

AI AND THE FUTURE OF WORK IN GREATER METROWEST

AN INDEX OF OCCUPATIONAL
EXPOSURE AND LOCAL HIRING
DEMAND



MetroWest Economic
Research Center
FRAMINGHAM STATE UNIVERSITY



August 2026



**FRAMINGHAM
STATE
UNIVERSITY**

Copyright 2026 Framingham State University



**MetroWest Economic
Research Center**
FRAMINGHAM STATE UNIVERSITY

METROWEST ECONOMIC RESEARCH CENTER

Framingham State University
100 State Street, Framingham, MA 01701-9101
508-626-4033 | FAX: 508-626-4040
www.merc-fsu.org

***No material from this publication may be reproduced in whole
or in part without the permission of
MetroWest Economic Research Center (MERC),
Framingham State University.***

*MERC provides regional economic data.
MERC takes no position of policy issues at any level.*

Report by:

Hugh Qiu, PhD
Luis D. Rosero, PhD

With Research and Design Assistance From:

Yair Rachmany, BA Candidate
Rebecca Emken, BA
Ethan Langevin, BS
Ling Zhang, BA Candidate

TABLE OF CONTENTS

Executive Summary.....	1
Introduction and Literature Discussion.....	3
Methodology.....	5
1. AI Exposure Measures.....	5
2. Localizing AI Exposure to Greater MetroWest Communities.....	7
3. Measuring Local Labor Demand with Job-Posting Data.....	7
4. Posting-Weighted AI Exposure of Local Hiring Demand.....	8
5. Linking AI Exposure and Job-Posting Change.....	8
6. Occupation and Community Decomposition.....	9
7. Interpretation and Limitations.....	9
Summary of Findings.....	10
Aggregate Greater MetroWest Impact.....	11
AI Exposure and Local Hiring Demand.....	13
Community-by-Community Impact Narratives.....	15
Ashland Impact.....	15
Framingham Impact.....	16
Holliston Impact.....	17
Hopkinton Impact.....	18
Hudson Impact.....	19
Marlborough Impact.....	20
Natick Impact.....	21
Northborough Impact.....	22
Sherborn Impact.....	23
Southborough Impact.....	24
Sudbury Impact.....	25
Wayland Impact.....	26
Westborough Impact.....	27
Workforce and Planning Implications.....	28
Limitations and Additional Data Needs.....	29
Summary.....	30
References.....	31
Appendix: Python Code for Figure Generation.....	33

Executive Summary

The Greater MetroWest AI Index examines how occupational exposure to artificial intelligence intersects with local job-posting trends across the region. The report combines resident occupational exposure data with Lightcast job-posting data to identify where AI-related task change may matter most and whether local hiring demand is beginning to shift.

In Greater MetroWest, the post-pandemic hiring boom gave way not only to lower posting volumes but also to changes in the composition of hiring demand. Unique job postings rose from 64,051 in 2019 to 81,205 in 2022, then declined to 60,725 by 2025. That left regional postings 25.2% below the 2022 peak and 5.2% below the 2019 baseline. The decline reflects two forces: a normalization from unusually strong hiring conditions and a shift in the occupational mix of jobs employers are trying to fill.

The next question is whether this pullback was concentrated in communities with higher AI exposure. The evidence does not show a simple pattern. Across communities, the correlation between resident AI exposure and the 2022–2025 change in postings is 0.36. Higher-exposure communities were not mechanically associated with deeper posting declines. In several cases, smaller professional-resident communities proved more stable than larger employer centers, suggesting that community size, employer base, and sector mix matter as much as resident exposure.

The stronger AI-related signal appears not in which communities lost the most postings, but in how the composition of employer demand changed. The posting-weighted AI exposure of local hiring demand fell from 0.455 in 2019 to 0.389 in 2025. This suggests that Greater MetroWest employers posted relatively fewer jobs in some highly AI-exposed digital, professional, customer-service, marketing, finance, and administrative occupations, while healthcare, education, therapy, and local service roles became more prominent in the regional posting mix.

These results should therefore be read as an early-warning and workforce-planning signal, not as proof of AI-caused job loss. Interest rates, the unwinding of the pandemic hiring surge, sector-specific cycles, employer location decisions, remote work, and healthcare demand all shape local job postings. The value of the Greater MetroWest AI Index is that it helps identify where AI-related task change could matter most, where hiring demand is already shifting, and where local institutions may need to prepare workers and employers for adaptation.

The central takeaway is that Greater MetroWest does not face a uniform AI displacement story. It faces a workforce transition challenge: helping exposed office, analytical, customer-service, and professional workers adapt while supporting growth in healthcare, education, therapy, and other locally resilient sectors.




Metric	Value	Interpretation
2019 postings	64,051	Baseline year
2022 postings	81,205	Peak year in this data
2025 postings	60,725	-25.2% vs. 2022
Correlation: AI beta vs. 2022-2025 change	0.36	No simple negative community-level relationship
Aggregate posting-weighted demand AI exposure	0.389 in 2025	Down from 0.455 in 2019

Source: Lightcast, 2023 American Community Survey employment data, and MERC calculations.

Several patterns define the Greater MetroWest story:

- Largest post-2022 posting pullbacks in absolute terms: Framingham (-7,503), Marlborough (-5,490), Westborough (-2,598), and Natick (-1,942).
- Most resilient post-2022 posting trends: Ashland (+18.8%) expanded through 2025, while Southborough (-3.9%), Northborough (-3.9%), and Sherborn (-7.8%) held close to their 2022 levels.
- Largest aggregate occupation declines: software developers, computer occupations, laborers and material movers, retail salespersons, customer service representatives, clinical laboratory technologists, marketing managers, computer support specialists, industrial engineers, financial analysts, and accountants.
- Largest aggregate occupation gains: registered nurses, physical therapists, radiologic technologists, occupational therapists, physical therapist assistants, postsecondary teachers, clinical and counseling psychologists, lawyers, surgical technologists, and mental-health counselors.
- Main planning implication: Greater MetroWest should treat AI exposure as a signal for workforce transition and task redesign, not as a simple job-loss forecast.

The Greater MetroWest



INTRODUCTION AND LITERATURE DISCUSSION

Artificial intelligence is becoming a general-purpose technology for the local economy. Its effects will not stop at technology firms; AI is beginning to reshape office work, professional services, customer service, marketing, analytics, software development, education, healthcare administration, and municipal operations. The International Monetary Fund (IMF) frames generative AI as a broad labor-market shock with both productivity potential and displacement risk, especially in advanced economies where employment is concentrated in cognitive-intensive occupations (Cazzaniga et al., 2024).

The Greater MetroWest AI Index starts from a practical local question: which communities are more exposed to AI because of their occupational structure, and do local job postings show early signs of changing labor demand? This approach follows the occupation-to-geography logic in the AI exposure literature. Felten, Raj, and Seamans (2021) construct AI exposure at the occupation level and show how it can be aggregated into industry and geographic exposure measures. Their framework provides the conceptual foundation for building a community-level index from local occupational structure.

The Anthropic Economic Index is also useful because it separates automation from augmentation. In consumer Claude.ai use, automation and augmentation are relatively balanced. In enterprise API use, however, 77% of business uses involve automation patterns, compared with about 50% for Claude.ai users. This distinction matters locally because the same AI capability can have very different effects depending on how employers deploy it. AI can raise worker productivity, but it can also support workflow redesign that reduces demand for some tasks or positions (Appel et al., 2025).

Together, this body of research points to five lessons for interpreting the Greater MetroWest AI Index and the local job-posting evidence.

First, AI exposure is not the same as displacement. Occupational AI exposure measures whether AI can affect tasks within an occupation; it does not determine whether employers will reduce jobs, raise productivity, or change skill requirements. The Yale Budget Lab emphasizes that exposure metrics identify where AI could have an impact, not which jobs AI will automate out of existence. It also finds that exposure measures broadly agree on which occupations are exposed, but disagree more on the magnitude of exposure among highly exposed occupations (Gimbel, Kendall, and Kulsakdinun, 2026). Manning and Aguirre (2026) reach a similar conclusion: exposure reflects potential task change, not inevitable job loss.

Second, AI exposure differs sharply by occupation. The Yale Budget Lab finds higher exposure in computer and mathematical occupations, legal work, sales, office and administrative support, arts and media, business and financial operations, and management. Manual and physical occupations such as construction, maintenance, production, and building and grounds cleaning tend to show lower exposure (Gimbel, Kendall, and Kulsakdinun, 2026). This pattern supports using an occupation-group chart for Greater MetroWest to show which parts of the resident workforce are most exposed to AI-enabled task change.

Third, geography matters because occupations are not evenly distributed across places. Brookings argues that generative AI differs from earlier automation because its exposure is more concentrated in educated, office-oriented, and higher-wage labor markets. Earlier automation shocks were more concentrated in routine manual or blue-collar work, while generative AI is more relevant to research, writing, coding, analysis, marketing content, data synthesis, and presentation work (Muro, Methkuppally, and Kinder, 2025). Treasury researchers similarly find that workers in cities are more exposed to AI than rural workers, and that exposure is higher among workers with bachelor's degrees, STEM backgrounds, and higher-return college majors (Schendstok and Schreiner Wertz, 2024). These findings make local occupational structure central to understanding community-level AI exposure.

Fourth, job postings provide an early signal of labor demand and skill change. Acemoglu, Autor, Hazell, and Restrepo (2020) use online vacancy data to show that AI-exposed establishments changed the skills requested in job postings, dropping some previously posted skills while adding new ones. More recent World Bank research using 285 million Lightcast postings from 2018Q1 to 2025Q2 finds that postings declined more in occupations with higher AI-substitution vulnerability after ChatGPT's release, especially entry-level roles, administrative support, and professional services (Liu, Wang, and Yu, 2025). This supports the use of job postings as a timely demand-side measure, while recognizing that postings are not the same as employment, hires, or layoffs.

Fifth, vulnerability depends on adaptive capacity. High-exposure professional workers may have stronger transferable skills, savings, and job-transition options. Clerical and administrative workers may face higher adjustment costs if AI exposure turns into displacement. Manning and Aguirre (2026) estimate that many highly exposed workers have above-median adaptive capacity, but they also identify 6.1 million U.S. workers who are both highly exposed and low in adaptive capacity, concentrated in clerical and administrative roles. The Council of Economic Advisers similarly distinguishes between AI exposure and AI vulnerability, arguing that workers are more vulnerable when high exposure is combined with lower job-performance requirements and weaker evidence of occupational adaptation (Council of Economic Advisers, 2024).

Recent survey evidence suggests that AI adoption is already becoming part of the Greater MetroWest workplace. In a 2026 MetroWest Economic Research Center survey of regional employers, most respondents reported permitting AI use for work-related tasks, particularly for writing and communication, research, and administrative functions. At the same time, relatively few employers reported plans to replace staff positions with AI or automated processes. These responses suggest that local adoption is currently focused more on productivity enhancement than direct labor substitution (MERC, 2026).

This report applies those lessons to Greater MetroWest by combining two complementary views of the local labor market. The first is a resident workforce exposure measure, which identifies where local residents work in occupations more exposed to AI-enabled task change. The second is a demand-side job-posting analysis, which examines how local employer hiring demand changed from 2019 to 2025 across communities and occupations. The evidence points to areas where workforce adaptation may be most important. This report does not claim that AI has already reduced employment in Greater MetroWest. Instead, it identifies where occupational exposure is concentrated, where hiring demand is changing, which communities may face greater adjustment pressure, and what workforce strategies could help the region adapt. The key question is not simply which jobs will disappear, but which workers, firms, and communities are prepared to use AI productively and which may need targeted support to adapt.



Methodology

This report connects national AI occupational exposure measures with local Greater MetroWest labor-market data. The workflow is straightforward: assign occupation-level AI exposure scores to local occupations, weight those scores by each community’s resident occupational structure, and then use Lightcast job-posting data to examine whether local hiring demand is shifting toward or away from more AI-exposed occupations.

The result is an exposure and early-warning framework, not a causal estimate of AI-related job loss. The analysis identifies where AI could affect local work tasks, where employer demand is changing, and which communities may face greater workforce adjustment pressure.

1. AI EXPOSURE MEASURES

The AI exposure scores used in this report are based on the task-based methodology developed by Eloundou, Manning, Mishkin, and Rock (2024), which evaluates whether large language models and related tools could reduce the time required to complete occupational tasks without reducing output quality. This task-based approach is consistent with the broader labor-market literature, which treats jobs as bundles of tasks rather than single, indivisible activities (Acemoglu, Autor, Hazell, and Restrepo, 2020; Eloundou et al., 2024).

The exposure framework distinguishes three levels of AI integration:

Measure	Interpretation	Usefulness
Alpha	Exposure to a general-purpose language model, such as ChatGPT, without additional tools.	Captures direct AI exposure.
Beta	Exposure to a language model with some complementary tools, software, or workflow integration.	Main index used in this report.
Gamma	Exposure assuming more complete AI integration with customized software and firm-specific tools.	Captures upper-bound exposure.

Source: Eloundou, Manning, Mishkin, and Rock (2024).

To estimate occupational AI exposure, this report uses the beta AI exposure measure as the main index because it represents a realistic middle case. Alpha may understate AI’s impact because many employers are already embedding AI into software and workflows. Gamma may overstate near-term exposure because it assumes more complete customization and firm-specific integration. Beta therefore provides a practical estimate of near- to medium-term exposure.

The report also distinguishes between DV ratings and human ratings. DV ratings are generated through the GPT-4-assisted evaluation process in Eloundou et al. (2024), while human ratings come from human annotators. The Greater MetroWest dataset compares both, but this report uses DV beta exposure as the primary measure because it provides a standardized occupation-level score that can be consistently applied across local occupational categories.

To translate national AI exposure estimates into local labor-market insights, the analysis connects occupation-level exposure measures with Greater MetroWest workforce and hiring data. The workflow is straightforward: assign occupation-level AI exposure scores to local occupations, weight those scores by each community’s resident occupational structure, and then use Lightcast job-posting data to examine whether local hiring demand is shifting toward or away from more AI-exposed occupations.

The result is an exposure and early-warning framework, not a causal estimate of AI-related job loss. The analysis identifies where AI could affect local work tasks, where employer demand is changing, and which communities may face greater workforce adjustment pressure.

Greater MetroWest AI Exposure by Major Occupation Group AI exposure Beta Measure




Figure 1: Greater MetroWest AI Exposure by Major Occupation Group

Source: MERC calculations using O*NET occupational classifications, 2023 American Community Survey employment data, and DV beta AI exposure scores from Eloundou et al. (2024).

Figure 1 shows that AI exposure is uneven across major occupation groups in Greater MetroWest. Exposure is highest in occupations that rely heavily on cognitive, analytical, communication, administrative, and digital tasks, including office and administrative support, computer and mathematical occupations, sales, business and finance, legal work, architecture and engineering, and management. Occupations that involve more physical, manual, or in-person work, such as construction, installation and repair, production, transportation, food service, grounds maintenance, and agriculture, show lower exposure.

This pattern is central to the report. A community's AI exposure score is not determined simply by whether it has technology companies. It is driven by the occupational structure of its residents. Communities with larger shares of residents in professional, managerial, analytical, computer-related, financial, sales, and administrative roles will naturally have higher resident AI exposure. Communities with more workers in manual, service, transportation, production, or in-person care roles will tend to have lower exposure.

The chart also reinforces the main interpretive caution: AI exposure is not automatic job loss. It measures where AI can affect tasks. Some exposed occupations may face substitution pressure, especially where tasks are routine and easy to automate, while others may see productivity gains through AI augmentation. The report therefore uses Figure 1 as a foundation for measuring exposure, then turns to local job-posting data to examine whether hiring demand is actually shifting across Greater MetroWest communities.

2. Localizing AI Exposure to Greater MetroWest Communities

To convert national AI exposure scores into a local Greater MetroWest AI Index, this report links occupation-level AI exposure scores to local occupational employment patterns. The community-level exposure file uses 2023 American Community Survey employment data for 13 Greater MetroWest communities and maps residents into broad occupational categories.

A practical challenge is that local demographic data often report broad categories, such as “Computer and Mathematical Occupations” or “Office and Administrative Support Occupations,” while AI exposure scores are available at the more detailed O*NET-SOC occupation level, such as “Software Developers” or “Customer Service Representatives.”

To bridge this mismatch, the methodology uses an occupational crosswalk. Detailed O*NET-SOC occupations are grouped into broader occupation families using SOC code prefixes, and the AI exposure score for each broad group is calculated as the average of the detailed occupations within that group.

The community-level AI exposure score is calculated as an employment-weighted average:

$$\text{Community AI Exposure}_i = \sum_o \text{EmploymentShare}_{i,o} \times \text{AIExposure}_o$$

$$\text{EmploymentShare}_{i,o} = \text{Employment}_{i,o} / \text{TotalEmployment}_i$$

In this formula, i represents the community and o represents the occupation group. A community receives a higher AI exposure score when a larger share of its residents work in high-exposure occupations, such as computer and mathematical occupations, sales, or office and administrative support.

This measure should be interpreted as resident occupational exposure, not employer AI adoption. It tells us how exposed a community's resident workforce is based on the occupations residents hold. It does not show whether local employers are using AI, whether residents work inside the same community, or whether workers have already been displaced.

3. Measuring Local Labor Demand with Job-Posting Data

To connect AI exposure with local economic change, the report uses Lightcast Q2 2026 all-occupation job-posting data for Greater MetroWest from 2019 through 2025. The data report annual unique postings by community and occupation.

The job-posting data are used in two ways. First, total annual postings by community measure each community's overall hiring trend. Second, occupation-level postings identify which roles are growing or declining within each community. This allows the report to ask whether local hiring changes are concentrated in more AI-exposed or less AI-exposed occupations.

Because postings are employer-side flow measures, they should not be read as employment levels, hires, layoffs, or job quality. They are most useful as a timely indicator of changing demand. The occupation-level analysis uses all occupations listed in the Lightcast export, while the posting-weighted AI exposure measure depends on successful matching between SOC occupation codes and AI exposure scores.

4. Posting-Weighted AI Exposure of Local Hiring Demand

The community AI exposure index measures the occupational exposure of residents. The job-posting analysis measures the exposure of local hiring demand. These concepts are related, but they are not the same. A community may have highly exposed residents while local employers post jobs concentrated in less exposed occupations, or vice versa. The distinction matters because job postings are assigned by employer or job location, whereas the resident AI exposure index is based on where workers live.

To measure the AI exposure of local hiring demand, the report calculates a posting-weighted exposure score:

$$\text{Demand AI Exposure}_{i,t} = \sum_o \text{PostingShare}_{i,o,t} \times \text{AIExposure}_o$$

$$\text{PostingShare}_{i,o,t} = \text{Postings}_{i,o,t} / \text{TotalPostings}_{i,t}$$

In these formulas, i represents the community, t represents the year, and o represents the occupation. $\text{PostingShare}_{i,o,t}$ is the share of all job postings in community i and year t that belong to occupation o . AIExposure_o is the AI beta exposure score assigned to occupation o . The demand AI exposure score is therefore a weighted average of occupational AI exposure, where occupations with more local job postings receive more weight.

This measure answers a demand-side question: are the jobs being posted in a community increasingly concentrated in high-AI-exposure occupations or in lower-exposure occupations?

A decline in posting-weighted demand exposure may indicate that local hiring is shifting away from more AI-exposed professional, digital, or administrative occupations. An increase may indicate that employers are still posting heavily in AI-exposed roles, possibly because AI is complementing workers rather than replacing them.

5. Linking AI Exposure and Job-Posting Change

The report uses descriptive correlation analysis to examine whether communities with higher resident AI exposure experienced different posting trends after 2022, the year ChatGPT became widely available. The main comparison is:

$$\text{PctChangePostings}_{i,2022-2025} = (\text{Postings}_{i,2025} - \text{Postings}_{i,2022}) / \text{Postings}_{i,2022}$$

In this formula, i represents the community. $\text{Postings}_{i,2022}$ is the number of unique job postings in community i in 2022, and $\text{Postings}_{i,2025}$ is the number of unique job postings in the same community in 2025. The formula measures the percentage change in job postings from 2022 to 2025. A negative value means postings declined after 2022, while a positive value means postings increased.

This change is compared with each community's AI beta exposure score. The analysis is descriptive and should not be interpreted as causal. A negative relationship would suggest that more exposed communities had weaker posting recovery or larger declines after 2022. A weak or positive relationship would suggest that AI exposure alone does not explain local hiring changes.

This caution is important because job postings are shaped by many factors beyond AI, including post-pandemic normalization, interest rates, healthcare demand, employer location decisions, remote work, demographic change, and community-specific industry mix.

6. Occupation and Community Decomposition

To make the index interpretable, the report decomposes each community's exposure by occupation group. This shows which occupations contribute most to a community's AI exposure, such as management, computer and mathematical occupations, business and financial operations, sales, or office and administrative support.

This decomposition matters because two communities may have similar overall AI exposure scores but different adjustment profiles. A community whose exposure is driven by management and professional occupations may have more adaptive capacity. A community whose exposure is driven by clerical, administrative, or customer-service roles may face greater adjustment risk if AI adoption reduces demand for routine office tasks.

7. Interpretation and Limitations

The Greater MetroWest AI Index should be interpreted as an exposure and early-warning tool, not as a direct measure of job loss.

Data source	Use in this report	Important limitation
Community-level AI Index	Measures resident occupational exposure to AI by community using occupation weights and beta exposure scores.	Measures exposure to task change, not actual AI adoption or job loss.
2023 ACS / local occupation structure	Provides the employment weights used to localize national AI exposure scores.	Occupations are reported in broad categories, requiring crosswalk mapping from detailed O*NET-SOC occupations.
O*NET-SOC occupation mapping	Links detailed occupation-level AI exposure scores to broader local occupation groups.	Averaging across detailed occupations can hide variation within broad occupation groups.
Community Occupations all-occupation postings, 2019–2025	Measures annual job postings by community and occupation using Lightcast data.	Occupation-level analysis uses the all-occupation Lightcast export, but it still measures postings rather than employment, hires, wages, layoffs, or job quality.
Posting-weighted demand exposure	Measures whether local hiring demand is concentrated in high- or low-exposure occupations.	Depends on successful matching between job titles and occupation-level AI exposure scores.
Correlation analysis	Compares community AI exposure with job-posting changes after 2022.	Descriptive only; does not establish causality.

In summary, the methodology connects three questions:

- Resident exposure: Which Greater MetroWest communities have residents working in more AI-exposed occupations?
- Hiring-demand exposure: Are local job postings concentrated in higher-exposure or lower-exposure occupations?
- Adjustment pressure: Do communities with higher resident exposure show different job-posting trends after 2022?

This structure moves the report beyond a simple community ranking. It shows how AI exposure enters the local economy through resident occupational structure, how employer demand is changing, and which communities may need targeted workforce-development strategies.

Summary of Findings

Community	AI beta	2025 postings	2019-2025	2022-2025	Peak year
Ashland	0.429	2,279	+81.4%	+18.8%	2025
Framingham	0.363	14,971	-14.2%	-33.4%	2022
Holliston	0.427	1,311	+29.8%	-24.3%	2021
Hopkinton	0.463	2,090	-29.2%	-33.2%	2021
Hudson	0.396	2,863	+26.0%	-23.4%	2022
Marlborough	0.36	12,956	-8.2%	-29.8%	2022
Natick	0.435	7,998	-17.7%	-19.5%	2022
Northborough	0.444	2,813	+24.7%	-3.9%	2021
Sherborn	0.461	369	+87.3%	-7.8%	2021
Southborough	0.446	2,610	+24.2%	-3.9%	2021
Sudbury	0.465	2,294	+34.3%	-12.1%	2021
Wayland	0.455	1,361	+4.9%	-23.0%	2021
Westborough	0.437	6,810	-11.9%	-27.6%	2021

Source: Lightcast, 2023 American Community Survey employment data, and MERC calculations.

The table shows that Greater MetroWest is not experiencing one uniform labor-market adjustment. Framingham, Marlborough, Natick, and Westborough remain the largest job-posting markets, but the largest post-2022 pullbacks occurred in Framingham, Hopkinton, Marlborough, and Westborough. Ashland is the clear exception, expanding through 2025, while Southborough and Northborough remained close to their 2022 levels.

The resident AI exposure ranking does not directly predict the post-2022 posting decline. Some higher-exposure communities were relatively resilient, while several lower-exposure large employer centers experienced sharper cyclical pullbacks. The implication is that community size, employer base, sector mix, and the post-pandemic hiring cycle are essential for interpreting AI exposure.

Aggregate Greater MetroWest Impact




Figure 2: Aggregate Greater MetroWest job postings, 2019-2025.

Figure 2 captures the region's hiring cycle. Greater MetroWest began with 64,051 postings in 2019, climbed to 81,205 in 2022, and fell to 60,725 in 2025. The 25.2% decline from the 2022 peak is the dominant aggregate pattern in the report.

The more important story is not only the size of the pullback, but its occupational composition. Several highly AI-exposed digital, analytical, finance, marketing, customer-service, and administrative roles fell sharply, while healthcare, therapy, education, and in-person service roles gained.






Figure 3: Largest all-occupation posting declines from 2022 to 2025.




Figure 4: Largest all-occupation posting gains from 2022 to 2025.

Figures 3 and 4 show the reallocation underneath the aggregate decline. Software developers alone declined by 2,068 postings from 2022 to 2025. Computer occupations, material movers, retail salespersons, customer service representatives, and several finance, marketing, engineering, and administrative roles also declined. The largest gains were concentrated in registered nursing, physical therapy, radiology, occupational therapy, postsecondary teaching, psychology, law, surgical technology, and mental-health counseling. The region is not simply posting fewer jobs; it is posting a different mix of jobs.

AI Exposure and Local Hiring Demand




Figure 5: Resident AI exposure index by community.

Figure 5 shows that Sudbury, Hopkinton, Sherborn, Wayland, Southborough, and Northborough have the highest resident AI exposure in the sample. This reflects where residents work by occupation, not necessarily the types of jobs posted by employers located in each community.




Figure 6: Resident AI exposure compared with post-2022 job-posting change.

Figure 6 links resident AI exposure to post-2022 posting change. The scatterplot does not show a simple negative relationship. The correlation between resident AI exposure and the 2022-2025 posting change is 0.36. In practical terms, some higher-exposure communities are smaller professional-resident communities with relatively stable local posting bases, while some lower-exposure communities are large employer centers that experienced larger cyclical pullbacks. This is a simple bivariate correlation and does not control for macroeconomic conditions, interest rates, sector mix, employer decisions, population change, or commuting patterns.




Figure 7: Posting-weighted AI exposure of local hiring demand.

Figure 7 provides the clearest demand-side signal. The posting-weighted AI exposure of the local hiring mix declined from 0.455 in 2019 to 0.389 in 2025. This pattern is consistent with fewer postings in software, data, marketing, customer service, finance, and administrative occupations, alongside stronger demand in healthcare, education, therapy, and local services. The pattern is descriptive rather than causal, but it helps identify where the hiring mix is moving.



Community-By-Community Impact Narratives

The community narratives below interpret each community's posting trend and occupation-level changes. Each graph shows total annual job postings from 2019 to 2025. The discussion focuses on the post-2022 period because 2022 is the aggregate peak in the dataset and closely precedes the widespread public adoption of generative AI tools.

ASHLAND IMPACT

Ashland is the region's clearest growth outlier. While most communities cooled after 2022, Ashland rose from 1,918 postings in 2022 to 2,279 in 2025, an 18.8% increase. Compared with 2019, postings were up 81.4%, and 2025 was the community's peak year in the dataset.

Its resident AI exposure score of 0.429 places Ashland 9th of 13 communities, a middle-range profile. On the employer-demand side, posting-weighted AI exposure was 0.301 in 2025, down 0.062 from 2019. The growth in postings therefore did not come mainly from high-AI-exposure office or digital roles.

The occupation table explains the difference. Ashland lost postings in retail supervision, material moving, software development, real estate, and driver/sales work. But these declines were more than offset by gains in etchers and engravers, registered nurses, coaches, welding-related roles, and production work.

Measure	Value	Note
Resident AI beta	0.429	Rank 9 of 13
2025 postings	2,279	Rank 9 by size
2019-2025 change	+81.4%	1,023
2022-2025 change	+18.8%	361
Peak year	2025	2,279 postings

Year	Total postings	Index (2019=100)	Demand AI exposure
2019	1,256	100.0	0.363
2020	1,406	111.9	0.334
2021	1,913	152.3	0.292
2022	1,918	152.7	0.371
2023	1,633	130.0	0.325
2024	1,980	157.6	0.288
2025	2,279	181.4	0.301

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

Ashland's recent strength appears to be rooted in specialized production, healthcare, recreation, and hands-on local demand rather than a revival of digital hiring.

Signal	Position	AI beta	2022 Job Listings	2025 Job Listings	Change in Job Listings
Decline	First-Line Supervisors of Retail Sales Workers	0.486	66	27	-39
Decline	Laborers and Freight, Stock, and Material Movers, Hand	0.042	54	17	-37
Decline	Software Developers	0.868	35	2	-33
Decline	Real Estate Sales Agents	0.438	27	5	-22
Decline	Driver/Sales Workers	0.333	30	10	-20
Gain	Etchers and Engravers	0.136	1	149	+148
Gain	Registered Nurses	0.380	48	148	+100
Gain	Coaches and Scouts	0.300	23	78	+55
Gain	Welding, Soldering, and Brazing Machine Setters, Operators, and Tenders	0.068	0	42	+42
Gain	Production Workers, All Other	0.055	15	56	+41

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

FRAMINGHAM IMPACT

Framingham remains the region's largest job-posting center, but it also absorbed the largest absolute pullback. Postings fell from 22,474 in 2022 to 14,971 in 2025, a 33.4% decline. The 2025 level was also 14.2% below 2019, indicating that the post-pandemic cooling brought Framingham below its pre-boom baseline.

Framingham's resident AI exposure score is 0.363, ranking 12th of 13 communities. That low resident score reflects the occupational structure of people living in Framingham, not the complexity of the employer base located there. Employer demand tells a different story: the posting-weighted AI exposure of listed hiring demand was 0.399 in 2025, down 0.053 from 2019.

Measure	Value	Note
Resident AI beta	0.363	Rank 12 of 13
2025 postings	14,971	Rank 1 by size
2019-2025 change	-14.2%	-2,474
2022-2025 change	-33.4%	-7,503
Peak year	2022	22,474 postings

Year	Total postings	Index (2019=100)	Demand AI exposure
2019	17,445	100.0	0.452
2020	16,894	96.8	0.439
2021	20,203	115.8	0.427
2022	22,474	128.8	0.434
2023	16,358	93.8	0.415
2024	15,333	87.9	0.419
2025	14,971	85.8	0.399

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

The occupational table shows a broad reallocation inside the region's largest market. Losses were not confined to one sector: software developers, computer occupations, clinical laboratory technologists, retail salespersons, and production workers all declined. Gains in physical therapy, postsecondary teaching, automotive service, law, and radiology show that Framingham is not simply contracting. Its hiring base is shifting away from parts of digital, lab, retail, and production demand toward healthcare, education, and local professional services.

Signal	Position	AI beta	2022 Job Listings	2025 Job Listings	Change in Job Listings
Decline	Software Developers	0.868	649	180	-469
Decline	Computer Occupations, All Other	0.736	382	146	-236
Decline	Clinical Laboratory Technologists and Technicians	0.232	291	61	-230
Decline	Retail Salespersons	0.342	510	348	-162
Decline	Production Workers, All Other	0.055	226	72	-154
Gain	Physical Therapists	0.191	96	224	+128
Gain	Postsecondary Teachers	0.486	82	204	+122
Gain	Automotive Service Technicians and Mechanics	0.082	103	145	+42
Gain	Lawyers	0.425	62	104	+42
Gain	Radiologic Technologists and Technicians	0.377	97	138	+41

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

HOLLISTON IMPACT

Holliston's current posting level remains above 2019, but the community has cooled meaningfully from its earlier peak. Postings reached 1,915 in 2021, fell to 1,732 in 2022, and then declined to 1,311 in 2025. That leaves Holliston 24.3% below 2022 but still 29.8% above 2019.

The community's resident AI exposure score is 0.427, ranking 10th of 13. Its posting-weighted demand exposure was 0.321 in 2025, only 0.021 below 2019, suggesting that the AI-exposure profile of local hiring demand has been relatively stable.

Measure	Value	Note
Resident AI beta	0.427	Rank 10 of 13
2025 postings	1,311	Rank 12 by size
2019-2025 change	+29.8%	301
2022-2025 change	-24.3%	-421
Peak year	2021	1,915 postings

Year	Total postings	Index (2019=100)	Demand AI exposure
2019	1,010	100.0	0.342
2020	1,455	144.1	0.311
2021	1,915	189.6	0.310
2022	1,732	171.5	0.355
2023	1,511	149.6	0.302
2024	1,213	120.1	0.334
2025	1,311	129.8	0.321

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

Holliston's adjustment is less about a digital pullback and more about shifts in education, production, logistics, and healthcare support. Declines were concentrated in teaching assistants, secondary school teachers, special education teachers, education administrators, and production workers. Gains in material moving, shipping and receiving, cost estimating, physical therapy, and retail supervision point to a more mixed local economy where logistics, applied business support, healthcare, and local management have become more visible.

Signal	Position	AI beta	2022 Job Listings	2025 Job Listings	Change in Job Listings
Decline	Teaching Assistants, Except Postsecondary	0.367	70	23	-47
Decline	Secondary School Teachers, Except Special and Career/Technical Education	0.333	57	14	-43
Decline	Production Workers, All Other	0.055	55	15	-40
Decline	Special Education Teachers, All Other	0.475	41	20	-21
Decline	Education Administrators, Kindergarten through Secondary	0.418	30	13	-17
Gain	Laborers and Freight, Stock, and Material Movers, Hand	0.042	33	62	+29
Gain	Shipping, Receiving, and Inventory Clerks	0.500	25	46	+21
Gain	Cost Estimators	0.500	1	15	+14
Gain	Physical Therapists	0.191	2	15	+13
Gain	First-Line Supervisors of Retail Sales Workers	0.486	11	23	+12

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

HOPKINTON IMPACT

Hopkinton combines high resident AI exposure with one of the sharpest posting declines. Postings fell from 3,128 in 2022 to 2,090 in 2025, a 33.2% decline. Compared with the 2021 peak of 3,307, the 2025 level was down 36.8%.

Measure	Value	Note
Resident AI beta	0.463	Rank 2 of 13
2025 postings	2,090	Rank 10 by size
2019-2025 change	-29.2%	-863
2022-2025 change	-33.2%	-1,038
Peak year	2021	3,307 postings

Its resident AI exposure score is 0.463, ranking 2nd of 13 communities. The demand-side shift is also pronounced: posting-weighted AI exposure fell to 0.378 in 2025, down 0.155 from 2019. That is one of the clearest declines in demand-side exposure in the report.

Year	Total postings	Index (2019=100)	Demand AI exposure
2019	2,953	100.0	0.533
2020	2,446	82.8	0.442
2021	3,307	112.0	0.406
2022	3,128	105.9	0.482
2023	1,958	66.3	0.392
2024	1,653	56.0	0.379
2025	2,090	70.8	0.378

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

Hopkinton illustrates how a high-exposure professional labor market can cool after the post-pandemic hiring boom. Software developers account for the largest drop, and the decline also includes marketing managers, computer occupations, industrial engineers, and customer service representatives. The gains are smaller and more community-facing: purchasing, physical therapy, stock and order fulfillment, home care, and food service. The local hiring mix has shifted away from digital and corporate functions toward support, care, and household-service demand.

Signal	Position	AI beta	2022 Job Listings	2025 Job Listings	Change in Job Listings
Decline	Software Developers	0.868	341	48	-293
Decline	Marketing Managers	0.500	99	15	-84
Decline	Computer Occupations, All Other	0.736	87	19	-68
Decline	Industrial Engineers	0.554	63	23	-40
Decline	Customer Service Representatives	0.568	46	7	-39
Gain	Buyers and Purchasing Agents	0.489	11	31	+20
Gain	Physical Therapists	0.191	0	19	+19
Gain	Stockers and Order Fillers	0.182	6	23	+17
Gain	Home Health and Personal Care Aides	0.150	28	43	+15
Gain	Fast Food and Counter Workers	0.097	21	35	+14

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

HUDSON IMPACT

Hudson's postings remain above the 2019 baseline, but the community gave back a sizable share of its post-pandemic increase. Postings rose to 3,738 in 2022 and then fell to 2,863 in 2025, a 23.4% decline from the peak but still 26.0% above 2019.

Hudson has one of the lower resident AI exposure scores in the sample, 0.396, ranking 11th of 13. Its posting-weighted demand exposure was 0.337 in 2025, down 0.052 from 2019, indicating a moderate move away from higher-exposure listed roles.

Measure	Value	Note
Resident AI beta	0.396	Rank 11 of 13
2025 postings	2,863	Rank 5 by size
2019-2025 change	+26.0%	590
2022-2025 change	-23.4%	-875
Peak year	2022	3,738 postings

Year	Total postings	Index (2019=100)	Demand AI exposure
2019	2,273	100.0	0.389
2020	2,825	124.3	0.378
2021	3,584	157.7	0.334
2022	3,738	164.5	0.394
2023	2,784	122.5	0.351
2024	2,710	119.2	0.339
2025	2,863	126.0	0.337

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

The occupation table shows that Hudson's pullback was broader than AI-exposed office work. Software developers, electrical engineers, industrial engineers, material movers, and production workers all declined, pointing to weakness across both technical and industrially related hiring. The gains were concentrated in radiology, MRI technology, physical therapy, food service, and operations management, suggesting a shift toward healthcare services and day-to-day local operations.

Signal	Position	AI beta	2022 Job Listings	2025 Job Listings	Change in Job Listings
Decline	Software Developers	0.868	116	6	-110
Decline	Electrical Engineers	0.562	71	3	-68
Decline	Laborers and Freight, Stock, and Material Movers, Hand	0.042	73	25	-48
Decline	Industrial Engineers	0.554	61	14	-47
Decline	Production Workers, All Other	0.550	83	38	-45
Gain	Radiologic Technologists and Technicians	0.377	0	30	+30
Gain	Fast Food and Counter Workers	0.097	26	49	+23
Gain	Physical Therapists	0.191	6	29	+23
Gain	Magnetic Resonance Imaging Technologists	0.322	0	22	+22
Gain	General and Operations Managers	0.481	28	46	+18

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

MARLBOROUGH IMPACT

Marlborough is the region's second-largest job-posting market and a major contributor to the aggregate decline. Postings fell from 18,446 in 2022 to 12,956 in 2025, a 29.8% decline. The 2025 level was also 8.2% below 2019.

Measure	Value	Note
Resident AI beta	0.360	Rank 13 of 13
2025 postings	12,956	Rank 2 by size
2019-2025 change	-8.2%	-1,151
2022-2025 change	-29.8%	-5,490
Peak year	2022	18,446 postings

Marlborough's resident AI exposure score is 0.360, the lowest in the sample. This does not mean local employer demand is low exposure; rather, it means the resident workforce has a lower occupational exposure profile. Employer demand remained more exposed than the resident score, with posting-weighted demand exposure at 0.430 in 2025, though that was 0.069 below 2019.

Year	Total postings	Index (2019=100)	Demand AI exposure
2019	14,107	100.0	0.499
2020	13,775	97.6	0.482
2021	17,383	123.2	0.450
2022	18,446	130.8	0.457
2023	13,623	96.6	0.437
2024	13,221	93.7	0.436
2025	12,956	91.8	0.430

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

Marlborough's table is important because it combines large scale with a clear occupational shift. The largest losses came from software developers, computer occupations, material movers, accountants and auditors, and financial managers. The strongest gains were in physical therapy, registered nursing, occupational therapy, food service, and HVAC. The pattern suggests that healthcare, local services, and skilled hands-on work remained resilient while corporate, digital, finance, and logistics-related hiring contracted.

Signal	Position	AI beta	2022 Job Listings	2025 Job Listings	Change in Job Listings
Decline	Software Developers	0.868	679	245	-434
Decline	Computer Occupations, All Other	0.736	702	337	-365
Decline	Laborers and Freight, Stock, and Material Movers, Hand	0.042	270	80	-190
Decline	Accountants and Auditors	0.560	253	83	-170
Decline	Financial Managers	0.485	252	105	-147
Gain	Physical Therapists	0.191	27	161	+134
Gain	Registered Nurses	0.380	525	632	+107
Gain	Fast Food and Counter Workers	0.097	78	119	+41
Gain	Occupational Therapists	0.289	9	50	+41
Gain	Heating, Air Conditioning, and Refrigeration Mechanics and Installers	0.105	29	66	+37

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

NATICK IMPACT

Natick shows a two-stage adjustment: a modest post-pandemic rise followed by a meaningful cooling. Postings peaked at 9,940 in 2022 and fell to 7,998 in 2025, a 19.5% decline. Compared with 2019, postings were down 17.7%.

Measure	Value	Note
Resident AI beta	0.435	Rank 8 of 13
2025 postings	7,998	Rank 3 by size
2019-2025 change	-17.7%	-1,725
2022-2025 change	-19.5%	-1,942
Peak year	2022	9,940 postings

Natick's resident AI exposure score is 0.435, ranking 8th of 13. Its posting-weighted demand exposure was 0.377 in 2025, down 0.060 from 2019, indicating a notable shift away from higher-exposure listed roles.

Year	Total postings	Index (2019=100)	Demand AI exposure
2019	9,723	100.0	0.437
2020	8,253	84.9	0.442
2021	9,055	93.1	0.369
2022	9,940	102.2	0.410
2023	7,709	79.3	0.400
2024	7,378	75.9	0.387
2025	7,998	82.3	0.377

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

The most visible adjustment is the contrast between software and retail weakness on one side and healthcare growth on the other. Software developers declined sharply, and retail salespersons, material movers, retail supervisors, and teaching assistants also fell. At the same time, registered nurses, physical therapists, occupational therapists, radiologic technologists, and nurse practitioners increased. Healthcare is the clearest source of new posting momentum in Natick's recent hiring mix.

Signal	Position	AI beta	2022 Job Listings	2025 Job Listings	Change in Job Listings
Decline	Software Developers	0.868	480	59	-421
Decline	Retail Salespersons	0.342	964	700	-264
Decline	Laborers and Freight, Stock, and Material Movers, Hand	0.042	121	30	-91
Decline	First-Line Supervisors of Retail Sales Workers	0.486	632	546	-86
Decline	Teaching Assistants, Except Postsecondary	0.367	162	85	-77
Gain	Registered Nurses	0.380	283	685	+402
Gain	Physical Therapists	0.191	27	160	+133
Gain	Occupational Therapists	0.289	4	63	+59
Gain	Radiologic Technologists and Technicians	0.377	23	81	+58
Gain	Nurse Practitioners	0.481	27	60	+33

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

NORTHBOROUGH IMPACT

Northborough is one of the more stable communities after 2022. Postings were 2,928 in 2022 and 2,813 in 2025, a decline of only 3.9%. Relative to 2019, postings were still up 24.7%.

Its resident AI exposure score is 0.444, ranking 6th of 13. The posting-weighted demand exposure score was 0.363 in 2025, only 0.030 below 2019, which suggests a moderate rather than dramatic shift in the exposure profile of local hiring.

Measure	Value	Note
Resident AI beta	0.444	Rank 6 of 13
2025 postings	2,813	Rank 6 by size
2019-2025 change	+24.7%	558
2022-2025 change	-3.9%	-115
Peak year	2021	3,582 postings

Year	Total postings	Index (2019=100)	Demand AI exposure
2019	2,255	100.0	0.393
2020	2,649	117.5	0.352
2021	3,582	158.8	0.297
2022	2,928	129.8	0.349
2023	2,788	123.6	0.337
2024	2,745	121.7	0.354
2025	2,813	124.7	0.363

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

Northborough's stability masks a reshuffling within the local job mix. Losses were concentrated in material moving, driver/sales work, automotive repair, diesel mechanics, and packaging, pointing to some cooling in logistics, transportation, and repair-related support roles. Gains in retail sales, merchandising, project management, physical therapy, and animal care show a more service- and coordination-oriented pattern of demand.

Signal	Position	AI beta	2022 Job Listings	2025 Job Listings	Change in Job Listings
Decline	Laborers and Freight, Stock, and Material Movers, Hand	0.042	208	131	-77
Decline	Driver/Sales Workers	0.333	37	15	-22
Decline	Automotive Service Technicians and Mechanics	0.082	31	10	-21
Decline	Bus and Truck Mechanics and Diesel Engine Specialists	0.000	29	9	-20
Decline	Packers and Packers, Hand	0.100	36	17	-19
Gain	Retail Salespersons	0.342	126	157	+31
Gain	Merchandise Displayers and Window Trimmers	0.305	21	44	+23
Gain	Project Management Specialists	0.550	23	44	+21
Gain	Physical Therapists	0.191	1	19	+18
Gain	Animal Caretakers	0.273	4	20	+16

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

SHERBORN IMPACT

Sherborn has the smallest posting base in the sample, so changes should be interpreted cautiously. Even so, the community remained relatively resilient after 2022. Postings declined from 400 in 2022 to 369 in 2025, a 7.8% drop, while remaining 87.3% above 2019.

Measure	Value	Note
Resident AI beta	0.461	Rank 3 of 13
2025 postings	369	Rank 13 by size
2019-2025 change	+87.3%	172
2022-2025 change	-7.8%	-31
Peak year	2021	485 postings

Sherborn's resident AI exposure score is 0.461, ranking 3rd of 13 communities. Its posting-weighted demand exposure was 0.331 in 2025, up 0.032 from 2019. Because the posting base is small, this demand-side exposure measure can move meaningfully when only a few occupations change.

Year	Total postings	Index (2019=100)	Demand AI exposure
2019	197	100.0	0.299
2020	410	208.1	0.332
2021	485	246.2	0.223
2022	400	203.0	0.422
2023	333	169.0	0.381
2024	285	144.7	0.341
2025	369	187.3	0.331

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

The occupation table does not point to a single dominant source of weakness. Declines were scattered across software development, retail sales, statistics, material moving, and security. The gain side is more coherent: registered nurses and several therapy or health-technologist roles increased. Sherborn's local demand therefore appears increasingly care-oriented, even as some small-scale professional, retail, and support categories softened.

Signal	Position	AI beta	2022 Job Listings	2025 Job Listings	Change in Job Listings
Decline	Software Developers	0.868	16	3	-13
Decline	Retail Salespersons	0.342	15	3	-12
Decline	Statisticians	0.725	10	0	-10
Decline	Laborers and Freight, Stock, and Material Movers, Hand	0.042	13	3	-10
Decline	Security Guards	0.167	8	0	-8
Gain	Registered Nurses	0.380	5	73	+68
Gain	Physical Therapists	0.191	0	8	+8
Gain	Cardiovascular Technologists and Technicians	0.357	1	9	+8
Gain	Recreational Therapists	0.318	1	8	+7
Gain	Occupational Therapy Assistants	0.267	0	7	+7

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

SOUTHBOROUGH IMPACT

Southborough also held close to its 2022 level. Postings declined only 3.9%, from 2,715 in 2022 to 2,610 in 2025, and remained 24.2% above 2019. The 2025 level was only 5.1% below the community's 2021 peak.

Measure	Value	Note
Resident AI beta	0.446	Rank 5 of 13
2025 postings	2,610	Rank 7 by size
2019-2025 change	+24.2%	509
2022-2025 change	-3.9%	-105
Peak year	2021	2,751 postings

The resident AI exposure score is 0.446, ranking 5th of 13 communities. On the demand side, posting-weighted AI exposure fell from 0.469 in 2019 to 0.379 in 2025, a sizable decline. This means that Southborough's overall posting level was relatively stable, but the exposed share of its hiring mix moved lower.

Year	Total postings	Index (2019=100)	Demand AI exposure
2019	2,101	100.0	0.469
2020	2,307	109.8	0.434
2021	2,751	130.9	0.364
2022	2,715	129.2	0.425
2023	2,128	101.3	0.425
2024	2,122	101.0	0.389
2025	2,610	124.2	0.379

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

The table shows the churn underneath that stability. Declines in medical secretaries and administrative assistants, material movers, truck drivers, software developers, and pharmacy technicians indicate weakness across administrative, logistics, digital, and selected healthcare-support roles. Gains in registered nursing, physical therapy, radiology, special education, and alarm-system installation point to stronger demand for clinical, education, safety, and hands-on service work.

Signal	Position	AI beta	2022 Job Listings	2025 Job Listings	Change in Job Listings
Decline	Medical Secretaries and Administrative Assistants	0.667	49	7	-42
Decline	Laborers and Freight, Stock, and Material Movers, Hand	0.042	62	26	-36
Decline	Heavy and Tractor-Trailer Truck Drivers	0.279	58	30	-28
Decline	Software Developers	0.868	41	15	-26
Decline	Pharmacy Technicians	0.426	52	28	-24
Gain	Registered Nurses	0.380	99	355	+256
Gain	Physical Therapists	0.191	7	62	+55
Gain	Radiologic Technologists and Technicians	0.377	2	53	+51
Gain	Special Education Teachers, All Other	0.475	47	85	+38
Gain	Security and Fire Alarm Systems Installers	0.167	1	25	+24

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

SUDBURY IMPACT

Sudbury has the highest resident AI exposure score in the sample, but its local posting pattern is more service-oriented than that ranking might suggest. Postings fell from 2,611 in 2022 to 2,294 in 2025, a 12.1% decline, while remaining 34.3% above 2019.

Measure	Value	Note
Resident AI beta	0.465	Rank 1 of 13
2025 postings	2,294	Rank 8 by size
2019-2025 change	+34.3%	586
2022-2025 change	-12.1%	-317
Peak year	2021	2,762 postings

Sudbury's resident AI exposure score is 0.465, ranking 1st of 13. Yet its posting-weighted demand exposure was 0.348 in 2025, essentially unchanged from 2019. This contrast highlights the difference between resident exposure and local employer demand.

Year	Total postings	Index (2019=100)	Demand AI exposure
2019	1,708	100.0	0.356
2020	1,954	114.4	0.348
2021	2,762	161.7	0.313
2022	2,611	152.9	0.374
2023	2,011	117.7	0.373
2024	1,929	112.9	0.348
2025	2,294	134.3	0.348

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

Sudbury is therefore a useful reminder that high resident AI exposure does not necessarily mean local postings are concentrated in high-exposure occupations. Declines were mainly in home health and personal care, material moving, coaching, tutoring, and teaching assistant roles. Gains came from physical therapists, occupational therapists, licensed practical nurses, merchandising, and tax preparers. The local demand profile looks more like a community-service labor market with steady healthcare and therapy demand than a software- or corporate-office hiring base.

Signal	Position	AI beta	2022 Job Listings	2025 Job Listings	Change in Job Listings
Decline	Home Health and Personal Care Aides	0.150	66	28	-38
Decline	Laborers and Freight, Stock, and Material Movers, Hand	0.042	47	11	-36
Decline	Coaches and Scouts	0.300	52	25	-27
Decline	Tutors	0.543	48	22	-26
Decline	Teaching Assistants, Except Postsecondary	0.367	58	32	-26
Gain	Physical Therapists	0.191	16	117	+101
Gain	Occupational Therapists	0.289	3	49	+46
Gain	Licensed Practical and Licensed Vocational Nurses	0.250	27	61	+34
Gain	Merchandise Displayers and Window Trimmers	0.305	19	53	+34
Gain	Tax Preparers	0.625	0	28	+28

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

WAYLAND IMPACT

Wayland cooled sharply after its 2021 peak, but the 2025 level remained slightly above 2019. Postings fell from 1,767 in 2022 to 1,361 in 2025, a 23.0% decline. Compared with the 2021 peak of 1,898, postings were down 28.3%.

Measure	Value	Note
Resident AI beta	0.455	Rank 4 of 13
2025 postings	1,361	Rank 11 by size
2019-2025 change	+4.9%	64
2022-2025 change	-23.0%	-406
Peak year	2021	1,898 postings

Wayland's resident AI exposure score is 0.455, ranking 4th of 13 communities. Posting-weighted demand exposure was 0.329 in 2025, down 0.073 from 2019, indicating a moderate shift away from higher-exposure listed roles.

Year	Total postings	Index (2019=100)	Demand AI exposure
2019	1,297	100.0	0.404
2020	1,317	101.5	0.345
2021	1,898	146.3	0.300
2022	1,767	136.2	0.350
2023	1,292	99.6	0.334
2024	1,243	95.8	0.356
2025	1,361	104.9	0.329

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

The occupation table shows broad, small-scale cooling rather than a single sector shock. Software developers declined, but so did home health and personal care aides, food preparation workers, receptionists and information clerks, and dishwashers. The gains in trucking, coaching, medical assisting, merchandising, and counseling suggest that demand shifted toward transportation, health support, recreation, and other community-facing services.

Signal	Position	AI beta	2022 Job Listings	2025 Job Listings	Change in Job Listings
Decline	Software Developers	0.868	32	3	-29
Decline	Home Health and Personal Care Aides	0.150	56	27	-29
Decline	Food Preparation Workers	0.037	32	10	-22
Decline	Receptionists and Information Clerks	0.533	34	15	-19
Decline	Dishwashers	0.000	25	6	-19
Gain	Heavy and Tractor-Trailer Truck Drivers	0.279	8	37	+29
Gain	Coaches and Scouts	0.300	4	17	+13
Gain	Medical Assistants	0.359	14	25	+11
Gain	Merchandise Displayers and Window Trimmers	0.305	9	19	+10
Gain	Counselors, All Other	0.275	3	13	+10

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

WESTBOROUGH IMPACT

Westborough remains one of Greater MetroWest's major job centers, but it experienced a substantial post-2022 contraction. Postings fell from 9,408 in 2022 to 6,810 in 2025, a 27.6% decline. Compared with the 2021 peak of 9,970, the 2025 level was down 31.7%.

Measure	Value	Note
Resident AI beta	0.437	Rank 7 of 13
2025 postings	6,810	Rank 4 by size
2019-2025 change	-11.9%	-916
2022-2025 change	-27.6%	-2,598
Peak year	2021	9,970 postings

The resident AI exposure score is 0.437, ranking 7th of 13 communities. Posting-weighted demand exposure was 0.409 in 2025, down 0.063 from 2019. This indicates a moderate shift away from higher-exposure listed positions, though Westborough's employer base still includes many professional, technical, and operational roles.

Year	Total postings	Index (2019=100)	Demand AI exposure
2019	7,726	100.0	0.472
2020	7,822	101.2	0.457
2021	9,970	129.0	0.434
2022	9,408	121.8	0.451
2023	7,577	98.1	0.416
2024	7,287	94.3	0.426
2025	6,810	88.1	0.409

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

Westborough's occupation table highlights a sharp retreat in office-based and digital hiring alongside continued demand in technical and operational areas. Software developers, accountants and auditors, marketing managers, customer service representatives, and computer occupations all declined. Gains in physical therapy, engineering, production, machine assembly, and management suggest that healthcare, engineering, manufacturing support, and operations remained important sources of resilience.

Signal	Position	AI beta	2022 Job Listings	2025 Job Listings	Change in Job Listings
Decline	Software Developers	0.868	320	123	-197
Decline	Accountants and Auditors	0.560	171	63	-108
Decline	Marketing Managers	0.500	154	53	-101
Decline	Customer Service Representatives	0.568	214	114	-100
Decline	Computer Occupations, All Other	0.736	164	66	-98
Gain	Physical Therapists	0.191	29	122	+93
Gain	Engineers, All Other	0.523	80	158	+78
Gain	Production Workers, All Other	0.055	146	184	+38
Gain	Engine and Other Machine Assemblers	0.095	20	58	+38
Gain	Managers, All Other	0.491	93	129	+36

Source: MERC calculations from Lightcast job-posting data and the Greater MetroWest community-level AI exposure index.

WORKFORCE AND PLANNING IMPLICATIONS

The findings point to a workforce transition issue rather than a narrow displacement story. AI exposure does not mean that jobs will automatically disappear, but the literature suggests that AI can reshape tasks, change skill requirements, complement some workers, and reduce demand for some routine or entry level tasks. For Greater MetroWest, the practical value of the AI Index is to provide a neutral evidence base for monitoring where occupational exposure, job posting shifts, and workforce adjustment pressure may be emerging.

Several planning implications follow from the data and the broader AI research.

Develop applied AI and digital skills for exposed office and professional workers. Administrative support, customer service, marketing, finance, business operations, and related professional roles are among the areas where AI may alter daily tasks. Useful training would be practical rather than abstract, including AI literacy, data literacy, Excel, Power BI or Tableau, basic Python, and effective use of frontier AI tools such as ChatGPT, Gemini, and Claude. The focus is not only prompt engineering, but also task definition, context building, output verification, confidentiality protection, and responsible use of AI inside daily workflows.

Monitor entry level pathways in AI exposed occupations. Early career roles may face greater pressure because many entry level tasks involve writing, research, documentation, customer communication, data cleaning, reporting, and basic analysis. Local colleges, employers, and workforce partners may use this evidence to examine whether internships, first jobs, student projects, and applied coursework are preparing new workers to work effectively with AI rather than compete directly against it.

Treat healthcare and education growth as areas for AI augmentation. The job posting data show stronger momentum in healthcare, therapy, nursing, education, and related service roles. These sectors should not be viewed as outside the AI transition. They are likely to require AI supported documentation, scheduling, patient communication, instructional support, analytics, compliance, and administrative workflow tools. The relevant workforce question is not only whether these sectors are growing, but whether workers and employers are prepared to use AI in ways that improve productivity, service quality, and worker capacity.

Use community specific employer engagement to interpret the data. Large job centers such as Framingham, Marlborough, Natick, and Westborough may benefit from better information on employer adoption, changing skill requirements, and future hiring plans. Smaller and more residential communities may require a different lens, focused on resident skill resilience, commuting access, remote work readiness, and career transition options. Employer conversations would help clarify whether observed posting changes reflect AI adoption, sector cycles, remote hiring, business relocation, or other local factors.

Continue the Greater MetroWest AI Index as a regional labor market indicator. For MERC and Framingham State University, the index can support ongoing regional learning by tracking occupational exposure, job postings, skill demand, and community level workforce risks over time. Possible complementary activities include employer roundtables, student applied AI projects, collaboration with MassHire and local chambers, AI skill workshops, and an annual Greater MetroWest AI labor market dashboard. Framed this way, the index does not prescribe a specific policy response. It provides information that local stakeholders can use to understand how the labor market is changing.

MERC's recent local employer survey results provide additional context for these findings. Most surveyed employers reported permitting AI use for work-related tasks, particularly for writing, communication, research, and administrative functions. At the same time, relatively few reported plans to replace workers with AI or automated processes (MERC, 2026). Taken together, these responses suggest that AI adoption in Greater MetroWest remains in a relatively early stage, with employers focused more on productivity enhancement and workflow support than direct labor substitution. Whether that pattern persists will depend in part on the pace of adoption and the extent to which AI becomes embedded in everyday business operations.

Timing will shape how disruptive the transition feels. If AI adoption unfolds gradually, workers, employers, colleges, and training providers will have more time to adjust. If adoption accelerates quickly, some workers and career pathways may come under pressure before local institutions can respond. Regular updates to the Greater MetroWest AI Index would help the region monitor those changes and prepare evidence based responses.

LIMITATIONS AND ADDITIONAL DATA NEEDS

This report is an early descriptive analysis. It is useful for identifying patterns and planning questions, but it does not prove causality. Several limitations are important.

Data source or method	Use in this report	Important limitation
All-occupation Lightcast postings	Measures annual unique postings by community and occupation.	Still measures postings, not actual employment, hires, wages, layoffs, or job quality.
Resident AI exposure	Measures occupations held by residents using ACS occupation structure.	Does not measure local employer adoption or whether residents work inside the town.
Demand AI exposure	Measures AI exposure of occupations employers are posting locally.	Depends on SOC-code matching and excludes unmatched or non-substantive categories.
Correlation analysis	Compares resident exposure with 2022–2025 posting change.	Descriptive only; does not control for interest rates, sector mix, population, commuting, or employer decisions.
Annual data	Shows the broad post-pandemic cycle.	Cannot precisely identify the timing of the ChatGPT shock or monthly changes in employer behavior.

The next version of the report would be stronger with posting-level Lightcast data including employer name, NAICS industry, detailed SOC code, education and experience requirements, wages, skills, remote-work status, and posting dates. ACS/LEHD commuting data and OEWS employment and wage data would help connect resident exposure with workplace demand. A local employer AI adoption survey would directly measure whether AI is being used to substitute for labor, augment workers, or change skill requirements.



SUMMARY

The Greater MetroWest AI Index shows that AI exposure is uneven across communities and occupations, but the local job-posting evidence does not support a simple story that higher-exposure communities are losing the most job demand. Instead, the region is experiencing a broader post-2022 job-posting normalization combined with a meaningful shift in occupational mix.

The strongest aggregate signal is the decline in posting-weighted AI exposure of local hiring demand from 2019 to 2025. Greater MetroWest's recent hiring mix is less concentrated in AI-exposed digital, analytical, customer-service, finance, marketing, and administrative roles and more concentrated in healthcare, education, therapy, and local services. For workforce planning, this is the key insight: AI risk is not evenly distributed, and the response should be targeted adaptation rather than generalized alarm.

The practical recommendation is to move from awareness to implementation. Greater MetroWest needs an AI-ready workforce strategy that strengthens applied AI literacy, supports workers in exposed clerical and analytical roles, connects healthcare and education growth with AI augmentation, and links employers with students, faculty, MassHire, chambers of commerce, and community partners.



REFERENCES

- Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2020). AI and Jobs: Evidence from Online Vacancies. NBER Working Paper No. 28257.
- Appel, R., McCrory, P., Tamkin, A., McCain, M., Neylon, T., and Stern, M. (2025). The Anthropic Economic Index Report: Uneven Geographic and Enterprise AI Adoption. Anthropic.
- Atkinson, T., and Yamco, S. (2026). Young Workers' Employment Drops in Occupations with High AI Exposure. Federal Reserve Bank of Dallas.
- Cazzaniga, M., Jaumotte, F., Li, L., Melina, G., Panton, A. J., Pizzinelli, C., Rockall, E., and Tavares, M. M. (2024). Gen-AI: Artificial Intelligence and the Future of Work. International Monetary Fund Staff Discussion Note.
- Council of Economic Advisers. (2024). Potential Labor Market Impacts of Artificial Intelligence: An Empirical Analysis. Executive Office of the President.
- Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). "GPTs are GPTs: Labor Market Impact Potential of LLMs." *Science*, 384, 1306.
- Felten, E., Raj, M., and Seamans, R. (2021). "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses." *Strategic Management Journal*, 42, 2195–2217.
- Gimbel, M., Kendall, J., and Kulsakdinun, R. (2026). Labor Market AI Exposure: What Do We Know? The Budget Lab at Yale.
- Gmyrek, P., Berg, J., Kamiński, K., Konopczyński, F., Ładna, A., Nafradi, B., Rosłaniec, K., and Troszyński, M. (2025). Generative AI and Jobs: A Refined Global Index of Occupational Exposure. ILO Working Paper 140. International Labour Organization.
- Huang, Y. (2024). The Labor Market Impact of Artificial Intelligence: Evidence from US Regions. IMF Working Paper WP/24/199.
- Kneebone, E., and Holmes, N. (2025). On-the-Job Exposure to AI Among Lower-Income Workers. Federal Reserve Bank of San Francisco Community Development Research Brief.
- Kochhar, R. (2023). Which U.S. Workers Are More Exposed to AI on Their Jobs? Pew Research Center.
- Liu, Y., Wang, H., and Yu, S. (2025). Labor Demand in the Age of Generative AI: Early Evidence from the U.S. Job Posting Data. World Bank Policy Research Working Paper 11263.
- Manning, S. J., and Aguirre, T. (2026). How Adaptable Are American Workers to AI-Induced Job Displacement? NBER Working Paper No. 34705.

Maslej, N., Fattorini, L., Perrault, R., Gil, Y., Parli, V., Kariuki, N., Capstick, E., Reuel, A., Brynjolfsson, E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Niebles, J. C., Shoham, Y., Wald, R., Walsh, T., Hamrah, A., Santarlasci, L., Betts Lotufo, J., Rome, A., Shi, A., and Oak, S. (2025). Artificial Intelligence Index Report 2025. Stanford Institute for Human-Centered Artificial Intelligence.

MERC (MetroWest Economic Research Center). (2026). 495/MetroWest Employers Survey, June 2026. MERC Insights Series. Forthcoming.

Muro, M., Methkupally, S., and Kinder, M. (2025). The Geography of Generative AI's Workforce Impacts Will Likely Differ from Those of Previous Technologies. Brookings Institution.

Scavette, A., Jenkins, D., and Ding, L. (2025). Occupational Exposure to Generative Artificial Intelligence in the Third Federal Reserve District. Federal Reserve Bank of Philadelphia.

Schendstok, M., and Schreiner Wertz, S. (2024). Occupational Exposure to Artificial Intelligence by Geography and Education. U.S. Department of the Treasury, Office of Economic Policy Working Paper 2024-02.

Appendix: Python Code for Figure Generation

The following Python code generates the report figures using the all-occupation Lightcast Q2 2026 job-posting workbook and the town-level AI exposure file.

```
import pandas as pd
import numpy as np
import matplotlib.pyplot as plt
from matplotlib.ticker import FuncFormatter
from pathlib import Path
import re

# =====
# 1. File paths and settings
# =====
# This version is designed for the all-occupation Lightcast Q2 2026 export.
# The workbook contains one annual sheet per community, such as "Ashland Annual".
# Each annual sheet includes a row named "Total Across All Occupations" and
# separate occupation rows underneath it.

JOB_FILE = "Community Occupations--All Occupations Postings Table-2016-2025.xlsx"
AI_FILE = "Town_level_Data_AI_Index_UPDATED.xlsx"

OUTDIR = Path("metrowest_ai_report_figures")
OUTDIR.mkdir(exist_ok=True)

YEARS = list(range(2019, 2026))
TOTAL_ROW_LABEL = "Total Across All Occupations"

# =====
# 2. Helper functions
# =====
def clean_town(x):
    """Standardize town/community names."""
    if pd.isna(x):
        return np.nan
    x = str(x).replace(", MA", "").strip()
    return {"Marlborough": "Marlborough"}.get(x, x)

def normalize_title(x):
    """Normalize occupation titles for title-based checks if needed."""
    if pd.isna(x):
        return np.nan
    x = str(x).lower().strip()
    x = x.replace("&", "and")
    x = re.sub(r"^[a-z0-9\s]", "", x)
    x = re.sub(r"^\s+", "", x)
    return x
```

```

def soc5(x):
    """Return the 6-digit SOC form used in the Lightcast export, e.g., 15-1252."""
    if pd.isna(x):
        return np.nan
    m = re.search(r"\d{2}-\d{4}", str(x))
    return m.group(0) if m else np.nan

def weighted_average(group, value_col, weight_col="Postings"):
    """Posting-weighted average that safely handles empty or zero-weight groups."""
    g = group.dropna(subset=[value_col]).copy()
    total_weight = g[weight_col].sum()
    if total_weight == 0:
        return np.nan
    return np.average(g[value_col], weights=g[weight_col])

def add_bar_labels(ax, values, fmt="{:,.}", negative=False):
    """Add readable bar labels with padding so extreme labels do not overlap axis text."""
    xmin, xmax = ax.get_xlim()
    span = xmax - xmin
    pad = span * 0.01
    for i, value in enumerate(values):
        label = fmt.format(int(value))
        if negative:
            ax.text(value + pad, i, label, va="center", ha="left", fontsize=8)
        else:
            ax.text(value + pad, i, label, va="center", ha="left", fontsize=8)

# =====
# 3. Load all-occupation Lightcast annual job-posting data
# =====
xl = pd.ExcelFile(JOB_FILE)
annual_sheets = [s for s in xl.sheet_names if s.endswith(" Annual")]
towns = [clean_town(s.replace(" Annual", "")) for s in annual_sheets]

posting_rows = []
town_total_rows = []

for sheet_name in annual_sheets:
    town = clean_town(sheet_name.replace(" Annual", ""))
    df = pd.read_excel(JOB_FILE, sheet_name=sheet_name)
    df["Town"] = town
    df["Occupation"] = df["Occupation"].astype(str).str.strip()
    df["SOC5"] = df["SOC"].apply(soc5)

    # Use the explicit total row for town-level totals.
    # Do not sum all occupation rows for town totals, because the export includes
    # a total row plus occupation rows; summing both would double-count totals.
    total_row = df[df["Occupation"].eq(TOTAL_ROW_LABEL)]
    if total_row.empty:
        raise ValueError(f"Missing total row in {sheet_name}")

```

```

for year in YEARS:
    year_col = f"{year} Total Unique Postings"
    town_total_rows.append({
        "Town": town,
        "Year": year,
        "Total_Postings": pd.to_numeric(total_row[year_col].iloc[0], errors="coerce")
    })

# All occupation rows, excluding the explicit total row.
occupation_rows = df[~df["Occupation"].eq(TOTAL_ROW_LABEL)].copy()
for year in YEARS:
    year_col = f"{year} Total Unique Postings"
    temp = occupation_rows[["Town", "SOC", "SOC5", "Occupation", year_col]].copy()
    temp = temp.rename(columns={year_col: "Postings"})
    temp["Year"] = year
    temp["Postings"] = pd.to_numeric(temp["Postings"], errors="coerce").fillna(0)
    posting_rows.append(temp)

town_totals = pd.DataFrame(town_total_rows)
position_rows = pd.concat(posting_rows, ignore_index=True)
position_rows = position_rows[position_rows["Postings"].notna()].copy()

# =====
# 4. Load AI exposure data and create SOC-level lookup
# =====
ai = pd.read_excel(AI_FILE, sheet_name="occ_level")

# Detailed occupation-level AI beta scores.
occ_ai = ai[["O*NET-SOC Code", "Title", "dv_rating_beta"]].dropna(subset=["O*NET-SOC Code", "dv_rating_beta"]).copy()
occ_ai["SOC5"] = occ_ai["O*NET-SOC Code"].apply(soc5)
occ_ai["dv_rating_beta"] = pd.to_numeric(occ_ai["dv_rating_beta"], errors="coerce")

# Lightcast uses six-digit SOC codes. The AI source uses detailed O*NET-SOC codes.
# We average detailed O*NET-SOC scores to the six-digit SOC level for matching.
ai_by_soc = (
    occ_ai
    .dropna(subset=["SOC5", "dv_rating_beta"])
    .groupby("SOC5", as_index=False)["dv_rating_beta"]
    .mean()
)

# Community-level resident AI exposure from the existing town-level AI index file.
town_ai_records = []
for col in ai.columns:
    if str(col).startswith("dv_rating_beta - "):
        town = clean_town(str(col).replace("dv_rating_beta - ", "").strip())
        value = pd.to_numeric(ai[col], errors="coerce").sum()
        town_ai_records.append({
            "Town": town,
            "AI_Exposure_Beta": value
        })

```

```

})
town_ai = pd.DataFrame(town_ai_records)

# =====
# Figure 1: Greater MetroWest AI exposure by major occupation group
# =====
spade_df = pd.read_excel(SPADE_FILE, sheet_name="Sheet1")
spade_df = spade_df.dropna(subset=["title", "dv_rating_beta"]).copy()

label_map = {
"Management occupations": "Management",
"Business and financial operations occupations": "Business & finance",
"Computer and mathematical occupations": "Computer & math",
"Architecture and engineering occupations": "Architecture & engineering",
"Life, physical, and social science occupations": "Life & social sciences",
"Community and social service occupations": "Social services",
"Legal occupations": "Legal",
"Educational instruction, and library occupations": "Education & library",
"Arts, design, entertainment, sports, and media occupations": "Arts & media",
"Health diagnosing and treating practitioners and other technical occupations": "Healthcare practitioners",
"Health technologists and technicians": "Health technologists",
"Healthcare support occupations": "Healthcare support",
"Firefighting and prevention, and other protective service workers including supervisors": "Protective service",
"Law enforcement workers including supervisors": "Law enforcement",
"Food preparation and serving related occupations": "Food & serving",
"Building and grounds cleaning and maintenance occupations": "Grounds maintenance",
"Personal care and service occupations": "Personal care",
"Sales and related occupations": "Sales",
"Office and administrative support occupations": "Office & admin",
"Farming, fishing, and forestry occupations": "Agriculture",
"Construction and extraction occupations": "Construction",
"Installation, maintenance, and repair occupations": "Installation & repair",
"Production occupations": "Production",
"Transportation occupations": "Transportation",
"Material moving occupations": "Material moving",
}

spade_df["label"] = spade_df["title"].map(label_map).fillna(spade_df["title"])

# Remove only the visible text label for Protective service.
# The data point remains in the radar chart.
spade_df.loc[spade_df["label"] == "Protective service", "label"] = ""

labels = spade_df["label"].tolist()
values = spade_df["dv_rating_beta"].astype(float).tolist()

angles = np.linspace(0, 2 * np.pi, len(labels), endpoint=False).tolist()
angles_closed = angles + angles[:1]
values_closed = values + values[:1]

```

```

fig = plt.figure(figsize=(12, 10))
ax = plt.subplot(111, polar=True)

ax.set_theta_offset(np.pi / 2)
ax.set_theta_direction(-1)

ax.plot(
    angles_closed,
    values_closed,
    linewidth=2,
    marker="o",
    markersize=4,
    label="AI exposure beta"
)

ax.fill(
    angles_closed,
    values_closed,
    alpha=0.20
)

ax.set_xticks(angles)
ax.set_xticklabels(labels, fontsize=10)

ax.set_ylim(0, 1.0)
ax.set_yticks([0.2, 0.4, 0.6, 0.8, 1.0])
ax.set_yticklabels(["0.2", "0.4", "0.6", "0.8", "1.0"], fontsize=9)
ax.set_rlabel_position(225)

ax.set_title(
    "Greater MetroWest AI Exposure by Major Occupation Group\nDV Rating Beta Measure",
    fontsize=18,
    fontweight="bold",
    pad=32
)

ax.legend(
    loc="lower right",
    bbox_to_anchor=(1.25, -0.05),
    frameon=True
)

fig.text(
    0.5,
    0.02,
    "Source: MERC calculations using O*NET occupational classifications, "
    "2023 American Community Survey employment data, and DV beta AI exposure scores from Eloundou et al. (2024).",
    ha="center",
    fontsize=9
)

```

```
plt.tight_layout(rect=[0, 0.05, 1, 1])
plt.savefig(OUTDIR / "metrowest_ai_exposure_radar_chart_dv_beta.png", dpi=300, bbox_inches="tight")
plt.close()

# =====
# Figure 2: Aggregate Greater MetroWest job postings, 2019–2025
# =====
aggregate_trend = town_totals.groupby("Year", as_index=False)["Total_Postings"].sum()

fig, ax = plt.subplots(figsize=(9, 5), constrained_layout=True)
ax.plot(aggregate_trend["Year"], aggregate_trend["Total_Postings"], marker="o", linewidth=2)
for _, row in aggregate_trend.iterrows():
    ax.annotate(f"{int(row['Total_Postings']):,}", xy=(row["Year"], row["Total_Postings"]), xytext=(0, 8), textcoords="offset points",
    ha="center", fontsize=8)
ax.set_title("Figure 2: Aggregate Greater MetroWest Job Postings, 2019–2025")
ax.set_xlabel("Year")
ax.set_ylabel("Total unique job postings")
ax.grid(True, alpha=0.3)
plt.savefig(OUTDIR / "aggregate_metrowest_job_postings.png", dpi=300, bbox_inches="tight")
plt.close()

# =====
# Figure 3: Largest all-occupation posting declines, 2022–2025
# =====
position_change = (
    position_rows
    .groupby(["Occupation", "Year", as_index=False])["Postings"]
    .sum()
    .pivot_table(index="Occupation", columns="Year", values="Postings", fill_value=0)
)

for year in YEARS:
    if year not in position_change.columns:
        position_change[year] = 0

position_change["Change_2022_2025"] = position_change[2025] - position_change[2022]
top_declines = position_change.sort_values("Change_2022_2025").head(12)

fig, ax = plt.subplots(figsize=(10.5, 6.5), constrained_layout=True)
bars = ax.barh(top_declines.index, top_declines["Change_2022_2025"])
min_value = top_declines["Change_2022_2025"].min()
ax.set_xlim(min_value * 1.15, 0)
add_bar_labels(ax, top_declines["Change_2022_2025"].values, negative=True)
ax.set_title("Figure 3: Largest All-Occupation Posting Declines, 2022–2025")
ax.set_xlabel("Change in unique job postings")
ax.set_ylabel("Occupation")
ax.grid(axis="x", alpha=0.3)
plt.savefig(OUTDIR / "largest_all_occupation_declines_2022_2025.png", dpi=300, bbox_inches="tight")
```

```
plt.close()

# =====
# Figure 4: Largest all-occupation posting gains, 2022–2025
# =====
top_gains = position_change.sort_values("Change_2022_2025", ascending=False).head(12)

fig, ax = plt.subplots(figsize=(10.5, 6.5), constrained_layout=True)
bars = ax.barh(top_gains.index, top_gains["Change_2022_2025"])
max_value = top_gains["Change_2022_2025"].max()
ax.set_xlim(0, max_value * 1.15)
add_bar_labels(ax, top_gains["Change_2022_2025"].values, negative=False)
ax.set_title("Figure 4: Largest All-Occupation Posting Gains, 2022–2025")
ax.set_xlabel("Change in unique job postings")
ax.set_ylabel("Occupation")
ax.grid(axis="x", alpha=0.3)
plt.savefig(OUTDIR / "largest_all_occupation_gains_2022_2025.png", dpi=300, bbox_inches="tight")
plt.close()

# =====
# Figure 5: Resident AI exposure index by community
# =====
ai_rank = town_ai.sort_values("AI_Exposure_Beta")

fig, ax = plt.subplots(figsize=(9, 6), constrained_layout=True)
bars = ax.barh(ai_rank["Town"], ai_rank["AI_Exposure_Beta"])
for i, value in enumerate(ai_rank["AI_Exposure_Beta"]):
    ax.text(value + 0.003, i, f"{value:.3f}", va="center", ha="left", fontsize=8)
ax.set_xlim(0, ai_rank["AI_Exposure_Beta"].max() * 1.12)
ax.set_title("Figure 5: Resident AI Exposure Index by Community")
ax.set_xlabel("Resident AI exposure index, beta measure")
ax.set_ylabel("Community")
ax.grid(axis="x", alpha=0.3)
plt.savefig(OUTDIR / "resident_ai_exposure_index_by_community.png", dpi=300, bbox_inches="tight")
plt.close()

# =====
# Figure 6: Resident AI exposure vs. post-2022 posting change
# =====
town_trend = town_totals.pivot_table(index="Town", columns="Year", values="Total_Postings", aggfunc="sum").reset_index()
town_trend.columns.name = None
town_trend["Pct_Change_2022_2025"] = (town_trend[2025] - town_trend[2022]) / town_trend[2022]
town_trend = town_trend.merge(town_ai, on="Town", how="left")
scatter_df = town_trend.dropna(subset=["AI_Exposure_Beta", "Pct_Change_2022_2025"]).copy()
```

```

x = scatter_df["AI_Exposure_Beta"]
y = scatter_df["Pct_Change_2022_2025"]
corr = x.corr(y)

fig, ax = plt.subplots(figsize=(11.5, 6.5), constrained_layout=True)
ax.scatter(x, y, s=120, alpha=0.9)

if len(scatter_df) >= 2:
    slope, intercept = np.polyfit(x, y, 1)
    x_line = np.linspace(x.min(), x.max(), 100)
    ax.plot(x_line, slope * x_line + intercept, linewidth=2)

ax.axhline(0, linewidth=1)

offsets = {
    "Ashland": (12, 4),
    "Framingham": (12, 4),
    "Holliston": (12, -8),
    "Hopkinton": (12, -8),
    "Hudson": (12, 4),
    "Marlborough": (12, 4),
    "Natick": (12, -8),
    "Northborough": (12, 12),
    "Sherborn": (12, 8),
    "Southborough": (12, -8),
    "Sudbury": (12, 4),
    "Wayland": (12, 4),
    "Westborough": (12, 4),
}

for _, row in scatter_df.iterrows():
    town = row["Town"]
    dx, dy = offsets.get(town, (8, 5))
    arrowprops = dict(arrowstyle="-", linewidth=0.8) if town in ["Northborough", "Southborough"] else None
    ax.annotate(
        town,
        xy=(row["AI_Exposure_Beta"], row["Pct_Change_2022_2025"]),
        xytext=(dx, dy),
        textcoords="offset points",
        fontsize=10,
        ha="left",
        va="center",
        bbox=dict(boxstyle="round,pad=0.15", fc="white", ec="none", alpha=0.75),
        arrowprops=arrowprops,
    )

ax.set_title(f"Figure 6: Resident AI Exposure vs. Job-Posting Change, 2022–2025\nCorrelation = {corr:.2f}")
ax.set_xlabel("Resident AI exposure index, beta measure")

```

```

ax.set_ylabel("Job-posting change, 2022–2025")
ax.yaxis.set_major_formatter(FuncFormatter(lambda val, pos: f"{val:.0%}"))
ax.grid(True, alpha=0.3)
plt.savefig(OUTDIR / "resident_ai_exposure_vs_posting_change.png", dpi=300, bbox_inches="tight")
plt.close()

# =====
# Figure 7: Posting-weighted AI exposure of local hiring demand
# =====
# This figure uses only postings that can be matched to a valid six-digit SOC AI beta score.
# Unmatched or non-substantive categories remain in total posting counts but are excluded
# from the posting-weighted AI exposure calculation.
position_ai = position_rows.merge(ai_by_soc, on="SOC5", how="left")
matched = position_ai.dropna(subset=["dv_rating_beta"]).copy()

demand_ai_by_year = pd.DataFrame([
    {
        "Year": year,
        "Posting_Weighted_AI_Exposure": weighted_average(group, "dv_rating_beta")
    }
    for year, group in matched.groupby("Year")
])

coverage_by_year = (
    position_ai.groupby("Year")["Postings"].sum().rename("All_Postings")
    .to_frame()
    .join(matched.groupby("Year")["Postings"].sum().rename("Matched_Postings"))
    .reset_index()
)
coverage_by_year["Coverage"] = coverage_by_year["Matched_Postings"] / coverage_by_year["All_Postings"]

demand_ai_by_year = demand_ai_by_year.merge(coverage_by_year, on="Year", how="left")

fig, ax = plt.subplots(figsize=(9, 5), constrained_layout=True)
ax.plot(demand_ai_by_year["Year"], demand_ai_by_year["Posting_Weighted_AI_Exposure"], marker="o", linewidth=2)
for _, row in demand_ai_by_year.iterrows():
    ax.annotate(f"{row['Posting_Weighted_AI_Exposure']:.3f}", xy=(row["Year"], row["Posting_Weighted_AI_Exposure"]), xytext=(0,
8), textcoords="offset points", ha="center", fontsize=8)
ax.set_title("Figure 7: Posting-Weighted AI Exposure of Local Hiring Demand")
ax.set_xlabel("Year")
ax.set_ylabel("Posting-weighted AI exposure")
ax.grid(True, alpha=0.3)
plt.savefig(OUTDIR / "posting_weighted_ai_exposure_of_local_hiring_demand.png", dpi=300, bbox_inches="tight")
plt.close()

```



**FRAMINGHAM
STATE
UNIVERSITY**