest CZs.

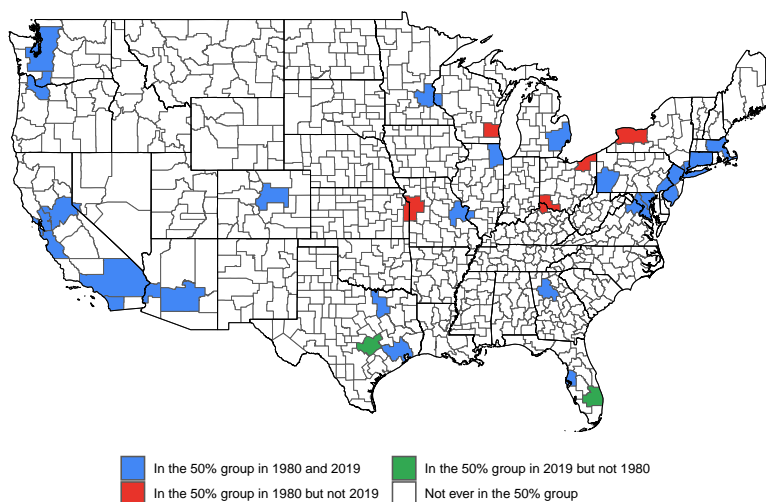


Figure 2.6: The minimum number of U.S. CZs accounting for 50% of total Personal Income in 1980 and 2019 (Alaska and Hawaii have no CZs in either group)

The third adaptation pursued by individuals was changing their political orientation. Changing political orientation was a long process in which college graduates moved from the Republican to the Democratic Party and non-graduates moved from the Democratic to

<sup>12</sup>From the 1940s through the early 1980s, Kansas City had one of the largest garment industries in the nation. See <https://www.kcur.org/community/2017-09-01/museum-expansion-highlights-garment-manufacturing-history-in-kansas-city>.

the Republican Party. Each move involved cultural and economic factors but both were connected to the manufacturing shock.

Non-graduates' shift to the Republican Party was initially driven by Democrats' lack of response as manufacturing jobs disappeared. In the early 1980s, the lost jobs were unionized durable goods jobs in the upper Midwest.<sup>13</sup> By the 1990s, the lost jobs were non-union textile and wooden furniture jobs in the in the mid-South. In the case of North Carolina factory workers, Democratic office holders were actively supporting free trade agreements that were costing jobs in textiles, wooden furniture and tobacco (Neff et al. 2002; Choi et al. 2024). As Neff et al. (2002) write, "the list of Democratic free-trade advocates through the 1990s reads like a Tar Heel Who's Who..." Leading the charge from the White House was President Bill Clinton and his support for the "Washington Consensus", a set of ideas that embraced globalization regardless of its effect on manufacturing jobs (Levy & Temin 2009).

Democratic support for expanded trade and the "New Economy" represented a sharp turning away from their traditional blue-collar base. Rank and file workers viewed Clinton's support for NAFTA as a betrayal (Frank 2016; Choi et al. 2024). As Newman & Skocpol (2023) note in reference to Western Pennsylvania, the closing of a factory often destroyed social and political networks: "Union members expressed loyalty and gave support because they expected these institutions to have their backs and act as partners to them and their families over the long term, in times of both fun and struggle" (p.3). The authors emphasize that this loyalty (and expectation) had often extended to the Democratic Party.

The manufacturing shock was increasing separation between college and high school graduates along multiple dimensions. Between 1979 and 1985 alone, the average hourly earnings of men with a high school diploma fell from \$17.37 to \$16.37 (-8%) while the average earnings of men with four years of college rose from \$24.35 to \$25.15 (+5%). The difference between the unemployment rates of high school and college graduates increased from 6.5

---

<sup>13</sup>For a discussion of the loss of unionized manufacturing jobs, see Hirsch & Macpherson (2003). As the authors show, between 1979 and 1985 the fraction of unionized manufacturing jobs in the U.S. fell from 37.6% to 26.9%. For a more general discussion of the changing geography of manufacturing jobs, see Eriksson et al. (2021).

percentage points (1980) to 9 percentage points. By 2000, 61.7% of college graduates lived in the 50 largest MSAs while 49.2% of high school graduates lived outside the 50 largest MSAs.

Historically, college graduates had been more liberal on social issues than high school graduates (e.g., views on abortion). By the 1990s college graduates, perhaps influenced by geographic clustering and “echo chamber effects” were increasingly liberal on all economic issues and moving to the Democratic Party (Sunstein 2002; Bishop 2009; Parker 2019). The rural-urban political divide was nothing new in America, but the recent intensity of the divide was another downstream effect of the manufacturing shock.

To quantify the shift in voting behavior along educational lines we can plot presidential election shares over time. In 1980, Ronald Reagan’s margin over Jimmy Carter was 23 points among college graduate voters and 8.1 points among non-graduate voters. Figure 2.7 shows how starting after Reagan’s second term (1988), college voters gradually moved toward the Democratic Party while non-college voters moved to the Republican Party. By 2016, 66% of non-graduate voters and 71% of white non-graduate voters supported Donald Trump (Jones 2018).




Figure 2.7: Republican share of presidential election vote by education group

We can characterize these downstream political trends empirically by regressing a CZ's 2016 Republican presidential vote share on its initial endowments. Table 2.1 reports the results from a regression of the later Republican vote on a (continental) CZ's 1980 share of employment in manufacturing (Column 1) and its 1980 college-educated share of the adult population (Column 2), controlling for urban/rural status and the baseline Republican vote share. Initial manufacturing intensity is a strong positive predictor of a shift towards the Republican Party, while higher initial education levels are a strong negative predictor. Every 10 additional points in a CZ's 1980 manufacturing share predict about a 2.5 point increase in its Republican vote share by 2016. Every 10 additional points in its 1980 college share predict about a 19 point decrease in its downstream Republican vote share.

Table 2.1: The relationship between 1980 CZ manufacturing and college shares and growth in the Republican presidential vote share

	Republican vote share (2016)	
	(1)	(2)
Manuf. share of emp. (1980)	0.245*** (0.0558)	
College share of pop. (1980)		-1.862*** (0.0947)
Republican vote share (1980)	0.579*** (0.0467)	0.662*** (0.0355)
Urban area	-0.114*** (0.0100)	-0.0369*** (0.00867)
Constant	0.318*** (0.0322)	0.513*** (0.0220)
Observations	725	726
$R^2$	0.304	0.534

Standard errors in parentheses

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

Together, the three avenues of adjustment for people (and the relative non-adjustment of places) became a highway to alignment, with geographic location now becoming a strong predictor of education, income, and political outlook among the population. The result was a sharper distinction between red states and blue states, and between urban and rural areas within states: national polarization.

### 3 The geography of occupational exposure to LLMs

The core operation of a Large Language Model involves analyzing a sequence of words and predicting the next word in the sequence (Wolfram 2023). The quality of the prediction depends critically on the length of the word sequence being analyzed, which, in turn, depends on the capacity of computer resources and the quantity of data on which the model is being trained (Vaswani et al. 2023). In recent years, expanded computer and data resources have allowed the development of models that solve complex statistical problems, write computer code, and can interpret and respond (with varying degrees of accuracy) to commands like this: “Please summarize the scholarly contribution of this economics research article for a peer-reviewed journal based on the text of the introduction.”<sup>14</sup> Over time, software with these capabilities will have a widespread labor market impact.

Having demonstrated that geographic variation in exposure to the 1980 production shock was a central feature shaping its downstream effects and the three avenues of labor market adjustment, we now map the geography of the projected shock from LLMs to anticipate that potential impact. To do this, we begin with the measure of occupation-level exposure to LLMs developed by Eloundou et al. (2023). The measure analyzes the task composition of each occupation to assess what fraction of those tasks an LLM, or software built on top of an LLM, could speed up by at least 50% (or replace entirely). The measure does not indicate

---

<sup>14</sup>“This research article contributes to the understanding of the labor market impact of large-language models by providing a detailed analysis of geographic disparities in exposure to technological shocks. Using the historical decline of U.S. manufacturing employment as a comparative framework, it predicts how regions and populations might adapt to the LLM shock through education, migration, and political behavior, offering both a literal and figurative map for identifying areas most at risk and potential avenues for adjustment and policy intervention.” (ChatGPT, June 4, 2024).



whether that means an occupation will be replaced with a computer model or become more productive in conjunction with one,<sup>15</sup> but it does point towards which occupations are most likely to experience some form of direct labor demand impact from the shock.<sup>16</sup> A key distinction between the 1980 shock and the LLM shock is that the first order effects of the former were on manual occupations such as textile manufacturers, while the latter is expected to primarily affect office jobs such as accounting and back-end administration.

Because detailed occupational data is not publicly available at the county or CZ level, we instead measure geographical exposure to LLMs by core based statistical area (CBSA) as defined by the U.S. Census Bureau. CBSAs include metropolitan statistical areas and micropolitan statistical areas, and must have at least one urban area with a population of 10,000 or greater. Hereafter, we refer to the 261 such areas for which we can calculate the occupational composition using CPS data, which are primarily metropolitan areas, as MSAs.

To calculate MSA-level occupation shares, we first pool the basic monthly CPS samples from 2021 through 2024. We then link each occupation in the CPS data to its LLM exposure score in the data from Eloundou et al. (2023), where exposure is measured by SOC code. To obtain an MSA-level exposure score, we then take a weighted average of the exposure scores of the occupations within each MSA, weighted by that occupation’s share in local employment.

Figure 3.1 maps the MSAs in the continental U.S. that are in the top 20% based on occupational exposure to the LLM shock. Many of these MSAs are on the East or West Coast, while a few areas with large concentrations in professional services are in the middle of the country. The list of the most exposed MSAs includes Washington, DC; Boulder and Denver, Colorado; San Jose and San Francisco, California; New York City; and Austin, Texas. This pattern aligns closely with the geographic analysis of AI exposure more broadly

---

<sup>15</sup>Elimination of a human worker is the limiting case of a technology that will enable any affected job to be performed with less labor. For example, automated teller machines did not replace bank workers but did lead to fewer per branch.

<sup>16</sup>We specifically use Eloundou et al.’s measure  $\gamma$ . All three of the measures of LLM occupational exposure used in the the paper yield comparable results regarding the geographic dispersion of exposure.

in Muro et al. (2019) using the Webb (2019) data, which identifies exposure by comparing occupational task information to a text analysis of the content of patents related to AI. Felten et al. (2021) also identify cities like Boston and Washington as having high general AI employment exposure according to an industry-focused measure.

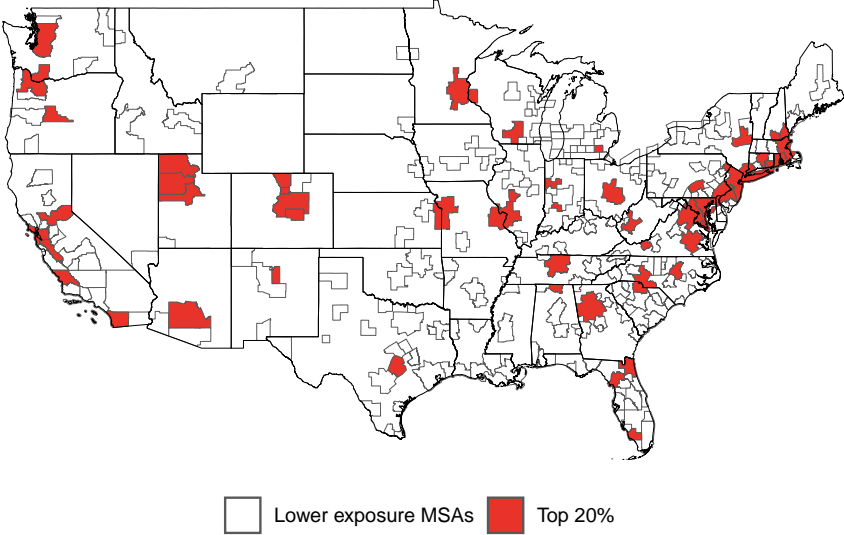


Figure 3.1: Map of the MSAs with the highest occupational exposure to LLMs

What do the MSAs with the greatest exposure to LLMs have in common? In particular, we are interested in whether they share population concentrations by education level or voting behavior. First, we can correlate our MSA-level exposure ratings with average years of education. When we do so, we obtain a correlation coefficient of  $\rho = 0.658$ . The tight relationship is further illustrated by Figure 3.2, a scatterplot with the share of college graduates in an MSA on the x-axis and our MSA-level measure of exposure to LLMs on the y-axis with a line of best fit added.

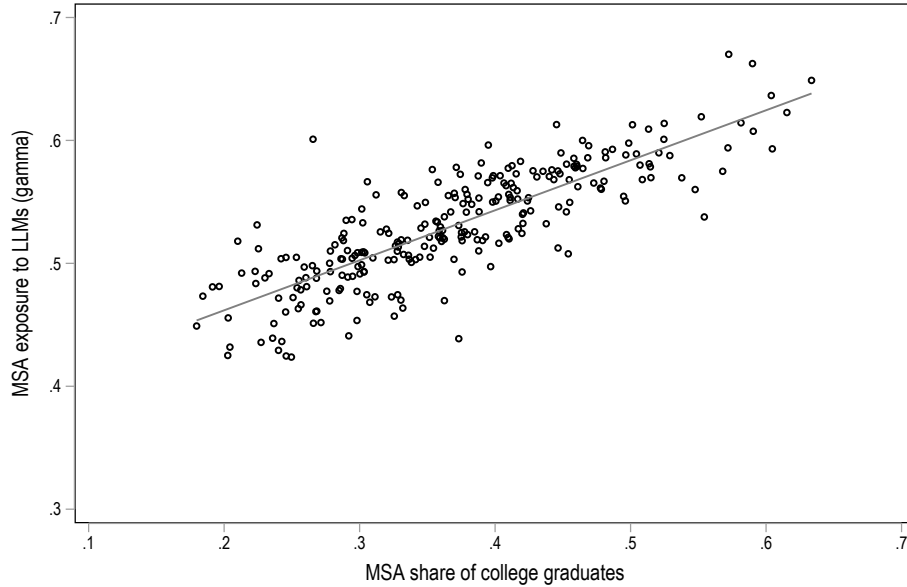


Figure 3.2: Scatterplot of the relationship between the share of college graduates in an MSA and its exposure to LLMs

A similar exercise comparing MSA LLM exposure and the 2020 Republican presidential vote share reveals that the most exposed MSAs tended to vote Democrat. The correlation between the two measures is  $\rho = -0.386$ . Although the relationship is not as tight as the relationship between education and LLM exposure, the scatterplot in Figure 3.3 shows that in general the MSAs with higher Republican vote shares are less exposed than other MSAs (at least for the present).




Figure 3.3: Scatterplot of the relationship between the 2020 Republican presidential vote share in an MSA and its exposure to LLMs

## 4 Predicting the three avenues of adjustment to LLMs

From the analysis in Section 2, one feature of U.S. geographic areas is that they adapt more slowly to major shocks than do the people who live (or used to live) there. As the latter half of the 20th century saw a large change in what the nation produced, MSAs with large initial college populations were well suited to these trends and grew rapidly in population and per capita Personal Income. Areas with smaller college shares saw population and per capita income grow more slowly (or even decline). The result was persistence across both variables.

A rank-rank comparison of commuting zones between 1980 and 2019 on the dimension of population ( $\rho = 0.98$ ) and per capita Personal Income ( $\rho = 0.71$ ) reveals relatively few changes in ordering, and many of the changes in per capita income rank are a function of several smaller CZs with high recent fracking activity. Under the hypothesis that places exhibit persistence in the organization of local production while people make greater adjustments following a national production shock, we next apply our model of the three avenues

of adjustment developed by studying the 1980 manufacturing shock to anticipate the downstream patterns we might expect following the LLM shock in terms of changes in educational attainment, migration patterns, and voting behavior.

#### **4.1 Avenue 1: educational attainment**

Starting in 1980, the loss of manufacturing jobs contributed to a shifting demand for labor that raised the return to a college education. A central adjustment to this shift was an increasing fraction of young cohorts who enrolled in college. Where the loss of manufacturing jobs reduced demand for non-college workers, an LLM shock is likely to reduce demand for college graduates, a conclusion supported in both the literature and the correlation between MSA education levels the measure we construct in Section 3.

Seventy-five percent of Americans already believe that a college degree is not very important for finding a well-paying job in today’s economy, and only 22% feel the cost of college is worth it (Fry et al. 2024). Because of these views and slowing growth among the number of college-age persons, some universities are already facing enrollment pressures. Should a shock to the demand for college employment further reduce the incentive to have a four-year degree, it is reasonable to predict both a decrease in relative supply of college graduates and a growing number of four-year colleges in financial trouble.

A decrease in educational attainment may not necessarily increase disparities, if LLMs can substitute for labor market skills and experience (Noy & Zhang 2023; Brynjolfsson et al. 2023). Autor (2024) interprets this possibility as a positive development that will “restore the middle class” by reducing the college/non-college wage gap. The argument is that “better” jobs will become newly available to the types of workers who fell behind due to the previously rising skill premium. In fact, whereas traditionally a college degree has been a prerequisite for many white-collar jobs, a recent analysis of trends in job postings reveals that employers are already beginning to remove a college degree as a job requirement, with about a fourfold increase from 2014 to 2023 in the annual rate of posted positions dropping the requirement (Fuller et al. 2024).

An alternative prediction is that the LLM shock will widen a different wage gap. Rather than the college/non-college framework, consider an LLM/non-LLM framework among holders of a college degree. Particularly if LLMs tend to substitute for certain occupations while complementing others, there may be a form of skill-biased technical change occurring among those with higher education. College graduates are not a homogeneously skilled demographic group, and there already exists significant variation in the returns to a college education by major (Altonji et al. 2012). The shift in labor demand induced by the LLM shock may well further widen the gap, especially in the MSAs with the greatest occupational exposure. As an illustrative example, the *Wall Street Journal* reports that in 2023, new job listings for workers within the information technology industry who have AI skills rose by 43%, whereas overall IT job listings declined by 31% (Rattner 2024). A similar trend has been documented in the academic literature (Acemoglu et al. 2022). If the effect of LLMs is to reduce the wage premium for a portion of college graduates, part of the rebuilding of a middle class as Autor (2024) envisions might develop not by bringing up workers from the bottom of the wage distribution, but by bringing down workers from the top.

There may even be three groups to consider among college graduates: those directly involved in LLM work, who will likely be the biggest “winners”; those whose jobs are complemented by LLMs, who may become more productive and therefore experience wage gains, and those who are replaced by AI. It is this last group in particular that we predict is most likely to out-migrate in search of new labor market activities and more affordable living costs given potentially slower wage growth for their skill sets.

## **4.2 Avenue 2: migration patterns**

Supposing that the subset of geographic areas identified in Section 3 are the ones whose labor markets are most directly affected by the LLM shock, and supposing as well that the organization of local production tends to be persistent over time, then in these areas there will be a group of workers who see their opportunities shrinking, and these workers will likely be college graduates who do not have the human capital to take advantage of AI tools.

The 1980 national production shock encouraged a prevailing trend of college graduates tending to migrate away from small towns and towards large cities. It is unlikely that the river will begin to flow in the other direction. As made clear in the literature, college graduates are attracted to destinations with existing concentrations of college graduates, and also value the consumption amenities that cities have to offer (Glaeser et al. 2001). In addition, “new work”—jobs resulting from innovation—is more likely to arise in locations that are dense in college graduates and diverse in industry (Lin 2011; Kim et al. 2024). Under that framework, college graduates in the third group, who see their wage premiums and labor market opportunities decline in the large urban areas where exposure to the LLM shock is the greatest, might reasonably be predicted to move away, but to other cities where there are already many college graduates and exposure is lower.

To generate a list of potential destinations, we first sort the metropolitan areas in our dataset for which we also have a measure of college attainment and population by LLM exposure (the mean value is 0.53). We then filter the list for MSAs with a population above the median value of about 410,000 and with at least a 35% share of college graduates, which is just below the median concentration among all MSAs in our data. A priori, an MSA might have low exposure but also low opportunity for college graduates, hence the selection based on already having a critical mass of graduates active in the local labor market. Finally, we look at the cost of living based on median list prices per square foot of houses in the MSA, keeping MSAs below \$225 per square foot (approximately the median national value) as of January 2024.<sup>17</sup> For reference, the correlation between log median list price per square foot and LLM exposure across MSAs is  $\rho = 0.31$ , and the average of the median list price per square foot for the MSAs in the top 20% of LLM exposure depicted in Figure 3.1 is \$304.

Table 4.1 presents the top locations that simultaneously possess relatively low LLM exposure with relatively high populations and education levels and affordable housing. In general, the areas on the list tend to be small-to-medium-sized cities in the South and Midwest, with

---

<sup>17</sup>Housing price data comes from the Federal Reserve Bank of St. Louis FRED database. See <https://fred.stlouisfed.org/release/tables?eid=1138280&rid=462>.

Dayton, Ohio at the top. At the state level, in addition to Ohio, which also has Toledo on the list, Georgia (Augusta, Savannah), Pennsylvania (Scranton, Allentown), South Carolina (Greenville, Columbia), and Louisiana (New Orleans, Baton Rouge). To help further narrow the list, the second from rightmost column reports the cumulative population growth rate from 2020 through 2023 as estimated by the U.S. Census Bureau.<sup>18</sup> Current population growth could serve as a leading indicator for MSAs that are already desirable destinations because of local amenities or growing economic opportunity. Looking at the MSAs on the list with the highest rates of recent population growth, Augusta, Savannah, Greenville, Chattanooga, Oklahoma City, Columbia, and Houston all appear to be potentially attractive destinations for “non-AI” college graduates displaced by the LLM shock. To preview the discussion of voting behavior, the rightmost column of Table 4.1 lists the fraction of voters in each MSA supporting the Republican candidate in the 2020 presidential election. On balance, migration involves people moving from large Democratic-leaning MSAs to smaller Republican-leaning MSAs. We return to this point in the next section.

---

<sup>18</sup>See <https://www.census.gov/data/tables/time-series/demo/popest/2020s-total-metro-and-micro-statistical-areas.html>



Table 4.1: List of the top large metropolitan areas with high education levels, low employment exposure to LLMs, and affordable housing costs

Rank	MSA	College	Exposure	Population	Growth	Housing	Republican
1	Dayton-Kettering-Beavercreek, OH	37.0%	0.503	814,363	0.0	\$124	53.6%
2	Augusta-Richmond County, GA-SC	35.2%	0.505	629,429	3.0	\$154	51.8%
3	Savannah, GA	40.3%	0.516	424,935	5.0	\$223	48.8%
4	Scranton-Wilkes-Barre, PA	36.1%	0.518	569,413	0.3	\$135	52.6%
5	Allentown-Bethlehem-Easton, PA-NJ	37.6%	0.518	873,555	1.4	\$200	50.0%
6	Greenville-Anderson-Greer, SC	41.0%	0.520	975,480	5.1	\$183	63.4%
7	York-Hanover, PA	35.8%	0.522	464,640	1.8	\$169	61.5%
8	Greensboro-High Point, NC	38.0%	0.523	789,842	1.7	\$173	48.2%
9	Toledo, OH	40.8%	0.523	600,141	-1.0	\$146	45.8%
10	Chattanooga, TN-GA	37.5%	0.525	580,971	3.3	\$218	62.2%
11	New Orleans-Metairie, LA	35.9%	0.526	962,165	-4.5	\$181	39.5%
12	Jackson, MS	41.7%	0.528	610,257	-1.6	\$150	47.5%
13	Oklahoma City, OK	36.0%	0.530	1,477,926	3.7	\$173	57.0%
14	Buffalo-Cheektowaga, NY	42.1%	0.532	1,155,604	-1.0	\$183	43.6%
15	Baton Rouge, LA	35.6%	0.534	873,661	0.4	\$163	54.9%
16	Akron, OH	37.9%	0.542	698,398	-0.5	\$140	47.0%
17	Louisville-Jefferson County, KY-IN	38.8%	0.542	1,365,557	0.3	\$173	50.1%
18	Milwaukee-Waukesha, WI	44.7%	0.546	1,560,424	-0.9	\$208	44.3%
19	Columbia, SC	38.3%	0.548	858,302	3.5	\$154	46.5%
20	Little Rock-North Little Rock-Conway, AR	37.6%	0.549	764,045	2.1	\$154	52.6%
21	Rochester, NY	45.5%	0.550	1,052,087	-1.2	\$162	42.9%
22	Birmingham, AL	39.8%	0.550	1,184,290	0.3	\$160	57.4%
23	Houston-Pasadena-The Woodlands, TX	40.1%	0.551	7,510,253	5.0	\$178	49.0%

Note: growth is population growth from 2020-2023. Housing is median list price per square foot, April 2024. Republican is the 2020 Republican presidential vote share.

### 4.3 Avenue 3: changing political orientation

In Section 2, we described how individual adjustments to the manufacturing shock increased the geographic alignment of education, income, and political orientation. Increased alignment led to sharper distinctions between red states and blue states, and between urban and rural areas within states—distinctions that advanced political polarization (Bishop 2009). By contrast, individual adjustments to LLMs may well undo some of this alignment.

The potential for de-alignment comes from a political division that cuts across traditional party lines: persons who want AI development and diffusion to be regulated versus persons who want AI development to proceed without restrictions: Team Apocalypse versus Team Utopia (Tiku 2023).

As of 2024, most people are not yet on either team, but that will change as LLMs and other AI diffuse through the economy and begin to disrupt employment. During the manufacturing shock, blue collar Democrats who lost their jobs frequently moved to the Re-

publican Party (Figure 2.7). Since the 1990s, Democrats and Republicans have significantly diverged on a number of ideological issues (Desilver 2022), including abortion and global warming (Ajasa et al. 2023). In this political climate, it is hard to imagine the loss of a job by itself is enough to cause any Republican or Democrat to switch parties in the near future, a conclusion supported by recent political science literature (Newman & Skocpol 2023).<sup>19</sup>

Over the longer run conditions are ripe for displaced workers from both parties to support a populist version of Team Apocalypse and AI regulation.<sup>20</sup> The most visible proponents of deregulated AI are a set of very rich and very vocal individuals.<sup>21</sup> This visibility is politically important since U.S. politics has shifted over time from who you support to who you oppose (Finkel et al. 2020) and 60% of U.S. adults say wealthy individuals don't pay their fair share of taxes (Oliphant 2023). As AI diffuses through the economy, political competition suggests that people will run for office on regulating an industry in which some people earn very high incomes to develop software that eliminates other peoples' jobs.

Earlier, we saw how LLM vulnerability is greater for “information processing” occupations with high educational requirements. On balance, MSAs most likely to be disrupted by LLMs are larger/more urban ( $\rho = 0.359$ ), have higher shares of college graduates ( $\rho = 0.658$ ) and have higher proportions of Democrats ( $\rho = 0.359$ ). Given the concentration of college graduates in urban MSAs, the push for AI regulation might appear to be stronger among urban Democrats than rural Republicans. In fact, the pattern of LLMs primarily affecting more educated workers may be temporary: Researchers are already exploring how LLMs can reshape robotics in ways that will ultimately affect blue-collar work as well (Holoubek 2024; Somers 2024). Over time, we can anticipate pro-regulation factions in each party to consider making common cause to increase their political power. This cross-party alliance and the previously noted migration of population from large Democratic-leaning MSAs to smaller

---

<sup>19</sup>In a survey article, Margalit (2019) writes: Economic shocks—e.g., job loss or sharp drop in income—exert a significant and theoretically predictable, if often transient, effect on political attitudes. In contrast, the effect on voting behavior is more limited in magnitude and its manifestations less understood.

<sup>20</sup>Team Apocalypse began as an informal group of computer scientists who feared that unregulated artificial intelligence could extinguish humanity (Stokel-Walker 2024).

<sup>21</sup>Examples include Marc Andreessen, Elon Musk and Peter Thiel.

Republican-leaning MSAs may both serve to moderate some of the geographic separation of the last 40 years.

## 5 Conclusion: déjà vu all over again?

On the surface, the downstream effects of LLMs appear more benign than the downstream adjustments to the manufacturing shock. The falling rate of return to a four-year degree will open opportunities for large numbers of persons. The migration from larger to smaller MSAs will not result in places left behind. Such political alignment as occurs will cut across current political dividing lines.

The caveat is that speed of adjustment matters. The loss of jobs from the manufacturing shock largely occurred in two bursts: the back-to-back recessions of 1980-82 (-1.6 million manufacturing jobs; -8%) and the 2000-2012 China shock and financial crisis (-5.3 million manufacturing jobs, -30%). The rapid loss of manufacturing jobs was not the only factor in today's political polarization but it was a central factor both in the U.S. and globally (Autor et al. 2020a; Hill 2021; Guriev & Papaioannou 2022).

Economic adaptation depends in part on societal adaptation. For this reason, the political economist Karl Polanyi argued that a central responsibility of government is to regulate the pace of change, asserting that the speed of adoption determined the social impact:

“whether the dispossessed could adjust themselves to changed conditions without fatally damaging their substance, human and economic, physical and moral; whether they would find new employment in the fields of opportunity indirectly connected with the change; and whether the effects of increased imports induced by increased exports would enable those who lost their employment through the change to find new sources of sustenance” (Polanyi 1944).

Word-processing software was a shock to the typist occupation but the pace of adoption was slowed by users having to purchase computers and install the software. The number of persons employed as “Word Processors and Typists” declined from one million in 1980 to

33,000 today but there was no huge disruption because the decline was spread over 44 years.<sup>22</sup> David (1990) points to similar constraints on electrification in the early 1900s. Compared to these earlier episodes, the dissemination of LLM-based services through the web and the cloud, absent regulation, will be much faster and, like electrification, very broad.

There are now many proposals to regulate AI. So far, the majority have focused on “red team” exercises that test safety and privacy, and proposals on the ownership of intellectual property.<sup>23</sup> The recent White House Executive Order on Artificial Intelligence<sup>24</sup> has a brief section titled “Supporting Workers” but there is no discussion of dissemination speed—an admittedly difficult topic—while the Executive Order’s call to support workers’ ability to bargain collectively is more realistic in Europe where workers are more unionized than in the United States (Kochan et al. 2024).

Software companies have long argued that they can regulate themselves.<sup>25</sup> There are reasons to be skeptical. In recent decades, many software technologies have tended toward a few dominant firms driven by network externalities. The dominant firms and their management became extremely wealthy as a result.

AI software—in particular LLMs—reinforces this tendency toward dominant firms and concentrated money. Developing and training an LLM is extremely expensive both in terms of computational time and in data and power requirements. One result is that university researchers are priced out of developing this technology (Nix et al. 2024). A second result is that the LLM industry is dominated by a few firms. These firms have to answer to their shareholders, who expect returns. While firms may speak about social responsibility, they will be under enormous pressure to develop products as fast as possible, attempting to deflect

---

<sup>22</sup>See <https://www.bls.gov/opub/mlr/1992/08/art3full.pdf> and <https://fred.stlouisfed.org/series/LEU0254503100A>.

<sup>23</sup>See <https://sd11.senate.ca.gov/news/bipartisan-vote-senate-passes-senator-wieners-landmark-ai-safety-and-innovation-bill> for one such example in California. A broader scope of regulation is described in Anderljung et al. (2023).

<sup>24</sup>See <https://www.whitehouse.gov/briefing-room/statements-releases/2023/10/30/fact-sheet-president-biden-issues-executive-order-on-safe-secure-and-trustworthy-artificial-intelligence/>.

<sup>25</sup>See <https://www.youtube.com/watch?v=CalhOA66ekA> for one such interview with Sam Altman, CEO of OpenAI.

countervailing government regulation if necessary.<sup>26</sup> Given the firms' economic power and the importance of money in current political life, the outcome of a government/industry contest is not easy to predict.

If history does not repeat itself exactly, it has been known to rhyme.<sup>27</sup> Though the manufacturing shock may have increased economic activity in the aggregate, it left a persistent wake of harm across specific swaths of the United States, with downstream effects that undermined the stability of the system. If it is possible to learn from this so as to dampen the effects of an LLM shock, it would be prudent to do so.

---

<sup>26</sup>The situation brings to mind the H.L. Mencken quote, "When somebody says it's not about the money, it's about the money."

<sup>27</sup>The observation is often credited to Mark Twain.

## References

- Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1), S293–S340.
- Ajasa, A., Clement, S., & Guskin, E. (2023). Democrats and Republicans deeply divided on extreme weather, Post-UMD poll finds. *Washington Post*. 23 Aug 2023. Accessed 10 Jun 2024.  
URL <https://www.washingtonpost.com/climate-environment/2023/08/23/extreme-weather-climate-change-poll/>
- Altonji, J. G., Blom, E., & Meghir, C. (2012). Heterogeneity in human capital investments: High school curriculum, college major, and careers. *Annu. Rev. Econ.*, 4(1), 185–223.
- Anderljung, M., Barnhart, J., Korinek, A., Leung, J., O’Keefe, C., Whittlestone, J., Avin, S., Brundage, M., Bullock, J., Cass-Beggs, D., et al. (2023). Frontier AI regulation: Managing emerging risks to public safety. *arXiv preprint arXiv:2307.03718*.
- Autor, D. (2024). Applying AI to rebuild middle class jobs. Tech. rep., National Bureau of Economic Research.
- Autor, D., Dorn, D., & Hanson, G. (2013). The geography of trade and technology shocks in the united states. *American Economic Review*, 103(3), 220–225.
- Autor, D., Dorn, D., & Hanson, G. (2021). On the persistence of the China shock. *Brookings Papers on Economic Activity*, (pp. 381–447).
- Autor, D., Dorn, D., Hanson, G., & Majlesi, K. (2020a). Importing political polarization? the electoral consequences of rising trade exposure. *American Economic Review*, 110(10), 3139–83.
- Autor, D., Goldin, C., & Katz, L. F. (2020b). Extending the race between education and technology. *AEA Papers and Proceedings*, 110, 347–351.
- Bartik, T. J. (1993). Who benefits from local job growth: migrants or the original residents? *Regional Studies*, 27(4), 297–311.
- Berry, C. R., & Glaeser, E. L. (2005). The divergence of human capital levels across cities. *Papers in Regional Science*, 84(3), 407–444.
- Bishop, B. (2009). *The Big Sort: Why the clustering of like-minded America is tearing us apart*. Houghton Mifflin Harcourt.
- Blinder, A. S. (2022). *A monetary and fiscal history of the United States, 1961–2021*. Princeton University Press.

- Bound, J., & Holzer, H. J. (2000). Demand shifts, population adjustments, and labor market outcomes during the 1980s. *Journal of Labor Economics*, 18(1), 20–54.
- Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). Generative AI at work. Tech. rep., National Bureau of Economic Research.
- Brynjolfsson, E., Mitchell, T., & Rock, D. (2018). What can machines learn and what does it mean for occupations and the economy? *AEA Papers and Proceedings*, 108, 43–47.
- Card, D., & Lemieux, T. (2001). Can falling supply explain the rising return to college for younger men? a cohort-based analysis. *Quarterly Journal of Economics*, 116(2), 705–746.
- Choi, J., Kuziemko, I., Washington, E. L., & Wright, G. (2024). Local economic and political effects of trade deals: Evidence from NAFTA. *American Economic Review*, 114(6), 1540–75.
- Compton, J., & Pollak, R. A. (2007). Why are power couples increasingly concentrated in large metropolitan areas? *Journal of Labor Economics*, 25(3), 475–512.
- Connor, D. S., Kemeny, T., & Storper, M. (2023). Frontier workers and the seedbeds of inequality and prosperity. *Journal of Economic Geography*, (p. lbad018).
- Council of Economic Advisors (2024). The 2024 Economic Report of the President.
- Dauth, W., Findeisen, S., Moretti, E., & Suedekum, J. (2022). Matching in cities. *Journal of the European Economic Association*, 20(4), 1478–1521.
- David, P. A. (1990). The dynamo and the computer: an historical perspective on the modern productivity paradox. *American Economic Review*, 80(2), 355–361.
- Desilver, D. (2022). The polarization in today’s Congress has roots that go back decades. *Pew Research Center*. 10 Mar 2022. Accessed 5 Jun 2024.  
URL <https://www.pewresearch.org/short-reads/2022/03/10/the-polarization-in-todays-congress-has-roots-that-go-back-decades/>
- Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). GPTs are GPTs: An early look at the labor market impact potential of large language models. *arXiv preprint arXiv:2303.10130*.
- Eriksson, K., Russ, K. N., Shambaugh, J. C., & Xu, M. (2021). Trade shocks and the shifting landscape of US manufacturing. *Journal of International Money and Finance*, 114, 102407.
- 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.
- Felten, E. W., Raj, M., & Seamans, R. (2018). A method to link advances in artificial intelligence to occupational abilities. *AEA Papers and Proceedings*, 108, 54–57.
- Finkel, E. J., Bail, C. A., Cikara, M., Ditto, P. H., Iyengar, S., Klar, S., Mason, L., McGrath, M. C., Nyhan, B., Rand, D. G., Skitka, L. J., Tucker, J. A., Van Bavel, J. J., Wang, C. S., & Druckman, J. N. (2020). Political sectarianism in america. *Science*, 370(6516), 533–536.
- Florida, R. (2002). *The rise of the creative class*. Basic books.
- Frank, T. (2016). *Listen, Liberal: Or, what ever happened to the party of the people?*. Macmillan.
- Freeman, R. (1976). *The overeducated American*. Academic press.
- Fry, R., Braga, D., & Parker, K. (2024). Is college worth it? *Pew Research Center*.  
URL <https://www.pewresearch.org/social-trends/2024/05/23/is-college-worth-it-2/>
- Fuller, J. B., Sigelman, M., & Martin, A. (2024). Skills-based hiring: The long road from pronouncements to practice.
- Ganong, P., & Shoag, D. (2017). Why has regional income convergence in the U.S. declined? *Journal of Urban Economics*, 102, 76–90.
- Ganzel, B. (2009). Farm boom of the 1970s. Accessed 7 Feb 2022.  
URL [https://livinghistoryfarm.org/farminginth70s/money\\_02.html](https://livinghistoryfarm.org/farminginth70s/money_02.html)
- Glaeser, E. L., & Gyourko, J. (2005). Urban decline and durable housing. *Journal of Political Economy*, 113(2), 345–375.
- Glaeser, E. L., Kolko, J., & Saiz, A. (2001). Consumer city. *Journal of Economic Geography*, 1(1), 27–50.
- Goworowska, J., & Gardner, T. K. (2012). *Historical migration of the young, single, and college educated: 1965 to 2000*. US Census Bureau, Population Division.
- Graefe, L. (2013). Oil shock of 1978–79. Accessed 7 Feb 2022.  
URL <https://www.federalreservehistory.org/essays/oil-shock-of-1978-79>
- Guriev, S., & Papaioannou, E. (2022). The political economy of populism. *Journal of Economic Literature*, 60(3), 753–832.
- Hershbein, B., & Stuart, B. A. (2024). The evolution of local labor markets after recessions.



- American Economic Journal: Applied Economics*, 16(3), 399–435.  
URL <https://www.aeaweb.org/articles?id=10.1257/app.20220132>
- Hill, F. (2021). *There Is Nothing For You Here: Finding Opportunity in the Twenty-First Century*. Mariner Books.
- Hirsch, B. T., & Macpherson, D. A. (2003). Union membership and coverage database from the Current Population Survey: Note. *ILR Review*, 56(2), 349–354.  
URL <http://unionstats.com>
- Holoubek, L. (2024). Foundation models, internet-scale data, and the path to generalized robots. *Personal Blog*. 2 Feb 2024. Accessed 5 Jun 2024.  
URL <https://lukasholoubek.com/foundation-models-internet-scale-data-path-to-generaliz>
- Jones, B. (2018). An examination of the 2016 electorate, based on validated voters. *Pew Research Center - U.S. Politics & Policy*.  
URL <https://www.pewresearch.org/politics/2018/08/09/an-examination-of-the-2016-electorate-based-on-validated-voters/>
- Kim, G., Merritt, C., & Peri, G. (2024). Measuring and predicting “new work” in the United States: The role of local factors and global shocks. Tech. rep., National Bureau of Economic Research.
- Kochan, T. A., Armstrong, B., Shah, J., Castilla, E. J., Likis, B., & Mangelsdorf, M. E. (2024). Bringing worker voice into generative AI. *MIT Sloan Research Report*.
- Levy, F. (1999). *The new dollars and dreams: American incomes in the late 1990s*. Russell Sage Foundation.
- Levy, F., & Temin, P. (2009). Institutions and wages in Post-World War II America. In *Labor in the Era of Globalization*, chap. 1. Cambridge University Press.
- Lin, J. (2011). Technological adaptation, cities, and new work. *Review of Economics and Statistics*, 93(2), 554–574.
- Margalit, Y. (2019). Political responses to economic shocks. *Annual Review of Political Science*, 22, 277–295.
- Moretti, E. (2012). *The new geography of jobs*. Houghton Mifflin Harcourt.
- Muro, M., Whiton, J., & Maxim, R. (2019). What jobs are affected by AI? better-paid, better-educated workers face the most exposure. *Brookings Metropolitan Policy Program Report*.
- Neff, J., Rives, K., & Wagner, J. (2002). NAFTA’s Impact on North Carolina, part 3:

- Leaders Scramble as Job Losses Mount. *Raleigh News and Observer*. 20 Aug 2002.
- Newman, L., & Skocpol, T. (2023). *Rust Belt Union Blues: Why Working-Class Voters Are Turning Away from the Democratic Party*. Columbia University Press.
- Nix, N., Zkrzewski, C., & De Vynck, G. (2024). Silicon Valley is pricing academics out of AI research. *Washington Post*. 12 Mar 2024. Accessed 10 Jun 2024.  
URL <https://www.washingtonpost.com/technology/2024/03/10/big-tech-companies-ai-research/>
- Notowidigdo, M. J. (2020). The incidence of local labor demand shocks. *Journal of Labor Economics*, 38(3), 687–725.
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192.
- Oliphant, J. B. (2023). Top tax frustrations for Americans: The feeling that some corporations, wealthy people don't pay fair share. *Pew Research Center*. 7 Apr 2023. Accessed 8 Jan 2025.  
URL <https://pewrsr.ch/3MIbaIX>
- Parker, K. (2019). The growing partisan divide in views of higher education. *Pew Research Center*.  
URL <https://pewrsr.ch/3mp2Q2k>
- Polanyi, K. (1944). *The Great Transformation*. Farrar & Rinehart.
- Pollard, E. (2019). A new approach to industry and occupation recoding in the CPS. *Federal Reserve Bank of Kansas City*. TB 19-02.  
URL <https://doi.org/10.18651/TB/TB1902>
- Rattner, N. (2024). AI talent is in demand as other tech job listings decline. *Wall Street Journal*. 7 Mar 2024. Accessed 27 Mar 2024.  
URL <https://www.wsj.com/tech/ai/ai-jobs-demand-tech-layoffs-5b7344c0>
- Rodríguez-Pose, A., Terrero-Dávila, J., & Lee, N. (2023). Left-behind versus unequal places: interpersonal inequality, economic decline and the rise of populism in the USA and Europe. *Journal of Economic Geography*, 23(5), 951–977.
- Somers, J. (2024). A revolution in how robots learn. *The New Yorker*. 25 Nov 2024.
- Stokel-Walker, C. (2024). OpenAI insiders' latest open letter is a warning against AI. will anyone listen? *Fast Company*. 4 Jun 2024. Accessed 5 Jun 2024.  
URL <https://www.fastcompany.com/91135744/openai-open-letter-fatigue>

- Sunstein, C. R. (2002). The law of group polarization. *Journal of Political Philosophy*, 10(2).
- Tiku, N. (2023). Doomsday to utopia: Meet AI's rival factions. *Washington Post*. 9 Apr 2023. Accessed 5 Jun 2024.  
URL <https://www.washingtonpost.com/technology/2023/04/09/ai-safety-openai/>
- U.S. Congress Joint Economic Committee (1986). The bi-coastal economy: Regional patterns of economic growth during the Reagan Administration: A staff study.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2023). Attention is all you need.
- Viscelli, S. (2016). *The Big Rig: Trucking and the Decline of the American Dream*. University of California Press.
- Webb, M. (2019). The impact of artificial intelligence on the labor market. *Available at SSRN 3482150*.
- Wolfram, S. (2023). What is ChatGPT doing ... and why does it work?  
URL <https://writings.stephenwolfram.com/2023/02/what-is-chatgpt-doing-and-why-does-it-work/>
- Yagan, D. (2019). Employment hysteresis from the Great Recession. *Journal of Political Economy*, 127(5), 2505–2558.